Fourier Transform Toolkit

Rob Dimeo
NIST Center for Neutron Research
robert.dimeo@nist.gov

CONTENTS

INTRODUCTION

Data from inelastic neutron scattering measurements is typically reduced into a quantity that is proportional to the scattering function, $S(Q,\omega)$. The exception to this is neutron spin-echo in which reduced data is proportional to the intermediate scattering function, $I(Q,t)$. The relationship between the two quantities is a Fourier Transform in the ω-t conjugate pair. There are an increasing number of systems being explored using multiple techniques including neutron spin-echo and researchers would like to be able to visualize and perhaps even analyze data in the time domain that was originally reduced to $S(Q,\omega)$. To that end we have developed the Fourier Transform Toolkit for use in the DAVE software system. In short, this toolkit allows users to transform their data from $S(Q,\omega)$ to $I(Q,t)$. Moreover, users can divide two data sets in the time domain in an effort to gain insight into data with the effects of instrumental resolution minimized. However extreme caution must be exercised when using this as a method of deconvolution. This deconvolution method is notorious for amplifying noise in the data.

Before describing the application in detail it is necessary to say a few words about the technique being applied here and the level of knowledge expected of the user of this application. This application uses filtering and the Fast Fourier Transform (FFT) to treat data. There are numerous sources of error associated with performing transformation on data. By using this application, you the user accept complete responsibility for knowing all of the pitfalls and sources of error associated with the FFT. Moreover you are responsible for all results that come from using this application to analyze your data. By reading this documentation you gain an understanding of how to use this application but it must not be done in lieu of having a thorough understanding of the FFT. There are numerous excellent texts on the FFT and you are expected to have working knowledge of this numerical technique and its suitability for your situation.

This application converts $S(Q,\omega)$ to $I(Q,t)$. Currently there is no capability in this application to transform $I(Q,t)$ to $S(Q,\omega)$. Prior to Fourier Transformation the user has a number of processing options available. For instance, the energy range over which the transform will be performed is adjustable. The user can specify an energy-dependent function to be subtracted off of the data, $S(Q,\omega)$, prior to transformation. Finally the data can be filtered with Hanning, Savitzky-Golay, Gaussian, or boxcar filters with adjustable bandwidths. Post transformed data can also be filtered. It is also possible to obtain estimates on the statistical uncertainty in the transformed data. The method used to do this is referred to as "Bootstrap Monte-Carlo" and a discussion of it can be found in Numerical Recipes[1]. Effectively a family of synthetic data is simulated using the Monte-Carlo method based on the original data set and each of the synthetic data sets are processed using an FFT. Statistics on each data point are calculated resulting in an error bar for each. The number of synthetic data sets computed for each real data set can be specified by the user.

[1] W.H.Press, B.P.Flannery, S.A.Teukolsky, and W.T.Vetterling, *Numerical Recipes:The Art of Scientific Computing (FORTRAN Version)*, Cambridge University Press (1989).

The operations on the data in the previous paragraph can be applied to $S(Q,\omega)$ for both a "sample" data set and the "resolution" function data set. One can then perform a deconvolution via division of the data sets and obtain an estimate on the resolution-corrected data.

If you have a question, feel free to contact Rob Dimeo via e-mail (robert.dimeo@nist.gov) or telephone ((301) 975-8135).

I would like to acknowledge the valuable input in the design of this software from Inma Peral, Vicky Garcia-Sakai, Paul Sokol, Zema Chowdhuri, Dan Neumann, Richard Azuah, Peter Peterson, Brian Toby, and John Copley.

OVERVIEW OF THE USER-INTERFACE

The user interface (UI) for this application is composed of three tabs labeled Sample, Resolution, and Deconvolution, as well as a plot window to the right of the tabs (see Fig. 1). You navigate among the various tabs by clicking on the tab itself. The control panels for Sample and Resolution, chosen by selecting one of the tabs, are identical in terms of the controls exposed to the user. The controls for the deconvolution are identical except that the FT Processing Parameters controls are absent, as are the post-filtering options and MC error estimate.

