FAST

User Guide for the Use of the

Satellite Data Tool (SDT) for Data Analysis

An Introduction to SDT

The following document is designed to help you get started using the Satellite Data Tool (SDT) to acquire and browse data from the FAST satellite. For details on the use of SDT and its associated utilities, see the documents referenced in the [More Details] section of this document.

SDT is a program developed to aid in the acquisition and display of data from a variety of satellites using a uniform user interface. Originally developed to support the CRRES and Polar missions, SDT now supports many other missions, including FAST and Cluster II. Below we describe what SDT is, how to set yourself up to use SDT, and how to run SDT.

SDT is actually a whole ensemble of programs that work together to allow for easy access to data from a variety of different spacecraft. All of the user interaction with the ensemble occurs through a user interface, or UI, program. Each satellite has a decommutation program that takes data from raw telemetry data files and converts it to usable scientific data. Other programs (SCM, DQH, etc) aid SDT and the decommutation programs in selecting and acquiring the correct data.

In the context of the FAST mission, SDT acts as a tool for acquiring fields, particle, and other data from mass storage (CD-ROMs or hard drives), and browsing that data to identify periods of interest for display or detailed analysis. The data from several different channels on a given spacecraft can all be shown on a common timeline, and the style and details of the presentation of each channel can be manipulated to varying degrees depending upon the type of data involved. Some calibration and reduction processes are supported within SDT, such as basic calibration of fields data and projection of fields into various coordinate systems. More extensive data reduction and calibration, such as extraction of magnetic field perturbations, precision fields despin, and calibration of particle data, is taken care of using an extensive library of IDL routines.

This page consists of several sections that give a user a basic knowledge of what's required to operate SDT. Like most pieces of data analysis software, the best way to learn SDT's use is to use it, exploring the menus and options. To that end, this document does not pretend to give an exhaustive documentation of the functionality of SDT; rather it serves to explain some of the less-obvious things required to run SDT successfully, along with some of the basic functionality required to allow a user to find their way around on their own.

SDT Setup

Software Distribution and System Requirements

SDT is distributed as part of the FAST SDT/IDL software package, available as a compressed tar archive on the [FAST FTP Site]. A computer running SunOS 5.x (Solaris 2.x) is required in order to run SDT. While not required, at least 256 Mbytes of RAM should be considered a minimum requirement of any system running SDT, and if one expects to be doing extensive fields analysis, 512 Mbytes to 1 Gbyte of RAM should be considered.

When installed, the FAST SDT/IDL software package typically resides in a directory structure rooted at /disks/fast/software. Some sites use a different root directory for the package, such as /usr/local/fast/software. If one's site is using a different root directory, then any references to /disks/fast/software should be replaced with the actual directory name.

Supporting Processes

SDT is supported by a batch job that downloads important spacecraft and telemetry database information from UCBSSL to the local installation. The setup of this batch job is beyond the scope of this document, but is essential to allow the proper functioning of SDT, Data Manager and associated programs in the FAST SDT/IDL software package.

System File Modifications

SDT stores the data it has acquired in shared-memory buffers so as to allow IDL or other external programs direct access to that data. The kernel of any machine running SDT needs to be properly configured in order to allow for the size and number of shared-memory buffers typically used by SDT during a given session. The following commands will need to be added to the /etc/system file of any machine that will be running SDT, and that machine re-booted in order for the changes to take effect:

```
* kernel configurations for FAST:
*
* increase maximum shared memory size 125 meg
set shmsys:shminfo_shmmax=125000000
* allow 300 shared memory identifiers
set shmsys:shminfo_shmmni=300
* allow 200 shared memory segments per process
set shmsys:shminfo_shmseg=200
```

Setup Script

A setup script is included with the FAST SDT/IDL software distribution. It is located at /disks/fast/software/delivery/setup (or setup.bourne if you prefer the Bourne shell to csh or tcsh; see the root directory

discussion above as well). When sourced in, either from your .cshrc file, or directly from the command line, the commands contained therein define or modify several different environment variables, allowing you to access the FAST SDT/IDL software.

