SCORE – A SAVE, COMPARE, AND RESTORE APPLICATION FOR SNAPSHOTING MACHINE SETTINGS

J. Galambos, J. Patton, T. Pelaia, ORNL, Oak Ridge TN, U.S.A.

Abstract

SCORE is an application used to save machine settings, compare live values with save values and restore saved values. Features include sorting by systems and subsystems, highlighting changed values and database storage of the information. The compare feature is useful for diagnosing problems, and the restore feature is useful in recovering good beam tune. Features of the application and performance experience with respect to database retrieval and live comparison will be discussed.

1. INTRODUCTION

Accelerators are inherently complicated devices consisting of multiple systems, each with many settable components. Total setpoints counts of many thousands of values are typical, and a single quantity set improperly can negatively impact beam operation. Efficiently managing the proper device settings is an important part of operations and beam study. A program called SCORE (Save Compare and Restore) has been developed at the Spallation Neutron Source for this purpose. At the earliest days of SNS commissioning it became obvious such a tool was direly needed. The EPICS BackUp and Restore Tool (BURT) [1] was evaluated but it lacked a compare feature, used a cumbersome file storage system and sometimes did not work for us. The beam study physicists and operators at SNS were queried for features they would like in a restore tool. As SNS was a new facility, the garnered requirements came from the combined experience of the multiple accelerator laboratories from which the SNS staff had come. SCORE was developed to satisfy these requirements and quickly became one of the most widely used application programs at SNS and the de-facto tool for machine setup restoration.

SCORE is a Java program written in the context of the XAL Framework [2]. Each set of saved data is stored as a collection in a database. Saved sets consist of multiple records consisting of signal names and values. When a set is selected, EPICS channel monitors are created for each signal. A Java GUI display offers the user a filterable view to compare the live machine values with the saved ones. The user may select arbitrary portions of the saved set for restoration to the machine. These features are described in more detail in the following sections.

2. DATA STRUCTURES

Internal Data Structure Within SCORE

User requirements included that the device parameters to be saved should be comparable to live values by a human before the restore is done, they should be filterable by major accelerator system (e.g. Ring, transfer line1, linac, ect.), and also filterable by sub-system (e.g. magnet, RF, Timing, etc.). Finally both setpoints and readbacks of comparable quantities should be displayed on the same row (e.g. magnet power supply current setpoint and readback). This is useful for example, since restoring a magnet power supply current setpoint to a saved value does not guarantee that it will return to the same value, esp. if the power supply is tripped either at the present time or was tripped when the original snapshot was taken.

Not all rows have both setpoint and readback values. For example some devices do not provide readback values for settable quantities, sometimes readback quantities are stored with no single direct settable signal for documentation of conditions (e.g. vacuum level), and sometimes important settable quantities need special permissions for change – yet are useful to compare to a good tune setup case. In the latter case, these special settable quantities are included only in the readback column, where they can be compared, yet are safe from any restore possibility.

Database Structures

A collection of the data structures described above represents a "Group" of data that can be utilized by SCORE. Example sets used at SNS are 1) the main magnet RF, timing, etc. settings used to directly setup the machine, 2) detailed diagnostic device calibration settings, and 3) detailed RF feedback and feed forward settings used by the RF group. This allows different groups or interested parties to create their own groups of quantities to check. A database table is used to define these groups.

Another database table keeps track of group snapshots that have been taken. This table is simple – it just has a primary key to the group discussed above, a timestamp when the snapshot was taken, a flag indicating whether this is the "golden" snapshot for this group and a text comment describing the group. It does not contain any actual values. This particular table is useful for browsing available snapshots over a prescribed time range and previewing the comments.

Of course when a snapshot is taken, the data must be saved and a third database table is used for this purpose. Each record in this table includes the timestamp, setpoint name, setpoint value, readback name and readback value. For historical reasons, all data is stored as strings in the database and converted to type Double in the Java program. An implicit assumption is that all data stored is real – if other types need to be stored, this table must be extended. This table has had many records committed to

date. Each "group" consists of order ~1000 rows – consider the number of major settable quantities in an accelerator. Also during beam tuning, many of these snapshots are taken during each day. None-the-less the query to retrieve a selected group snapshot is typically seconds.

A major advantage of database storage of the saved information is that there is no need to maintain archives of text files. The initial version of SCORE used a file storage method in a "directory hierarchy" and was much more cumbersome and lacked features discussed below.

3. JAVA PROGRAM

Retrieving a Snapshot

When the program starts, an "open" tab is active (or can be selected at anytime). This view is shown in Figure 1. The different snapshot group types are available in a pull-down list selection. Once the group type is selected a specific snapshot may be selected. One option is to simply select the "golden" save set for this group. Otherwise the user can browse for the desired snapshot. An adjustable time range is provided for refinement of the search. For a given time-span, all available group snapshots in this range are retrieved, and the timestamp and user comment of these snapshots are displayed. The user comment is especially useful and often the primary means of quickly remembering which save-set to use. Use of a database greatly simplifies implementation of this feature. The original SCORE deployment used text files for information storage, and this feature was not readily possible.