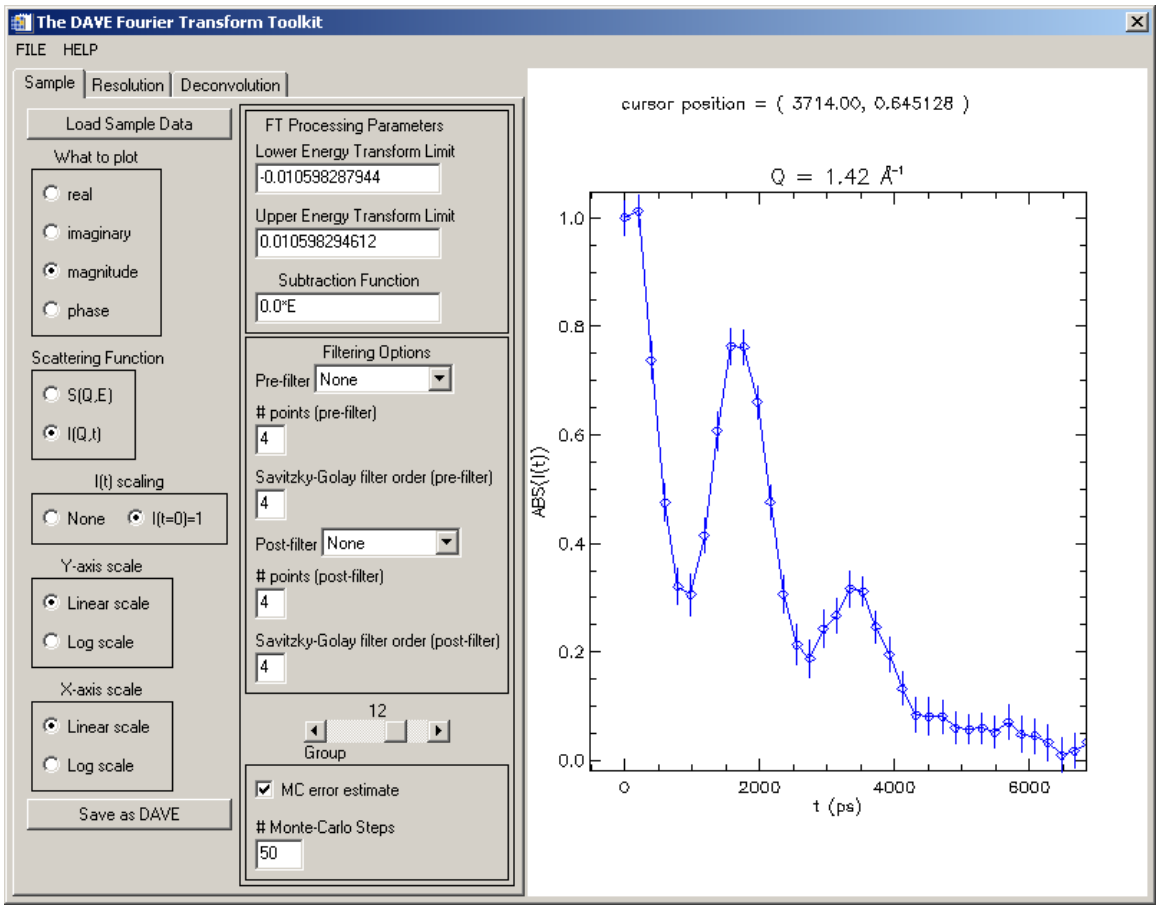

Figure 1 User interface with a "sample" data file loaded.

The Menu Bar

The menu bar at the top of the interface contains two menus: **FILE**, and **HELP**. The menu option under **FILE** is **Quit**. This choice exits the program returning the user to the main DAVE level.

The **HELP** menu contains options labeled **PDF documentation** and **About the Fourier Transform Toolkit**. The **Help** option launches the documentation you are now reading and **About**.... provides some information about who wrote it.

The Control Panel

The remainder of the controls are located on the control panel to the left of the plot window. These are explained in the remainder of this section.

