STAR ONLINE CONTROL SYSTEMS AND EXPERIMENT STATUS

W. Betts[*], M. DePhillips, V. Fine, L. Hajdu, J. Landgraf, J. Lauret, Brookhaven National Laboratory, Upton, NY 11973, USA

R. J. Porter, Univ. of Washington, Nuclear Physics Laboratory, Seattle, WA 98195, USA

M. Cherney, W. Waggoner, Creighton University, Omaha, NE 68178, USA

Abstract

The STAR (Solenoidal Tracker at RHIC [1]) experiment's operations consist of dozens of mission-critical sub-systems, each of them recording their own information but also having many interdependencies with each other. Disseminating such information to local or remote experts and systems often must be done in real or near-real time, providing such diverse information as the current collider and detector conditions, physics metadata or even various administrative information intertwined with shift responsibilities and duties or detailed operations reports. Furthermore, current and past states of facilities and detector systems are necessary for current running and planning, performance evaluations, quality assurance and ultimately, for the future data analysis. This implies that most of the information must therefore be maintained and archived over a long term for use as needed.

To achieve this daunting task, the STAR collaboration uses a variety of relatively inexpensive LAMP-based (Linux, Apache, MySQL and Perl/PHP/Python) and Java-based solutions and systems to collect and disseminate as much of this information as possible to STAR collaboration members, both in the local Control Room and worldwide. The strategy employed, developed around a "*three tier layer*" architecture design wherever possible, is able to keep data collection and analysis running efficiently. Several such tools will be described following along an intuitive path of a 'typical' user and collaborator as he interacts with these tools to learn about the status and control of the experimental systems in the course of acquiring physics data.

PREPARING FOR A SHIFT

An essential step toward a stable and efficient run period is the assignment of tasks according to best qualifications of the shift crew. Before coming for an experiment, a STAR collaborator would therefore refer to a set of documentation pertaining to the shift duties and required training. The STAR online website [2] (hosted by Apache on a Linux server) is therefore the starting point for the collaborator who wants to get involved in operations. He may directly refer to the shift sign-up tool and interface, providing a convenient way to view the the available (open) shift slots, the qualifications required for each shift and the other shift members for the period under consideration (cf. Fig 1). The shift sign-up is PHP-driven business logic that gathers part of its information
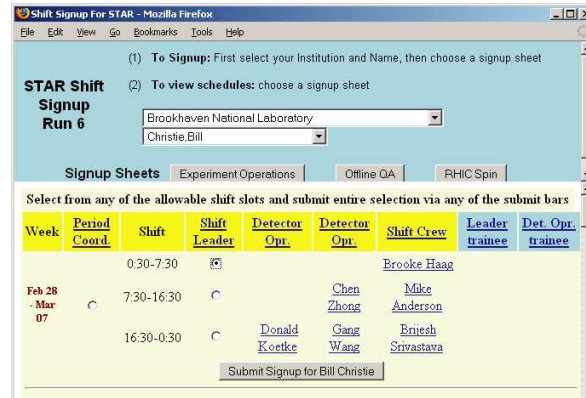

Figure 1: STAR's online shift sign-up. In this example, our collaborator is actively prevented from signing up from several open slots because he does not meet the prerequisites, and training was not available during the week shown.

from a collaboration-wide phone book (or roster), inventorying all collaborators and their hosting institutions. From there on, each shift and assignment is recorded and maintained by a shift-signup independent MySQL database, isolated from any other dependencies (a roll-out update of those records from the original collaboration roster will be done once every other day or on demand), providing a self-contained and secure way to hold the information. The purpose of the shift sign-up database is not only to keep the information of who has signed-up when, but also includes information detailing who is qualified for each shift position. The system therefore can prevent unqualified individuals from signing up for inappropriate shifts and does not allow an individual to sign-up for multiple shift positions within a single shift period. Keeping track of all shift duties within the shift-signup database also allows easy accounting of shift credits, a number of which are required of participating institutions for continued membership in the collaboration.

As his shift assignment nears, our collaborator can keep up with the more minute details of his upcoming experience and get a feel for what to expect by monitoring a set of documents maintained by the shift leaders that are posted online. Shift leaders can post detailed procedural guidelines and checklists for example and change them as conditions dictate.