Once a group snapshot is selected for retrieval, several actions are taken. A list of all systems and subsystems included in the snapshot is compiled for filtering purposes. The user can select which systems to display and information on each system is displayed on a separate tab. Within each tab, the display of information is ordered by subsystem. Filtering is also available for signal names (regular-expression). Also, monitors are opened for each signal (both readback and setpoints) to facilitate live value comparison, regardless of what fraction of the group has been selected for the initial display. The display filtering can be changed on-the-fly.

Comparing the Live and Saved Data

An example of a live display is shown in Figure 2. Each row of displayed information includes the setpoint signal name, setpoint saved value, setpoint live value, associated readback signal name, readback saved value and readback live value. The user can navigate different systems, subsystems etc. to look for values that are either set differently or are operating differently than when the group snapshot was taken. Of course visually looking for differences between many numbers in a large display is difficult for humans. There is a feature to highlight the live value displays in red if they differ by a prescribed fraction more than the corresponding saved value. Many difficulties in restoring a good tune setup are quickly resolved by browsing systems "looking for red".

The signal monitors may update at different rates, and if the display were always updated at each monitor change, the application would be slow, and the display difficult to read. We have implemented a timed display update every two seconds to avoid this problem.

Taking a Snapshot

The first (but optional) step in creating a new snapshot is creating a text description of the snapshot you are about to take. Pressing a button called "Snap n Save" results in a new snapshot being created. The live values from the monitors are copied to the "saved value" column and new database entries (as described in Section 2) are created. If a live value is undefined (e.g. the device is temporarily unplugged or communications not available so that no monitor can be established), the saved value is stored as NaN. The snapshot action is simple and fast. Snapshots are taken often during beam tuning activities when incremental improvements are being made amid steps backwards.

Restoring Values

Restoring values requires slightly more user actions than the snapshot action since this action affects the machine. First the user selects the rows for which restoration is desired, and then clicks a "Restore" button. A confirmation dialog prompts the user to ensure this action is desired, and if so the saved setpoint values in the selected rows are sent to the machine. At the same time, an automatic notice is posted in our electronic log book indicating which quantities have been restored. If any of the saved values have a NaN value, the restoration of these particular signals is suppressed.

4. EXPERIENCE

SCORE has proved to be a valuable aid in restoring machine setup and a useful aid in beam tuning. It fulfills a need for quickly understanding what has changed from a good tune. The compare feature is especially critical in this respect. It allows one to evaluate the present machine state with another saved state without resorting to blind restoration of values until things get better. Use of SCORE helped the rapid beam commissioning and power ramp-up progress over the last several years. Also the use of a database for the storing stored information has been positive. It offers features that are not possible with file storage and alleviates file maintenance.

At SNS beam availability is very important, and rapid restoration of production beam is crucial. SNS operators are able to restore the entire machine rather quickly and reliably, with help from with this tool. E.g. well tuned high power neutron production has been restored by operators alone in a few hours after a major power failure, with the help of SCORE.

