

## Description of the Software:

The **csx** program was developed for the display and analysis of *in vivo* spectroscopy and spectroscopic imaging data sets. The program was initially developed in the **OpenWindows** environment of Sun Microsystems Sun-4 or SPARC computers; however, the code has also been compiled for Silicon Graphics and LINUX systems, as well as Sun ULTRA SPARC systems operating under Solaris.

**csx (and csx2)** supports all commonly used spectroscopy processing, display and analysis functions, and are particularly optimized for processing of spectroscopic imaging data. The plotting of spectra and results is supported on Postscript compatible printers. Currently, the program processes single voxel, 1D or 2D SI files, single or multislice (3D SI is not supported). File formats are supported from the following systems: General Electric SIGNA (3.X, 4.X, 5.X and LX spectroscopy files (“N”, “G”, or “P” files; PROBE files)), General Electric Omega, Bruker DISNMR files (Aspect series), NMR1, SMIS (.mrd) and Philips (.sdat). In addition, a general import tool allows the import of data files in other formats (e.g. ASCII, short integer or integer, floating point or “double”, real or complex, any header offset or matrix size).

**imax (and imax2)** is an image display program, primarily intended for use in conjunction with **csx** and **csx2**. It includes facilities for the overlay of images with MRSI “grids” and “point-and-click” selection of spectra from regions of interest. It also allows region of interest measurements to be performed and combinations (add, subtract, logical combinations) of images to be calculated. The following file formats are supported: General Electric SIGNA (3.X, 4.X or 5.X image files), DICOM, and SMIS (.sur). In addition, a general import tool allows the import of data files in other formats (e.g. ASCII, short integer or integer, floating point or “double”, any header offset or matrix size).

## Distribution and Legal Matters:

The programs available for download here are provided free of charge for the non-profit use of individual users. Distribution to other users is not authorized. We make no warranties, express or implied, that these programs are free from error. They should not be relied upon for solving a problem whose incorrect solution could result in injury to a person or loss of property. If you do use the programs in such a manner, it is at your own risk. The authors disclaim all liability for direct or consequential damages resulting from the use of these programs.

## Which Files to Download:

4 compressed tar files are available for download:

- csx.tar.Z - contains binaries for Sun-4 type systems (pre-Solaris/ULTRA SPARC)
- csx2.tar.Z - contains binaries for Solaris/ULTRA SPARC systems
- csxsgi.tar.Z - contains binaries for SGI IRIX systems
- csxlinux.tar.Z - contains binaries for LINUX systems

## Installation:

After you have downloaded the appropriate tar file (remember, ftp in binary mode), you will need to uncompress and extract the tar file using a command of the type “tar -xvzf \*.tar.Z”. Generally, the software you receive will contain 1 sub-directories: (1) “bin” (2) “lib”. The “bin” directory contains the binary executables for the programs and some external routines which may be called by csx. The “lib” directory contains some run-time libraries which the programs may need in order to run on your system. Copy all of the files in the “bin” directory into **/usr/local/bin**. The contents of the lib directory should be put into **/usr/lib**. Make sure your “path” includes these entries.

Note that on LINUX systems you will need to download the OPENWINDOWS runtime libraries in order to operate these programs. The LINUX libraries can be downloaded from several sources, e.g. [http://www.linux.about.com/compute/linux/library/metalab/blmeta\\_libs\\_X\\_xview.htm](http://www.linux.about.com/compute/linux/library/metalab/blmeta_libs_X_xview.htm)

## Running the csx program

The program can be launched by typing **csx**. The active directory and filename can be typed in the windows in the upper-left corner of the main control panel. Note that a list (menu) of directories and filenames can be generated by placing the mouse over each of these items and holding the right-mouse button down. The file or directory will automatically be loaded when the right mouse button is released. Make sure that the file format is correctly selected for the type of data you are loading. You may also set a “template” so that only files which match the first letters of the entered template will be displayed when you select the menu. If you wish to reload a file for any particular reason, hit the "**Load**" button. Hit the "**Quit**" button to leave the **csx** program.

Most supported file formats automatically read file header information about sweep width, number of data points etc., but for some files it may be necessary to manually enter these values in the pop-up window which appears after the “Load” button is hit. Note that multi-slice files are processed one slice at a time - the slice number to be processed must be entered in the pop-up window.