Data Manager Configuration

SDT and its data-handling utility Data Manager, need to be told where you wish to store your data files. This is done through entries in a file called .datamgr_conf (note the leading period in the file name), located in one's home directory. A basic version of this configuration file is included as part of the FAST SDT/IDL software package, located in the following location:

/disks/fast/software/config/Datamgr/Datamgr.conf

Copy this file to the file .datamgr.conf in your home directory, then edit it according to the directions found near the end of the file, changing the entry for DATA_DIR to be the location of your data directory. Your data directory needs to reside on a file system with lots of space, or you need to be committed to a stringent program of data maintenance. The data files for a typical FAST orbit take up 150-200 Mbytes of disk space, and your data directory should have space for a couple of times that in order to cut down on time spent deleting old data to make way for new.

Running SDT

The process of running SDT consists of several steps:

*** Startup**
*** User Interface Configuration (UIcfg) Selection**
*** Data Selection and Acquisition**
*** Data Browsing**
*** Shutdown**

Each of these steps will be described below.

Startup

Starting SDT from the UNIX command line is easy; at your prompt, type sdt. As part of its operation, sdt will generate several diagnostic output files (at a minimum, the files generated will be as follows: errdqh, errfast, errUI, outdqh, outfast, outUI). Most users will run SDT from a particular sub-directory in which these files will reside, along with any personal UIcfg files (see UIcfg below).

There is currently a limit of one SDT session per machine imposed by the sdt script. This limit was originally set up to deal with lack of resources on early-1990's era workstations, as well as ownership problems with shared memory (?). It may be eliminated in the future.

Many users will receive a warning about support of OpenWindows code going away. Ignore it.

User Interface Configuration Selection and Modification

Once SDT starts (a process that can take up to a minute), a control window will popup. Two of the possible selections are of initial interest. If you wish to build a new User Interface configuration (UIcfg) from scratch, then you would choose the `Add Window' button.

Rather than do that, you can use one of the stock UIcfg files included with the FAST SDT/IDL distribution in order to get started (that's what we all did to learn!). WHERE TO GET THE STOCK UICFGS, HOW TO COPY THEM TO THE SDT OPERATING DIECTORY. Such pre-existing UIcfg files are accessed via the `Configuration' button. Selecting `Configuration' pops up a selection window. Clicking on the desired UIcfg in the Configuration list, and then clicking on the Select button will tell SDT which UIcfg you wish to work with. If you inadvertently end up in the Configuration selection, then you can escape using the Dismiss button.

If the data supporting the time interval that the chosen UIcfg was last looking at is still on-line, then SDT will ask one if one wishes to look at that data (the Okay selection), or bring up the raw UIcfg with empty plot panels (the `RawPlots' option) , or cancel one's request for the given UIcfg (the `Cancel' option). If one is returning to the analysis of a previously acquired time interval, one usually chooses the `Okay' option, while if one is just beginning to analyze a new time interval, one chooses the `Raw Plots' option.

Each UIcfg consists of a set of data channels to acquire and display, along with all the styling’s for each channels panel within the SDT display window, as well as the particular time span or orbit that the UIcfg was last used to view. Data channels are referred to as DQIs (data quantity identifiers) or DQDs (data quantity descriptors); both acronyms are used interchangeably in the documentation for SDT and the IDL data analysis routines. The channels included in a particular UIcfg can be modified using the Main and Plot Menus of the SDT Plot Window; such operations will be discussed in detail below.

Saving and Deleting UIcfgs

When one has created a new UIcfg, or made modifications to an existing one, one will want to save that UIcfg. This is done by typing the name of the UIcfg in the space provided in the Configuration Manager window, and clicking on the Save button.

If one wishes to remove a UIcfg from the list of possible choices (and delete the file corresponding to that UIcfg from disk), one can select the UIcfg by clicking in the selection list window, then pushing the Delete button.