---

*wbetts@bnl.gov

THE CONTROL ROOM

In the Control Room, naturally one finds the most intricate details of ongoing operations. The many detectors and assorted sub-systems all have their own control and readout systems. Helping to unify what could easily be a tremendous mishmash of systems, most detectors and much of the collider status information are monitored or controlled by a set of EPICS-based tools maintained collectively under the banner of the "*Slow Controls*" sub-system [3]. EPICS, hosted at Argonne National Laboratory [4], is a collaboratively developed set of Open Source software used to create distributed real-time control systems. Figure 2 shows an example of EPICS in action, displaying the top level GUI menu for the STAR TPC control system. Slow Controls additionally archives selected data from the various systems, and has mechanisms for graphical depiction of archived data. Shift crew members routinely spend time looking at and interacting with Slow Control screens such as in Fig. 2. However, once a data-taking run is underway, the detector states should be static and attention shifts to QA monitoring.

A variety of online monitoring systems provide feedback on the quality of data being collected. Of these, the most appealing to look at are typically the visualization tools that operate in near-real time, such as the wall-mounted plasma screens shown in Fig. 3 or the closely-related online event display, an example of which is shown in Fig. 4. Of course, such visual tools, powerful as they can be, are only a small part of the online QA process. A number of diagnostic tests and histograms are also made available, but typically require more in-depth knowledge to interpret.
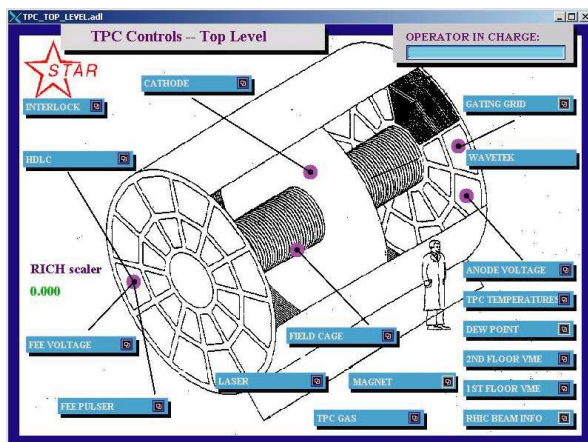

Figure 2: The EPICS top-level menu for the STAR TPC. Most of the detectors have similar EPICS GUI controls.

Figure 3: STAR Control Room plasma displays with near-real-time event displays, rendered with OpenGL.
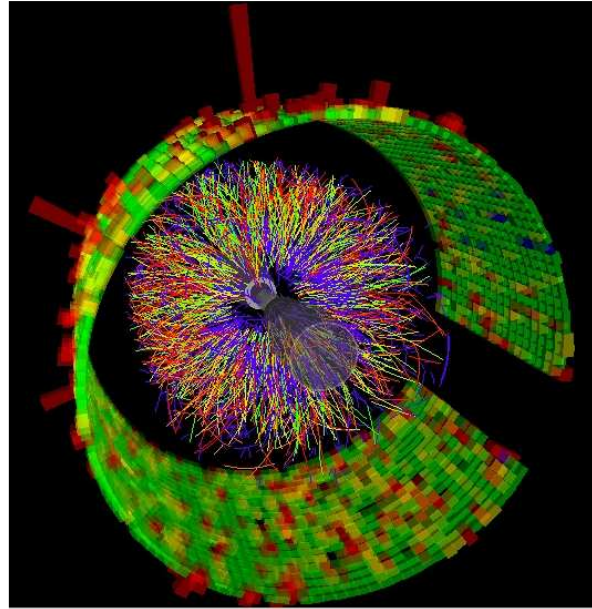

Figure 4: A sample image from the online event display. One can often quickly assess the health and status of various detectors from such images. For instance, noisy calorimeter sections or detectors failing to collect meaningful data are often quite obvious, such as the missing section of the barrel on the right.

Such event displays and online QA are made possible by continuously sampling real events and filling an event pool. Sampling occurs at a rate of a few Hertz, while the main event rate for tracking detectors is ~60-90Hz. From the event pool, various online QA analyses are done (which can include quick passes through particle tracking algorithms) allowing runs to be carefully but quickly checked for the apparent level of quality -- the shift crew can make the decision of marking the run good or bad as the data flows from the detector to the QA system. Event images such as shown above also make great educational and public relation tools for visitors to the Control Room and beyond -- for example, during running in 2005, for a period of several months, an artist arranged a pixelized version of live events displayed on the side of a building in Munich, Germany, live video images of which were available on the Web [5].