The general procedure in analyzing an SI data set is to first load the data set, then perform a 1D or 2D FFT (with respect to the distance dimensions) by hitting the “**2DFFT(X,Y)**” button. The 1- or 2D-FFT has various options for voxel shifting and/or filtering in the distance (k-space) domains; the default filtering is a Cosine function. At this stage the data set may be flipped vertically or horizontally or rotated right or left to match the orientation of a localizer MR image as needed. To apply the k-space filter without performing the subsequent FFT, hit the “**Apodize Only**” button; otherwise, select “**Apodize and 2DFFT**”.

After 1- or 2D-FFT, individual spectra from the matrix may be viewed and processed. The “**X Coord**” and “**Y Coord**” values may be changed (either by typing in a number, or by using the arrow icons (see below)) to view different traces from the matrix. The whole (or part) of the matrix may be displayed using the “Grid Plot” option from the display menu (hold down the right mouse button over the canvas region to see this menu). The region of the matrix which is plotted will be determined by the values of the **X Start**, **Y Start**, **X spectra to plot** and **Y Spectra to plot** items which are located at the bottom of the “**Display Options**” panel.

### **Vertical Scaling**

The vertical scaling of the FID (and of a spectrum) is controlled by the VS item at the bottom left of the main control panel. It can be set in three ways

- (1) hitting the "Auto Scale" check button will automatically adjust the vertical scale to fill the screen,
- (2) You can type in a scale number directly,
- (3) You can use the left and middle mouse buttons when over the arrow icon\* to adjust the scale up and down.

\* *arrow\_icon* = 

*Arrow icons are used for a number of purposes - scaling, phase correction etc. You can adjust the value they are associated with up and down by using the left and middle mouse buttons respectively (when the mouse is placed over the arrow icon).*

### **Time Domain Filtering**

To filter an FID, hit the **Filter** checkbox at the bottom of the main control panel. The default filtration is an exponential decay corresponding to the line-broadening factor displayed in the "**Time Domain Proc**" control panel. Alternative filter functions can also be selected - filtering options are exponential, Gaussian, convolution difference, left and right shifting, and zero-filling. Also, a time-domain convolution (high-pass) filter can be used (Marion, Ikura and Bax, JMR, **84**, 425-430 (1989)) for water-suppression (provided that the water signal is at zero frequency (i.e. in the middle of the spectrum)). The bandwidth of the (Gaussian) water suppression can also be specified in the control panel.

*For data sets which record spin echoes as opposed to FIDs, it is important to set the echo position (in units of number of points) so that the correct filter is used. Not that the default echo position is non-zero (128) for some File formats. (since most of the time I use spin-echo data)). If you want to phase correct spin echo data, select the “Echo?” checkbox.*

## **Fourier transformation and Phase Correction**

To view the Fourier transform of the FID, hit the "Fourier Transform" checkbox at the bottom of the main control panel. The spectrum can be scaled in the same way as the FID. To phase correct the spectrum, first of all use the **Zero Order phase correction** to properly phase a single peak (use the phase correction arrow icon). Place the cursor over this peak, and hit **p** on the keyboard. The phase of this peak will now be held fixed while the **first order phase** is varied in order to phase correct the rest of the spectrum. If you want, you can display the magnitude signal ( $= \sqrt{R^2 + I^2}$ ) by hitting the "Magnitude" check box in the main control panel. Alternatively, you can try checking the "AutoPhase" box, although there is no guarantee that this will work! Parameters effecting the AutoPhase routine can be changed using the "**Phase**" control panel.

There is also an option for applying first order phase-corrections in the distance domains of spectroscopic images if needed (note that this should not be necessary if the phase-encode gradients were correctly applied in the original data set (for an explanation, see Le Fur et al., MRM 29(4), 431-435, 1993). The best procedure for applying these corrections is to first phase correct a single spectrum in the normal way, and then hit the "**k**" key (with the mouse over the canvas - see below). The phase of this spectrum will then be fixed. Then change the value of the "**X**" or "**Y**" as phases needed until the whole matrix is correctly phased.