Load Sample Data

This button, when pressed, launches a dialog allowing you to load in a DAVE file. Note that your data must have equal energy bins. It is assumed that the data is in energy units of meV; however backscattering data collected in μeV is automatically converted to meV internally, requiring you to do nothing additional. It is important to recognize that the data loaded becomes the real part of a complex array.

What to plot

This set of four exclusive buttons allows the user to select which component of the transformed data he/she wishes to view in the plot. Options include the real, imaginary, magnitude, and phase of the data. Note that the data for S(Q,E) are assumed to be entirely real with no imaginary component.

Scattering Function

This set of two exclusive buttons allows the user to select which function to plot in the window. The choices are S(Q,E) and I(Q,t).

I(t) scaling

This set of two exclusive buttons allows the user to select whether or not to apply scaling to I(Q,t). If **None** is selected then no scaling is applied. If **I(t=0)=1** is selected then the data are normalized to unity at t=0.

Y-axis scale

This set of two exclusive buttons allows the user to choose between displaying the data on a linear or logarithmic y-axis.

X-axis scale

This set of two exclusive buttons allows the user to choose between displaying the data on a linear or logarithmic x-axis.

Save as DAVE

This button prompts the user to save the time-domain transformed data as the native DAVE format which can be further processed, visualized, and analyzed using the other programs in the DAVE suite. The application will write out four files corresponding to the real part, imaginary part, magnitude, and phase. The default root name for the data file is just the original DAVE filename with an extension: time_mag, time_phase, time_real, time_imaginary. For instance, if the original DAVE file read into the program is 20031201_01.dave then the four corresponding files will be:

20031201_01_time_mag.dave,

20031201_01_time_phase.dave,

20031201_01_time_phase.dave,

20031201_01_time_imaginary.dave.

Alternatively you have the option of changing the root name of the file from 20031201_01 to any other valid file name. In that case the application will append the four extensions discussed above to the new root name of your choice.

Lower/Upper Energy Transform Limit

In these two text fields the user can specify over what energy range to perform the Fourier Transform. The default limits are the limits of original data set.

Subtraction Function

The user can type in an expression in terms of the variable E (for energy) which will be subtracted from S(Q,E) prior to performing the FFT. The syntax for such a function is the standard IDL syntax. For example, if you wanted to subtract off the following cubic function

$3E^3-2E-1$

from S(Q,E) prior to applying the FFT then you would type in the following:

`3*E-2*E-1.0`.

Pre-Filter Filtering Options (*Sample and Resolution tabs only*)

Prior to applying the FFT the user can select a pre-filter from this drop-down menu. Available choices include boxcar average, Gaussian blur, Savitzky-Golay, Hanning, and None. The default is none. These filters all require a width which the user specifies in the control box labeled **# points (pre-filter)**. An additional filter parameter required for the Savitzky-Golay choice is the order. This is specified in the control box labeled **Savitzky-Golay filter order (pre-filter)**. If you are unfamiliar with these filters then you are encouraged to consult a reference such as Numerical Recipes[2] prior to using them.

Post-Filter Filtering Options (*Sample and Resolution tabs only*)

The same filters as described in the previous paragraph can be applied after application of the FFT using this identically-configured set of controls.

Group

For data that consists of multiple groups (Q-values for instance) the user can select which group is displayed in the plot window by changing the position of the slider control.

MC error estimate (*Sample and Resolution tabs only*)

If this checkbox is checked then a Monte-Carlo estimate of the errors propagated through the FFT will be performed. In this method, the data, S(Q,E) and the associated statistical uncertainties on the data are used to synthesize a "large" number of additional data sets. The number of synthetic data sets is specified by the field labeled **# Monte-Carlo Steps**. These synthetic data sets are all propagated through the FFT and the standard deviations of the resulting transformed data are computed and used as the estimates for the propagated error bars. The user has the flexibility to modify the number of synthetic data sets used in the Monte-Carlo calculation via the field labeled **# Monte-Carlo Steps**. Note that this

[2] W.H.Press, B.P.Flannery, S.A.Teukolsky, and W.T.Vetterling, *Numerical Recipes:The Art of Scientific Computing (FORTRAN Version)*, Cambridge University Press (1989).

algorithm only accounts for statistical uncertainty in the data. Other sources of error such as truncation, aliasing, and quantization are not included in the error bars.