Data Selection and Acquisition

Once you have chosen a UIcfg in SDT, you then need to select and acquire the data that you will browse using that UIcfg. This is done via the Data Manager utility, accessible through the Main Menu of the SDT Plot Window. You access the Main Menu on SDT by right-clicking with your mouse in any of the space outside of the plot panels of the SDT window (the neutral space). Data Manager allows you to load data from mass storage (CD-ROMs (local or on the

jukebox at UCBSSL) or hard drive) to your local data directory so that SDT can see the data. You can look for data on the basis of orbit number, time span, or more complex search criteria (operational or geophysical). Once one has specified the parameters of one's data request, one pushes the `Submit Query', and waits for Data Manager to report back on the results of the query.

Now for some unfortunate, but necessary jargon and acronyms. The data from FAST is organized into bundles of channels known as APIDs. Each APID consists of several different fields quantities, or a particular type of particle data, or spacecraft ephemeris/housekeeping data. SDT determines which APIDs are required to provide the DQIs selected in the UIcfg, and passes that information to Data Manager. Data Manager then determines where to acquire the desired data from (on-line storage, local or remote CD-ROMs), and reports back how many of the requested files are currently available on-line, and how many will need to be transferred.

Select `List Files ...' to see a list of the files to copy, and select `Get Files' to tell Data manager to go ahead and copy the files to your local data directory (this can be a good point at which to get some coffee...). Once Data Manager finishes copying the files, then `Dismiss' the file list, and click either of the `Return to SDT' buttons on the main Data Manager window.

SDT will then redraw, and state that it's getting data from the decommutator (meaning it's pulling data from the files and translating it). SDT's progress on this task will be shown by an updating percentage in the title bar of the SDT window. You can freely zoom in on particular time spans while the decommutation process is happening (see Browsing below).

If you add a new plot panel to your UIcfg during a SDT session, the data for that plot panel may not be available from the APIDs loaded to support your original data request, and you will get an empty plot panel because of this. If this occurs, one needs to return to Data Manager, resubmit the query, load any additional data files necessary to support the additional channels of data, and then return to SDT to continue with one's browsing of the data.

Browsing Data

Now that one has set up the Plot Window by selecting a UIcfg, and acquired some data to look at using Data Manager, one is ready to begin browsing around using the facilities available through the Plot Window of SDT.

There are two menus that are used within the Plot Window of SDT in order to change the time interval that is being examined, along with the styling’s of the individual plot panels within the Plot Window. The Main Menu is accessed by right-clicking anywhere within the neutral space surrounding the plot panels in the Plot Window. The Plot Menu for an
individual plot panel is accessed by right-clicking within a particular plot panel. Each sort of menu allows one to perform a variety of useful operations that will be described in some detail below. More details can be found in the following document:

/disks/fast/software/delivery/docs/SdtUse.

A basic set of operations that one needs to know in order to use SDT effectively is as follows:

- **Zooming and Panning**
- **Adding, Deleting, and Rearranging Plot Panels**
- **Plot Panel Stylings**
- **Printing**

Zooming and Panning

One can zoom in on particular time intervals using the zoom function, accessed through the Main Menu.

To zoom in on a particular time interval, one first left-clicks on the beginning of the time interval anywhere in the neutral space surrounding the plot panels. One then middle-clicks on the end of the time interval. The selected time interval will be denoted by vertical red lines across all of the plot panels on the page. Selecting the `Zoom' option of the main menu brings up a sub-menu; selection of `In' on that sub-menu expands the selected time interval to fill the entire width of the Plot Window.

To zoom out, select a region of the plot window as if one were zooming in. By selecting Zoom-Out from the Main Menu, one then compressed the current time interval to fit into the selected width of the Plot Window, the remainder of the Plot Window being filled with the data from surrounding times.

One can also zoom to a specific time span using the Zoom-Text option; a dialog box appears in which the exact date and time of the start and ends of the desired time interval can be specified.

If one left- or middle-clicks within a plot panel, rather than in the neutral space surrounding the plot panels, then one can accidentally zoom one plot panel independent of the other plot panels. This is usually counter-productive, leading to the data from different channels being out-of-sync. One can recover from this sort of an error by selecting the Zoom-All option. This zooms all of the plot panels out to the full time span available for the dataset under display.

One can specify the number of plot panels to display within the Plot Window using the Plots/Page option on the Main Menu.