## **Frequency Domain Filtering**

Baseline-correction is possible using the "**Frequency Domain Proc**" options. The procedure is to first hit the "Define Baseline" button. Then use the right mouse button to define points in the baseline. When the baseline has been defined, hit the "Done" button. The program then calculates a continuous baseline function using cubic spline interpolation. To view the baseline function, select "Baseline" from the canvas display menu (right mouse button in canvas). To view the baseline-corrected spectrum, hit the "Baseline Correct" checkbox in the Frequency Domain Processing control panel.

There is also an option to "zero" part of a spectrum (between the left and right cursors on the screen) in the "Frequency Domain Proc" options.

## **Analysis Options**

Spectra may be analyzed in several ways in the “csx” program. Integration is possible (see below), as well as curve fitting in both the time and frequency domains. The curve-fitting method may be selected using “Setup Curve Fitting” from the “Analyze” menu. Once the curve-fitting method and parameters have been selected, the next step is to do a peak picking (either manually or automatically) to let the program know which peaks to fit. This can be done using the “Curve-fit” sub-menu under “Analyze” (it is also possible to re-use a previous peak-picking). Once the peak picking has been performed, the curve-fitting runs automatically. The results of the curve-fitting routine can be viewed by selecting the “Results” sub-menu from the Canvas menu (see below - hold right mouse button down over “Canvas” (FID/spectrum display window)).

The time-domain curve-fitting routine (“NLLS”) is described in more detail below. SI data sets may be “Batch Analyzed” (see below).

Finally, under the “Miscellaneous” sub-menu, it is also possible to make a signal-to-noise (S/N) measurement. First, expand the spectrum so that the display contains the (tallest) peak which will be used for the S/N measurement. Then place the left and right cursors around the region of noise which will be used for the S/N measurement. Hit the measure S/N button. The reported value is the result of the equation:  $S/N = \text{peak height}/2 * (\text{rms}(\text{noise}))$  (see R Freeman, “A Handbook of Nuclear Magnetic Resonance”, Longman, 1988). The “Miscellaneous” sub-menu also allows line-width and peak-picking measurements to be made.

## **Integration**

To integrate a spectrum, choose “Integral” from the canvas display menu (use the right mouse button over the canvas region). The integral will be displayed over the top of the spectrum. A control panel appears underneath the main control panel which allows the control of the vertical scale for the integration, baseline tilt and offset correction (of the integral). To set the integral to zero at desired points in the spectrum, put the mouse at the appropriate point and type **z**. To clear all zero-points, hit **c** while the mouse is over the canvas. When the integral is properly adjusted, the values may be saved to a disk file by choosing “Save Integral” from the control panel. This file can be printed using the “Print Results” command.

## **Printing**

**csx** currently only supports systems that are equipped with Postscript printers. You can “**Setup Printer Options**” from the “**Print**” menu. A control panel lets you select the plot size (height and width in inches), orientation (“landscape” or “portrait”), and printer name. There is also an option to print to a file (in either Sun “Rasterfile” or Apple Macintosh “PICT” file formats) as opposed to the printer.

The print menu lets you either plot the spectrum at high resolution or do a low-resolution “Screendump”. You can also print out the results of the curve-fitting routine. Finally, the print menu

also lets you dump out the contents of the display buffer in floating point format (“Dump Display Buffer to File”).

### **Batch Analysis - Production of Metabolic Images**

It is also possible to analyze all or part of the SI matrix in a single step using the "**Batch Analysis**" option of the "**Analysis**" Menu. The region of the matrix which is analyzed is determined by the values of the **X Start**, **Y Start**, **X Width** and **Y Width** items which are located at the bottom of the “**Display Options**” panel (“**Sub-matrix Dimensions**”).

The general procedure is to expand the spectrum to display only a single peak of interest. Options are available to measure and/or report various combinations of the peak height, frequency, or integral (with or without baseline correction ("tilt")). The output file format can be a Signa file, short int, sun floating point or ASCII. The output file (whose name is specified by the "**Output Filename**" item) can also be interpolated (e.g. to create a 256x256 file from a 32x32 2D SI data set, the interpolation factor should be set to 8) with either smooth linear interpolation or "jagged" voxels, and either scaled automatically (from 0 to 10000) or manually using the specified **manual scaling factor**.

If outputting a Signa file it is also possible to specify various parameters which will be included in the file header such as the Study, Series and Image number, as well as text strings describing the study and series.