 # Monte-Carlo Steps (*Sample and Resolution tabs only*)

User-specified number of synthetic data sets to use in estimating the statistical uncertainty on the transformed data.

The Data Display Window

The data display window contains a plot of the data reflecting which tab is currently selected and which data processing options are currently selected in the control panel. Only one data file can be displayed at a time. The data display window allows zooming using the familiar rubber-band box method. This works by the user holding down the left button while the cursor is in the data display window, dragging the mouse over the region of interest, and releasing the mouse button. The result is a magnified view of the region of interest. In order to autoscale the plot, press the right mouse button while the cursor is anywhere in the data display window. Also the data coordinates are displayed at the top of the data display window as a coordinate pair: (x,y).

The Deconvolution Tab

Deconvolution is performed by dividing the sample I(Q,t) by the resolution I(Q,t). The error bars as determined via Monte-Carlo estimation for the sample and resolution are propagated through the division via the usual rule for error propagation:

$$z = \frac{u}{v}; \quad \sigma_z = z\sqrt{\left(\frac{\sigma_u}{u}\right)^2 + \left(\frac{\sigma_v}{v}\right)^2},$$

where $u=I_{sample}(Q,t)$ and $v=I_{res}(Q,t)$.

You need to be aware that this method of deconvolution has its own difficulties associated with it. In particular it has the tendency to amplify noise in the data.

If you wish to write the result of the deconvolution estimate to four DAVE files then you must ensure that the errors have been calculated. This is accomplished by making sure that the checkboxes labeled **MC error estimate** are checked for both sample and resolution and then you press the deconvolute button in the deconvolution tab.

A SIMPLE EXAMPLE

In order to demonstrate how one can use this application we will work through an example. In this example our goal is to obtain an estimate on the intrinsic intermediate scattering function for a tunneling system. To be specific, we wish to obtain an estimate on I(Q,t) with the effects of resolution removed to the extent possible with the FFT.

The data files for this example are (a) sample: 20010618_01dyn.dave and (b) resolution: 20010529_01dyn.dave. These data files were collected on the High Flux Backscattering spectrometer. The sample file is loaded by pressing the button labeled **Load Sample File** on the Sample tab. Likewise the resolution file is loaded by selecting the resolution tab and then pressing the **Load Resolution File**. For the purposes of this tutorial, you can select **Load example files** from the **File** menu and both example files will be loaded. When the sample file is loaded, the Fourier transform is performed. You can see that this has been computed by pressing the I(Q,t) selection on the **Scattering Function** exclusive button set and the **magnitude** choice in the **What to plot** button set. There are 16 detectors (groups) in this data set, each corresponding to a different wavevector transfer. Note that the dynamic range for the data is ±11 μeV but the program has converted the display units so that they are in terms of meV (±0.011 meV). This was done in an effort to provide consistency in the time units (pico-seconds). Also it is important to note that the original energy-domain data has been placed in the real part of the complex data and the imaginary part of the data is zero. You may examine the real and imaginary parts of the complex I(Q,t) as well its magnitude and phase by choosing the button labeled as such from the **What to plot** button set for the **Scattering function** selected as I(Q,t).

There are a number of different ways to view your data. It can be instructive to view your time-domain data in semi-logarithmic form with the time-axis on a logarithmic scale. You can do this in this application by selecting **Log scale** from the **X-axis scale** button set. Additionally you can choose to display the y-axis on a logarithmic scale by choosing **Log scale** from the **Y-axis scale** button set. You also have the ability to zoom into window by the usual rubber-band box type method. Hold down the left mouse button and drag the cursor to define a region into which you wish to zoom. Release the mouse button when you have defined the region. A right mouse click when the cursor is anywhere in the plot window will autoscale the display.