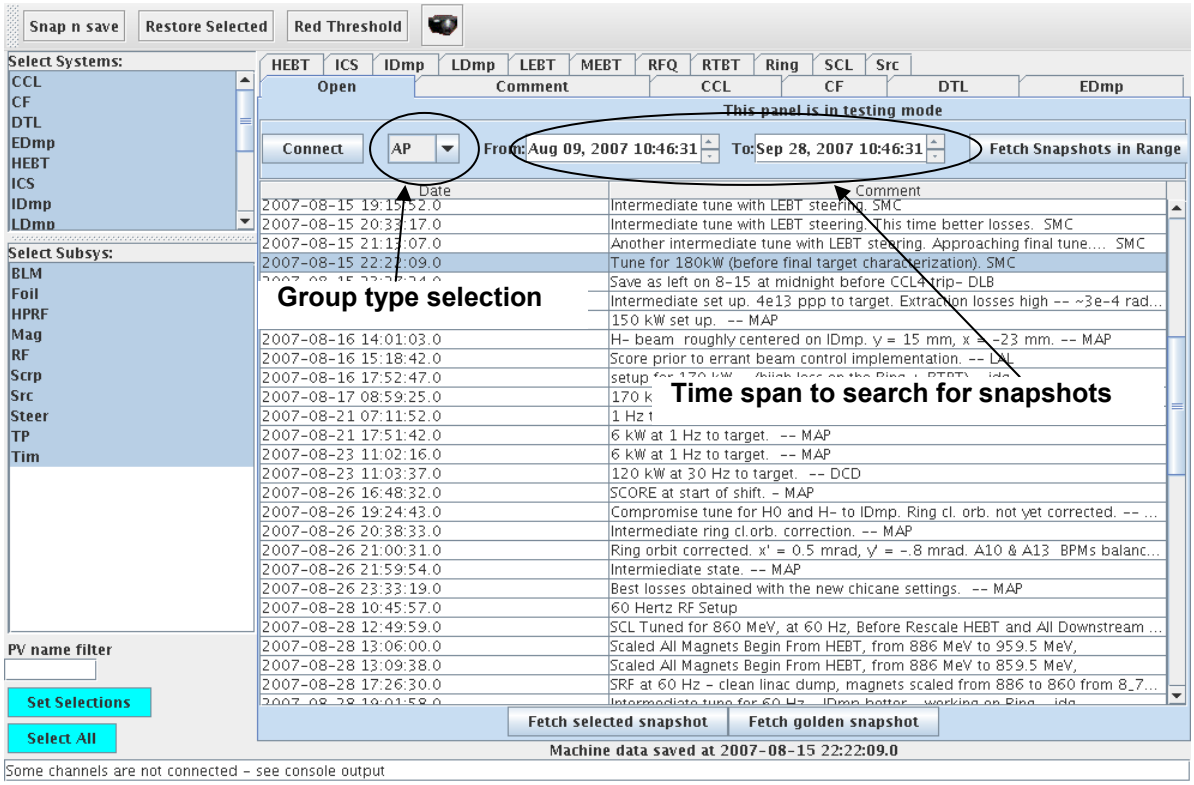

Figure 1: The user interface for viewing and selecting group snapshots. Each group snapshot has a timestamp and user comment for identification.

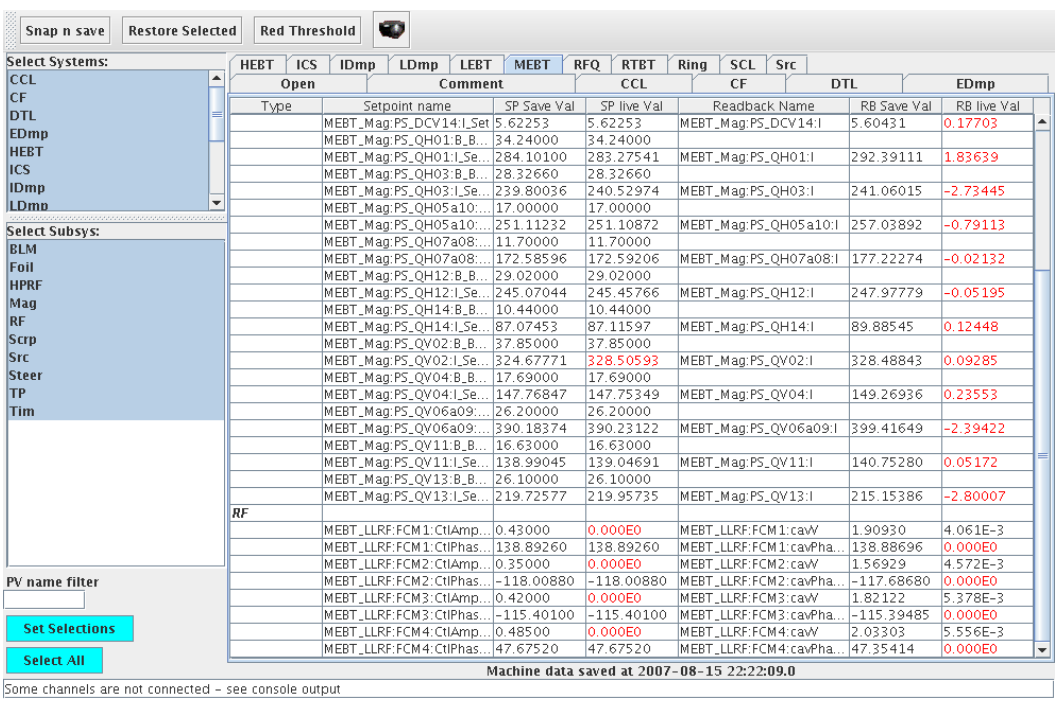

Figure 2: Example display of a selected system (MEBT) from a group snapshot. Live values that differ by more than a prescribed amount from the saved value are displayed in red.

REFERENCES

[1] BURT: Back Up and Restore Tool, http://www.aps.anl.gov/epics/extensions/burt/index.php .

[2] Thomas Pelaia II, Chris Allen, Sarah Cousineau, John Galambos, Jeff Holmes, Andrei Shishlo, Yan Zhang, Alexander Zhukov, P. Chu, "XAL Status", these proceedings.