Presentation of Data from the Fermilab Datalogging System

Iouri Ivanov, Kevin Cahill, Brian Hendricks
Fermilab
Accelerator Division Controls Department
MS 347
P.O. Box 500
Batavia, IL. 60510 USA

Presentation of data from the Fermilab datalogging system is realized as an application program running under the Fermilab accelerator control system. This program utilizes a data acquisition system which is comprised of more than forty datalogging processes supporting more than 15,000 channels of which 9,200 are in use. Presentation of data can be performed for up to eight devices simultaneously, and is possible in two modes: time plot and X versus Y, where both X and Y are devices. Performance of the system allows retrieval and plot rates of several thousand points per second. Also the user may plot simple arithmetical expressions of devices, plot integrated curves for each expression (or one device), plot fitting functions, and display some simple statistical values about the data.

1. GENERAL

Presentation of data from the datalogging system is an important task which allows a user to spot long term trends for both technological and physical processes of the accelerator. The present work has resulted in the creation of a console program that reflects the wishes of accelerator operators, engineers, and researchers to study these trends. Flexible plotting modes coupled with retrieval and plotting rates of several thousand points per second encourage users to make maximum advantage of the datalogging system.

The program consists of two main parts which are mutually linked. They are the data update and the user interface.

The data update portion includes: retrieval of data from the datalogging system, mathematical transformation of data for graphing, mathematical methods of analyzing the data, presentation of data in both graphical and text forms and data exportation using e-mail communication.

The user interface includes all options and buttons which allow the user to define the setup of parameters controlling the session. Those parameters are as follows: time interval of interest, list of devices or arithmetical expressions comprised of devices, the view of graphical presentation of such expressions, etc. Moreover, there are means which allow the user to simplify the task of defining such a setup having from 15 to 100 parameters per plot trace. There are also provisions for saving previous setups for future reuse. There is both a circular buffer of the most recent plot setups as well as a buffer of files into which favorite plot setups can be saved.

2. DATA STRUCTURE AND DATA COLLECTION

The fundamental unit of data is a device or channel. The datalogging system supports more than 15,000 channels on more than forty nodes (9,200 channels are in use).

Presentation of data can be performed for up to eight channels simultaneously and is possible as either a time plot or an X versus Y plot, where X and Y are both devices. Devices have a number of attributes. These include a name, an array index number, a datalogger node, a number of points to skip, and a minimum time between points. The last two attributes are used to decrease the number of points read for densely populated devices.

Some devices (or even all of them) may be used as operands in simple arithmetical expressions. Each expression occupies a grid in the graphical window. In this sense, it is convenient to consider a single device that isn't part of an expression as a trivial expression. All grids for one session are overlapped in a graphical window, and two such windows are available to the program. Plot grids have attributes as follows: initial and final values for the Y axis, the type of Y axis scaling - linear or logarithmic, the type of trace, the fitting function type - polynomial (up to third order) or exponential, and the presence/absence of statistical information. Statistical information includes integral, average, and standard deviation values.

All the grids together comprise a single plot. A single plot has its own title, number of divisions for the X and Y axes (the same for all grids), and data point matching algorithm. The data point matching algorithm does not matter for time plots and the trivial expression case. However, when two or more devices are involved in one expression (like a nontrivial arithmetic expression for a time plot or any expression for the X versus Y case) the method of data point matching makes a difference. The menu of choices includes interpolation and matching with another menu for specifying accuracy.

The time interval made up of a start time and an end time is an attribute of the entire plot. So, it is common for all grids.

Data collection for each device on the page is organized by calling a library routine which receives as parameters the device attributes (or their derivatives) and the desired time interval. This routine sends a request to the datalogger process on the node specified and receives data from it. The returned data is packed in two arrays of the same dimension. One contains values in engineering units, and the other contains corresponding timestamps in seconds. These arrays for all devices of a single plot are kept in program memory and are freed just before the next plot session starts.

3. PRESENTATION OF DATA

3.1 GRAPHICAL PRESENTATION OF DATA

Before collecting the data, the program creates up to eight overlapped grids on a graphical window and labels them with different, hard-coded colors which serve to identify and distinguish them.

After having completed the collection of all data involved in the expression, the program starts the plot. If there are two or more devices, the user has a choice of establishing a correspondence between points of different devices. The user may choose either "Interpolation" or "Matching". In both cases, the program identifies the device with the least number of points collected. In a loop for each of these data points, the program either calculates the interpolated point for any other device in the expression that corresponds to the timestamp of the selected device, or looks for the point with the closest timestamp within the limits of accuracy for that point in the matching case. Then the point coordinates are saved for plotting.

There are various plot trace options. These options include whether or not to connect points, what symbol to plot, symbol size, whether or not symbols are filled, and the number of symbols to plot.

The symbols available are circles, squares, triangles, diamonds, and crosses. Also, special ASCII-characters can act as symbols depending on the order number of the grid. These characters are X O I = N Z $ %. If a character is chosen as the plot symbol, the size and fill parameters are ignored. The number of plotted symbols per trace can range between 0 and 65535. It is also possible to have every point plotted as a symbol. The color of the symbol plotted corresponds to the color of the associated grid.