Before creating a metabolite image, it is necessary (usually) to correct for voxel-to-voxel frequency shifts (“susceptibility correction”). This allows a narrower frequency range for integration and therefore improves the signal-to-noise ratio of the metabolite image. The susceptibility correction is based on finding the frequency of some reference peak common to all the spectra; for instance in a proton brain SI data set, the NAA peak can sometimes be used, although it may also be possible to use the water signal, or the PCr signal in a  $^{31}\text{P}$  data set.

To set the reference peak, place the cursor over the peak of interest, and type "s". Note that the vertical position of the cursor sets a threshold below which peaks are ignored by the peak-picking routine. The program searches for a peak over the frequency range specified in the "**Susceptibility**" control panel. Once the susceptibility correction ("shift table") has been determined for each spectrum in the SI matrix, the correction is automatically applied. To "un-apply" the shift correction, select “Off” for the susceptibility correction switch in the "**Susceptibility**" control panel. Note that it is also possible to load a susceptibility correction from a pre-existing file (e.g. created by a field-mapping MRI pulse sequence) in this control panel also.

Batch analysis can be time-consuming, particularly for large data sets. If you want to interrupt a long analysis process, you should be able to do this by hitting the “STOP!” button in the main control panel.

## **Display Options**

The display options panel can be opened by hitting the "Display" button in the main control panel. The display can be annotated with a filename, date (the current date + the date which the data was recorded on), an axis either in Hz/ppm or msec (see below), a horizontal line indicating the baseline (Zero Line), and a box around the display (Box). "Plot Mode" specifies where the baseline appears in the display (see "Canvas Commands" section on how to set the vertical position of the baseline).

The pen size (thickness) can also be specified, as well as the spectrum expansion range (Left ppm and Right ppm). For the axis, it is possible to specify the increment between axis labels, the number of minor tick marks between major ticks, and the number of significant digits in the axis labels. In the same control panel, it is also possible to set the dimensions and resolution for plots, the printer to be used, and the resolution of PICT files (see below).

## **Canvas Commands**

A number of commands work by placing the mouse over the "canvas" (i.e. the FID/spectrum display area) and typing a letter. Here is a list of all the canvas commands:

- a** set position of axis (to mouse location)
- b** set position of (spectrum/FID) baseline (to mouse location)
- c** clear all integration zero points
- e** expand display horizontally (between left and right cursors)
- f** display full spectrum/FID
- k** set x,y voxel fulcrum for first order phase correction with respect to distance
- t** set a threshold for the automatic peak-picking routine
- s** start susceptibility correction
- p** set the position of the "fulcrum" for 1st order phase-correction
- h** set the reference frequency of the Hz/ppm scale (after typing h, the cursor automatically moves up to the "ppm" item in the main control panel - type in the ppm value which you want the reference set to)
- z** set a zero-point for the integration routine
- >** scroll expanded region right
- <** scroll expanded region left
  
- ctrl-h** move the cursor left 1 pixel

- ctrl-j**      move the cursor down 1 pixel
- ctrl-k**      move the cursor up 1 pixel
- ctrl-l**      move the cursor right 1 pixel

In addition, the mouse buttons perform various functions in the canvas:

- Left Mouse Button**            controls position of left cursor
- Middle Mouse Button**        controls position of right cursor
- Right Mouse Button** shows menu which controls display. (The right mouse button is also used for peak-picking and baseline correction when in these modes)

**Some Notes on Time-Domain Fitting**

The **csx** program supports a time-domain non-linear least squares (NLLS) fitting routine (Society of Magnetic Resonance in Medicine, 9th Annual Meeting, New York, NY, August 1990. P.B. Barker and S. Sibisi, “Non-linear least squares analysis of in vivo 31P NMR data”, p. 1089). The NLLS program was originally developed by Sibusiso Sibisi while on sabbatical at the California Institute of Technology in 1987-1988. The algorithm is similar to the VARPRO method of van der Veen et. al. (J.W.C. van der Veen, R. de Beer, P.R. Luyten and D. van Ormondt, *Mag. Res. Med*, **6**, 92-98, 1988) and the least-squares fitting method of Derby et. al. (K. Derby, C. Hawryszko, and J. Tropp, *Magn. Reson. Med.*, **12**, 235-240 (1989)). NLLS involves the iterative, two-step minimization of the difference ( $\chi^2$ ) between the model equation  $F(t)$  and the experimental FID