Electronic Shift Log

Configuring the detectors and verifying the quality of data are critical shift crew functions, but only half of the battle. Of utmost importance is recording what has transpired. While DAQ and Slow Controls record lots of information automatically, human decisions, observations and actions need to be recorded as well. Information could also come from remote collaborators and the traditional means of communicating such as email or phone were revealed to be inadequate, difficult to keep for future reference, and obviously not scalable as the experiment grows with the number of sub-systems. For those reasons, STAR has developed an online ShiftLog that eliminates the need for reams of paper logbooks. Fig. 5 shows sample ShiftLog client views. Two flavours of the client exists: a Java client GUI application connected to a backend database provides browsing and editing functions, while a browsable and searchable logbook is also served by our online web server. Log entries can be associated with individual sub-systems, effectively creating multiple logbooks in one.
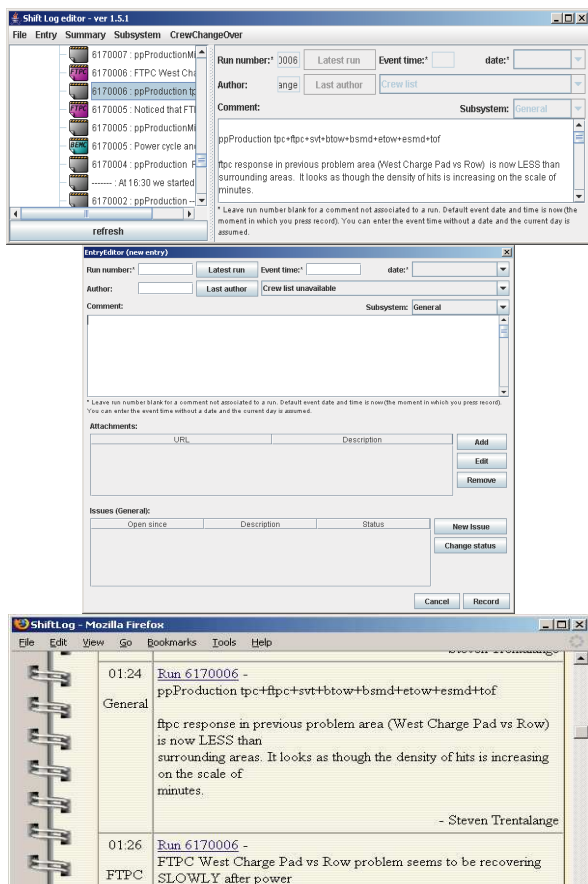

Figure 5: STAR's electronic logbook. At top is the Java client shown browsing existing entries. The middle image shows the input form for new entries. At bottom is the online browser showing a sample entry.

Using this tool, even remote sub-system experts can follow and respond to items, report problems detected via distributed monitoring tools and feed back information to the online crew without use of phone or emails. Additionally, this method of communication is also used by the offline QA system, adding reports and summary to the ShiftLog and providing delayed but valuable information to the shift crew without the need to consult mailing lists or other private communication. In other words, a single source of information provides complete coverage of all human and automated information gathering and presentation, so critical to the development of a shift. Other features of the system include automatic creation of links to associated run information, creation of Shift Summary reports, the inclusion of attachments, checklists for Shift Changeovers and even keeping track of crew members by checking the shift-sign up database. Such features help to insure smooth and informed transitions from crew to crew and maintain consistency of activities amongst the various sub-systems.

RUN SUMMARY AND METADATA

STAR's DAQ system automatically records copious amounts of information about the individual runs. Two frequently consulted tools for accessing that metadata in STAR are the intimately related Run Log Browser and Trigger Versioning site. The online PHP-based Run Log Browser (see Fig. 6) provides access to the metadata for each run, such as detector states, number of events by trigger, collider status, and file names and storage location. Additionally, the browser interface can create summary information covering multiple runs, allowing at a glance to determine for instance how many events
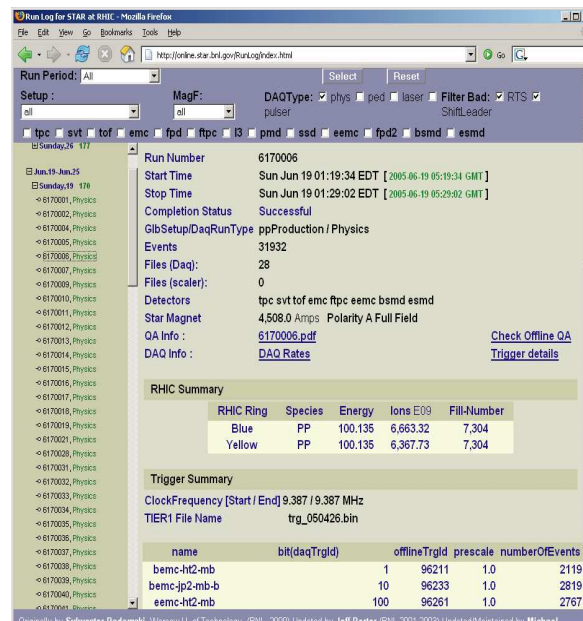