It is not necessary to propate error bars through the transformation in order to perform the deconvolution. However, if you wish to write the deconvoluted data estimate out to a DAVE file you must propagate the errors. Therefore we should now make sure that the error bars are propagated through the transform for the sample and resolution data sets. The default number of synthetic data sets used in the Monte-Carlo error estimation is 50 but you can change this in the box labeled **# Monte-Carlo Steps**. Simply check the box labeled MC error estimate and the Monte-Carlo method will proceed. A progress bar will appear indicating how close it is to completing the task. Depending on your computer and the size of the data sets, this step should take less than a few seconds. Do this step for both the sample and resolution data. When the errors have been propagated through the

transformation you can look at the data in magnitude I(Q,t) and see how the error bars now reflect the statistical uncertainty of the original data. This point is worth emphasizing: *the propagated error bars are related to the statistical uncertainty only, not the systematic error associated with using the FFT.*

Now that the error bars have been estimated for the sample and resolution files we are ready to perform the deconvolution estimate. Select the deconvolution tab and press the **Deconvolute** button. A progress bar will appear indicating how close it is to completing the task. Upon completion of the calculation the data will appear in the window. Do not be alarmed if the data "blows up" at long times. This is an effect of the resolution of the measurement instrument. You should always be aware of the limitations of the spectrometer and how resolution affects your ability to obtain reliable information on various time scales. For instance on this spectrometer the instrumental resolution is about 0.8 μeV. This corresponds to a Fourier time of about 5 ns (5000 ps). Therefore you should not be surprised to see this wild behavior for times around 5 ns and longer. This is expected.

In this example we did not choose any type of pre- or post-filtering. In general filtering of data is not advisable, especially when you are performing analysis on the post-processed data. If you do choose to use a pre-filtering option you should notice that the long-time behavior is more settled. Thus the deconvoluted (processed!) data will likely be more settled as well.

FREQUENTLY-ASKED QUESTIONS

Q: Can I use my own ASCII-formated data files with this software?

A: Yes but you will need to perform a few steps. First your data is assumed to be a function of energy in meV. Next your data must be formatted as 3-columns (x, y, dy). Then you must convert it to the DAVE format using the ASCII→DAVE data reduction application. You can find this under the Reduction menu item in the DAVE environment.

Q: Can I read in fits to data from PAN as a resolution function in this application?

A: Yes but you must be aware that the error bars on S(Q,E) for your fits from PAN are identical to the error bars on the data set to which the models were fit.

Q: The error bars are blowing up at large times for I(Q,t). Why is this?

A: Most likely you are seeing the effects of instrumental resolution being propagated through the FFT. As a seasoned user of FFTs you are already aware that there is nothing that you can really do about this other than measure the system on a different spectrometer with the same dynamic range and better instrumental resolution.

Q: How does the program propagate error bars through the FFT?

A: The application uses a Monte-Carlo algorithm. First the data, S(Q,E) and the associated statistical uncertainties on the data are used to synthesize a "large" number of additional data sets. These synthetic data sets are all propagated through the FFT and the standard deviations of the resulting transformed data are computed and used as the propagated error bars. The user has the flexibility to modify the number of synthetic data sets used in the Monte-Carlo calculation. Note that this algorithm only accounts for statistical uncertainty in the data. Other sources of error such as truncation, aliasing, and quantization are <u>not</u> included in the error bars.

Q: How much can I believe the error bars that are propagated through the FFT?

A: This is a difficult question to answer. The Monte-Carlo method (see previous question/response) only propagates statistical uncertainty. However there are other sources of error in the transformation that are not statistical in nature but relate to the resolution of

the data and dynamic range, for instance. See the answer to the previous question for more information.

Q: How does this program deal with the notorious "phase" problem in scattering?

A: This application sweeps the issue under the rug by assuming that the S(Q,E) data is real and that the imaginary part of the data is zero.

Q: How are the error bars handled when a filter is applied to the data?

A: The error bars are <u>not</u> modified when a filter is applied to the data.