$$\chi^2 = 0.5 \sum_{k=1}^m (F(t_k) - D_k)^2$$

$$F(t) = \sum_{j=1}^n a_j \exp(-1/T_j + i \phi_j)t + \phi_j$$

where  $D_k$  is the  $k$ th of  $m$  data points in the FID, and the  $j$ th of  $n$  oscillators has amplitude  $a_j$ , phase  $\phi_j$ , frequency  $\omega_j$  and decay time  $T_j$ . In the linear step, phases and amplitudes are fit whilst the decay times and frequencies are held fixed. In the non-linear step, the frequencies and decay times are fit whilst the linear parameters are held fixed. The starting values which are required for the non-linear parameters can be obtained from the normal Fourier Transform of the FID. Nearly any combination of parameters can be held fixed (amplitude, phase, frequency or decay time) if these values are already known. When a predefined stopping criteria is fulfilled (usually that the relative misfit ( $\chi^2$ ) does not decrease by more

than 0.1% over successive iterations) a synthesized FID is output from the fitted parameters for graphical confirmation of the success of the fitting routine. NLLS is implemented in FORTRAN 77 using LU decomposition of the normal equations for the linear fit, and the Levenberg-Marquadt method for the non-linear fit. Fitting time for a typical in vivo spectrum is approximately 30 seconds. Current (somewhat arbitrary ) limits on the program are (1) Maximum number of data points used for fit = 512, (2) Maximum number of peaks to be fit = 20.

It is not uncommon for the fitting program not to give satisfactory results the first time you use it. Often, the data must be carefully filtered and the proper starting values specified in order to produce a successful convergence. A number of options are available through the “Setup Curve Fitting” command of the “Analyze” menu. You can specify which parameters you want to fit (Frequency/Phase/Amplitude/Decay Time). Our experience has been that the most reliable results are obtained by carefully phase-correcting the spectrum, and then holding the phase fixed during the analysis. You can change the number of points to fit, the number of points to ignore at the beginning of the FID (“kiss-off”), the linewidth estimate (for automatic peak-picking), the number of non-linear iterates, and the global tolerance (i.e. the stopping criteria). (You should rarely need to change the “Reconstruction Fudge Factor” or the “Phase Angle Fudge Factor”). In addition, the fit may improve by using some convolution difference filtering to flatten the baseline, increasing the line-broadening to give a more Lorentzian lineshape, increasing (or decreasing) the number of peaks to fit, changing the starting values of each peaks linewidth, or specifying one peak with a very broad linewidth which will be used to fit the baseline. Simple, well-resolved spectra with flat baselines and good lineshapes will obviously be much easier to fit than poorly resolved, poorly shimmed spectra with undulating baselines. For spectra such as these, it may be difficult to produce satisfactory results.

*Note: **nlls** is a Fortran-77 program called by **csx**. Therefore, you need the runtime Fortran libraries (**libf77.so**) in order to run **nlls**.*

The results of the analysis program are contained in files with **.par** extensions. A utility program called **par\_to\_mac** is supplied which extracts certain parameters from multiple **.par** files and stores the output in a single, tab-delimited file which is suitable for export to other computer systems for further analysis. The **pimms** program performs a similar function for NLLS-batch analyzed SI data sets.

### **Some Notes on GE PROBE Spectra**

Probe data files are stored as GE Signa “raw” data files (filenames of the type P01234). The sequence first records 32 scans without the water-suppression pulse, and then a variable number of water-suppressed scans (specified by the “number of averages” variable). The storage of the data

depends on the phase-cycle used; if a N-step phase cycle was selected, every N scans are added and stored in memory as one trace.

In the main control panel of the **csx** program, select **File Format** as either **GE SIGNA 4.X Probe** or **5.X Probe** as appropriate: if 4.X, also specify whether data is “short” or “long” integer (“long” if extended dynamic range was used). Load in the file; for a 4.X file, you will need to specify the number of points and the spectral width, and you also need to specify the number of water spectra and the number of metabolite (water-suppressed) spectra. This will depend on the number of averages which you have used and the phase-cycle (“NEX”).