One can page up and down through the plot panels using the slider bar on the side of the Plot Window. One can also use the Up Arrow, Down Arrow, Page Up, and Page Down buttons of the numeric keypad in order to pan up and down through the plot panels a panel or a page at a time. One can also use the Left and Right Arrows to pan through the data a time interval at a time, or access the same functionality via the `Pan' option of the Main Menu.

One can prevent the redrawing of the Plot Window by selecting the `Lock/Unlock' toggle from the Main Menu; by locking the display, then panning up or down, one can avoid time-consuming redrawing of the Plot Window on slower workstations.

Adding, Deleting, and Rearranging Plot Panels

One can add, delete, or rearrange the plot panels making up a given UIcfg using options available on the Main and Plot Menus. To add a new plot panel, select Add Plot on the Main Menu. One will be presented with a series of sub-menus where the spacecraft (/FAST/) and particular type of data (/Fields-Survey, ESA-Burst/, etc.), and then particular channel (or DQD) of that sort of data (/V1-V4_S, ESA_Energy/, etc.). Once the type of data is selected, one is free to select as many different DQDs of that sort of data at one time as one wishes, speeding the construction of new UIcfgs.

Once the new DQDs have been selected, click on `Done,. SDT will then state that it is acquiring new data from the decommutator. It is often the case that when one adds a new DQD that data was not loaded as part of the original data acquisition, and rather than a new plot panel full of new data, one receives instead an empty plot panel. This indicates that either the given DQD is not available in the given Fields, ESA, or TEAMS mode for the time span or orbit at hand, or that the APID corresponding to the new DQD was not loaded in the original data request. After checking the Mode Sheets and Mode Tables to insure that the given DQD is actually available, one needs to resubmit the data request via Data Manager, and wait while more data files are downloaded to your DATA_DIR.

The order of plot panels can be rearranged using the Rearrange Plots option of the Main Menu. This option allows one to select blocks of plot panels and move them up or down through the UIcfg until things are just right.

Plot panels can be removed from the UIcfg either individually or in blocks using the Delete option of the Plot Menu. Right-clicking within a given plot panel brings up the menu, and selecting Delete will bring up a sub-menu, allowing one to select multiple plot panels to delete (Multiple), cancel the deletion (Dismiss), or delete the current panel .(Delete).

Plot Panel Styling's

The Plot Menu accessible by right-clicking within a given plot panel allows one to modify a variety of parameters that control how the data from a particular channel will be displayed. Most of the options are straightforward, and control obvious properties of the panel such as: point/line style y-axis and z-axis (for spectrogram-type data) limits, lin/log, etc. pitch angle, energy, or mass/charge selection for particle data (ESA and TEAMS).

Note that none of the setting of the plot panel options affects what data is available for more detailed data analysis using IDL.

Printing

The Print and Print Setup options of the Main Menu allow one to print the contents of the Plot Window to a printer or file in a variety of formats (postscript, gif, etc.), as well as save those printing parameters for later use. These control panels brought up by these menu options are straightforward, and allow for the selection of printing versus saving to file, output format, file name, plotting density (in the case of postscript; reduces the pixel density of the output at the expense of resolution).

Shut Down

Because SDT is a complex ensemble of programs, it pays to shut it down in a controlled fashion. The details of recovering from an uncontrolled shutdown can be found in the SdtUse, under `Cleanup By Hand', and basically consist of ways to ensure that all the shared memory allocated by SDT is properly returned to the operating system.

One can exit from SDT in a controlled fashion using the Exit option of the Control Window. One will be prompted as to whether one wishes to save the changes (if any) that one has made to the current UIcfg (plot panels, time span, etc.), or if one wishes to quit without saving the changes. Selecting Cancel allows one to select the Configuration-Save option; one can then select Quit and exit normally.

Authorship

SDT, the supporting software for basic calibration and data analysis, Data Manager, and the SDT/IDL interface have been written over the course of nearly a decade of work by the following authors: Ken Bromund, Jon Loran, Winston Teitler, Jack Vernetti (UCBSSL).