Figure 6: The STAR Run Log Browser is shown displaying summary information for a single physics data run.

were recorded in a given day or week. The Run Log Browser also provides cross references to a number of other tools, such as related ShiftLog entries and the automatically generated online and offline QA reports and information. Complementing the Run Log Browser, the Trigger Version page provides the intricate details about the many triggers used by STAR. Included are definitions of the event triggers used, trigger information by run (and vice versa), trigger rates and cumulative totals in tabular and graphical form (a sample of which is shown in Fig. 7) and estimates of total sampled luminosity.

While by no means complete, we covered the most important and commonly used tools for participating in and understanding STAR operations, from getting involved in data collection to knowing what is available for analysis. Having been on shift collecting data and documenting his activities, our typical collaborator will certainly want to spend some time analysing the data. The tools described herein will allow him to know how much data there is of various types that interest him and the details of its collection. He might even make suggestions for the focus of future data taking based on what he finds. From the available metadata and summaries, the efficiency of operations and projections into the future can also be extrapolated, aiding the complex planning that must be done to smoothly run so many interleaved sub-systems within a large collaboration existing at such a multi-faceted facility as RHIC.

REFERENCES

[1] http://www.star.bnl.gov/
[2] http://online.star.bnl.gov/
[3] D. Reichhold et al., Nucl. Instrum. Meth. A 499, 792 (2003)
[4] http://www.aps.anl.gov/epics/
[5] "*I see what I see not*" by Tim Otto Roth, http://www.photograms.net/kunstfassade/tor/bnl_e.html
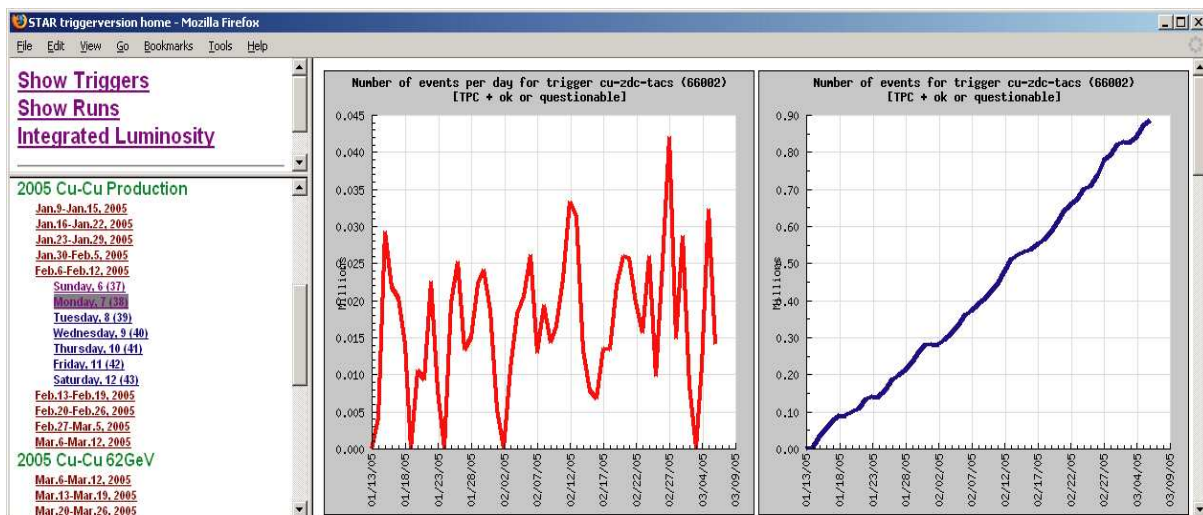
Figure 7: STAR's Trigger Version site, showing examples of PHP-generated graphs of a particular trigger type collected over the course of about two months. The blue graph on the right is an integration over the data shown in red during that time period.