Number of water scans = 32/phase-cycle

Number of metabolite scans = number of averages/phase-cycle

*Example: number of averages = 128, phase-cycle (“NEX”) = 2*

*Number of water scans = 16*

*Number of metabolite scans = 64*

After loading, the water FID will be at location X=0 (i.e. set the X Coordinate to 0), and the water-suppressed FID will be at X=1. The FIDs may now be processed by filtering and Fourier transformation in the same way as normal spectroscopy files.

## **Running the IMAX program**

The program can be launched by typing **imax**. The active directory and filename can be typed in the windows in the upper-left corner of the main control panel. Note that a list (menu) of directories and filenames can be generated by placing the mouse over each of these items and holding the right-mouse button down. You may enter a character string in the "**Template**" item to restrict the files displayed in the filename menu. The file or directory will automatically be loaded when the right mouse button is released. Make sure that the file format is correctly selected for the type of data you are loading. If you wish to reload a file for any particular reason, hit the "**Load**" button. In a series of images (for instance, consecutive image files such as I.001, I.002 etc...) it is possible to sequentially load files with file arrow icon. (Note: if skip is set to 2 (for instance), hitting the left mouse button on the file arrow icon will load files I. 001, I. 003, I. 005 etc.. sequentially. The middle mouse button decrements the file extension. Hit the "**Quit**" button to leave the **imax** program.

To change the brightness and contrast of an image, use the "**Level**" and "**Width**" arrow icons in the main control panel. Note that this works by changing the computer "colormap", i.e. the gray-scale of the entire monitor (all images) changes simultaneously. Once an image has been correctly windowed, it can be rescaled to fit a linear (0->255) grayscale by hitting the "**Apply LUT**" button; this allows images with different amounts of contrast to be windowed and viewed simultaneously.

If a display "**Size**" of 256 is selected, 4 images are displayed simultaneously in the canvas. If size is 512, only a single image is displayed (selected by the "**Window**" item in the control panel). The image can be flipped either vertically or horizontally (or both) if needed by hitting the appropriate check boxes. It is also possible to overlay grids in either the horizontal or vertical directions (or both - "**X Grid**" and "**Y Grid**"). The grid size is controlled by a panel which is displayed at the bottom of the screen when the "**FOV/Voxel Size**" button is hit. This panel allows the image **FOV**, **CSI FOV**, **CSI matrix size** to be entered. It also displays the slice location of the current image, as well as the location of the center of the image in terms of patient coordinates (Right/Left, Anterior/Posterior, Superior/Inferior). The grids may also be shifted by using the **Voxel Shift X** and **Y** items (entered in units of mm).

The main control panel also displays the current position of the cursor over the image (**X** and **Y** - in mm), and also reads off the cursor position in terms of which voxel in the SI grid the cursor is located (**X vox** and **Y vox**).

It is also possible to annotate the image with "boxes", e.g. drawn to depict the location of single voxel NMR studies performed, e.g., using the STEAM or PRESS pulse sequences. Position the cursor at the location of the box center, and then hit the "**b**" key. The size of the box is governed by the "**Box Size X**" and "**Box Size Y**" items (in mm). Up to 10 "boxes" may be displayed simultaneously. To remove a

box, type "c". To remove all the "boxes", type "a". The color of all overlays is governed by the "Overlay" item which can be used to select color (green), white, black or gray.

The annotated image can be saved as a disk file using the "**Save ->**" menu in the main control panel. Currently, the images are saved as 8-bit files without any header ("8-bit for Mac") with the suffix ".a" appended to the file name (e.g., if the original image filename is I.003, the annotated image is saved as I.003.a).

The **imax** program can also be used to set the voxel location to be displayed in the **csx** program. Simply click the left mouse button over the region of interest of the image. Note that this will only work reliably if both programs are launched from the same directory. The mechanism of communication between the 2 programs is a UNIX "socket" - you will see an entry in the working directory called "sock" while the programs are running. Occasionally, something can go wrong and the "sock" file cannot be opened; this is usually because (1) you do not have write-permission in the directory you are running the programs from, or (2) the "sock" file already exists. Solutions are (1) to go to a directory where you have write-permission (this is also a good idea if you want to create output files (even temporary ones)!, and (2) to delete the "sock" file with the UNIX command **rm sock**.

The **imax** and **imax2** programs run under 8-bit color. If you are running in 16 or 32-bit color ("true-color") on your systems you will need to run the "imax2\_tc" program.