Along with the curve of the expression, the integrated curve (multiplied by a user selected scale factor) and a fitted curve can be drawn . Both of those options are available for both time plots and X versus Y plots, where X is a device and Y is an expression. In the latter case, two options are supported. In the first one, X and Y are fitted independently, and their fitting functions define the equations of the parametrically defined fitting function. In the second option, the program fits just the resulting set of points for Y(X).

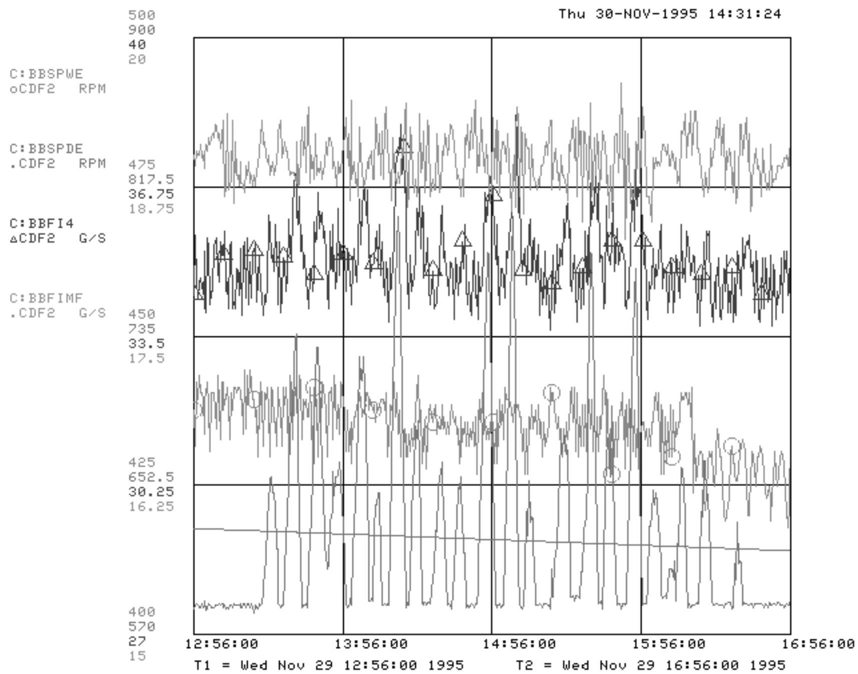

Fig. 1. An example of graphical data presentation with one exponential fit.

3.2 PRESENTATION OF DATA IN TEXT FORMAT

The numbers of points read for each device and plotted for each trace are displayed on the background text window. Also displayed there is the statistical information that includes the integral value, the average value, and the standard deviation value.

To explore the curve behavior in a manual way, the user may click on the "Calc Points" button. A popup window will be created to display the current grid coordinates of the cursor in the graphical window in the grid's units. The maximum and minimum values of expressions on the grid may be displayed as well after clicking on the "Min/Max" button.

After the plot has been completed the user may list the data in a popup window. There are options to list either the data actually read or the plotted data. To eliminate points from the data read, there exists an option to edit the collected data set (not the plotted data) and then to replot the edited data.

All statistical values are displayed on the graphical window and appear in the text window upon request.

The fitting function equations appear in a popup window upon request and can be displayed with different time units as the independent variable.

After the plot has finished the user may send the plotted data in pairs (time or X-value, expression value) to a user's e-mail address. The time values can be formatted in three different ways: regular format (day-time), seconds (since 1-Jan-72), or time in seconds after midnight on the first day.

4. USER INTERFACE

The user interface can be divided into five parts: individual attribute control, saving and recalling of plot setups, plot setup editor, presentation of text and graphical data and session control.

The first part presents a group of buttons which help the user to define attributes of individual devices, grids or the current plot setup.

The second part presents a group of buttons that help the user to work with saving and recalling setups that are associated with the circular buffer and the saved buffer of files.

The third part allows the user to edit, in a single action, several items of a single plot setup.

Presentation control consists of buttons which are responsible for plotting, listing and exportation of data and has been previously discussed.

Session control consists of two buttons named "Stop" and "Pause". The first of them cancels a plot, and the second suspends/continues a plot. These buttons only have an effect while plotting is in progress.

4.1 INDIVIDUAL ATTRIBUTE CONTROL

There are several groups of buttons designed to control different attributes of devices, grids, and single plots. There are some other attributes that are entered in special fields. These are the device name and the initial and final values of the grid. These attributes are read from their fields just after the command is issued to start the plot.

There are some individual attribute buttons for each device or grid. They control the datalogger node of the device and the type of fitting function. These attributes stay resident on the background text window. A shortage of display space makes it impossible to have an individual attribute button for each grid and for each attribute. Thus some attributes are not displayed on the background window and are displayed only if the operator clicks on the corresponding button. These are "Trace", "Symbol", "Skip", and "Units". The buttons related to devices and grids bring up a menu including all devices or corresponding grids and display the current state of the attributes. These buttons can be used both to display and to change attributes.

The buttons "Trace" and "Symbol" control the trace type of the plot and the type of symbols plotted, respectively. Their features have already been discussed.

The "Skip" button allows the user to choose a number of points to drop from the data read between accepted points. This removal of points is done by the application. There is another option to specify the minimum time interval between points. This work would be done by the datalogging process which has the benefit over the skip method of reducing network traffic.

The "Units" button allows the user to set a different Y axis scale for each grid. For user convenience, all "invisible" grid attributes can be displayed (and changed) by using a hidden button attached to each grid.

The buttons "X divisions" and "Y divisions" control the numbers of X and Y axis grid lines (for all grids).

Options to calculate statistical values are controlled by the buttons "Integrate", "Average" and "StdDev". Clicking on each button brings up a menu asking whether plotted curves are channels, integrals or both. When integral values have been chosen, the time units for the integrated curve and integral value should be input as well as a scale factor to adjust the presentation of the integrated curve in the graphical window.

There are several ways to change the time interval (T1, T2). First, the user can simply type in a new date and time over the old one and interrupt. Clicking on the "T2 now" button will immediately enter the current date and time for T2. Clicking on the "Interval" button allows the user to specify a time interval between T2 and T1 and enters the resultant date and time for T1. The "Inc/Dec" toggle button allows selection of the direction of time intervals.

4.2 SAVING AND RECALLING SINGLE PLOT SETUPS

There is a one to one correspondence between datalogger nodes and files containing circular plot setup buffers and save and recall files. These files are opened automatically after clicking on the "Data source" button and selecting the datalogger node desired.

Each plot setup is temporarily saved in a circular page buffer having thirty records. Each such setup can contain more than 100 parameters describing up to nine devices. The user can go quickly back to a plot setup used a few plots ago by clicking several times on the "Previous" button. Similarly, the "Next" button will bring the user to the next setup in the circular buffer.

More important setups can be saved in a save and recall file. Its capacity is sixty records. The current setup can be saved by using the "Save" button and inputting a descriptive title. To recall a record from the file, the user just clicks on the "Recall" button and then on the title of the setup desired in the file menu.

4.3 PAGE EDITOR

To give more flexibility for making setups, the concept of Enabled and Disabled devices was introduced. It provides an option to temporarily exclude (disable) or include (enable) some devices and their attributes into the current plot setup. This can be done for multiple devices at once.

If it is desired to have many devices with some of the same parameter values (node, initial and final values, etc.), the "Copy" button can be useful to accomplish this. The "Edit" button provides search and replace capabilities for device names. This is useful for setting up a plot when some device names contain common substrings.

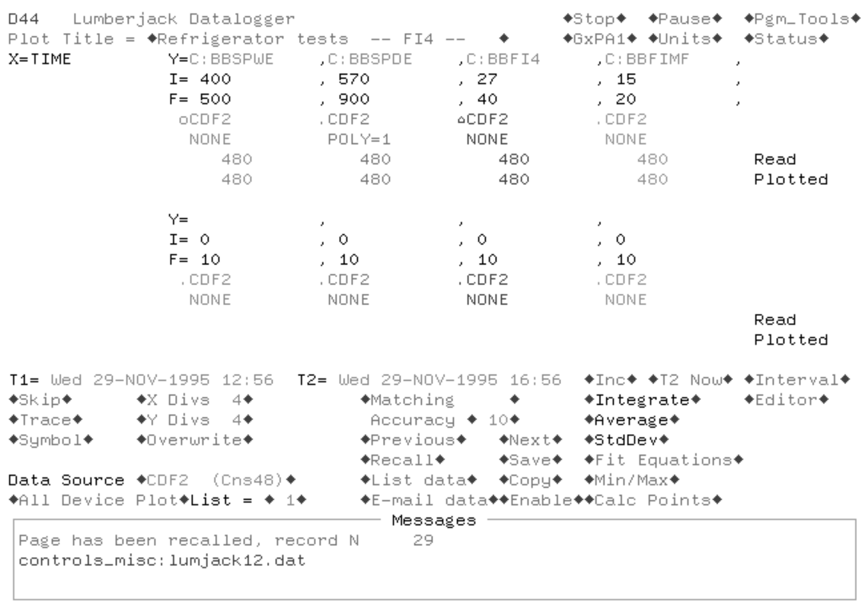

Fig. 2. Buttons and device attributes on the background text window.

5. SUMMARY

The program usage statistics show that this program enjoys great popularity among control room operators, accelerator physicists, engineers and administrators. This program covers most user needs for graphical data presentation, mathematical manipulation and data exportation, and its services continue to expand. A great deal of effort was expended to provide high speed data acquisition (thousands of points per second) for presentation. This was achieved by economizing the network traffic needed , modifying library routines to support the needs of the task, and optimizing the datalogging processes themselves. The next upgrades include the option to specify different time intervals for different devices and the option to display basic status values.