Tools for Trigger Rate Monitoring at CMS

Geoffrey N Smith[1], Charles N Mueller[1], Andrew S Wightman[1] on behalf of CMS

[1] University of Notre Dame, Notre Dame, IN (US)

E-mail: `gsmith15@nd.edu`, `charles.mueller@cern.ch`, `andrew.steven.wightman@cern.ch`

Abstract. In 2017, the LHC delivered an instantaneous luminosity of roughly $2.0 \times 10^{34} \mathrm{cm}^{-2}\mathrm{s}^{-1}$ to the Compact Muon Solenoid (CMS) experiment, with about 60 simultaneous proton-proton collisions ($<\mu>$) per event. In these challenging conditions, it is important to be able to intelligently monitor the rate at which data are being collected (the trigger rate). It is not enough to simply look at the trigger rate; it is equally important to compare the trigger rate with expectations. We present a set of software tools that have been developed to accomplish this. The tools include a real-time component - a script that monitors the rates of individual triggers during data-taking, and activates an alarm if rates deviate significantly from expectation. Fits are made to previously collected data and extrapolated to higher $<\mu>$. The behavior of triggers as a function of $<\mu>$ is then monitored as data are collected - plots are automatically produced on an hourly basis and uploaded to a web area for inspection. This same set of tools can also be used offline in data certification, as well as in more complex offline analysis of trigger behavior.

1. Introduction
The trigger system is extremely important to data acquisition at the Compact Muon Solenoid (CMS) [1] experiment at CERN. Roughly 500 separate algorithms combine to filter the approximately 40 MHz rate of collisions down to 1 kHz of data (corresponding to 1 GB/s) collected for offline analysis. Without an effective trigger system, physics at CMS would not be possible.

As the LHC pushes to higher beam intensities, CMS has to be ready to respond to emergencies if the trigger rates go out of expected range. The trigger rate is very sensitive to all aspects of the detector and how they operate, so it often provides the first indication of a detector problem. Therefore, it is very important to be able to intelligently monitor, characterize and visualize trends in trigger rates.

The trigger Field Operation Group at CMS has developed a set of software tools to accomplish this task, consisting of a suite of python modules and scripts that perform a variety of functions. Fits are made to the trigger rates in previous runs using linear and non-linear regression. These fits are then compared to the instantaneous trigger rate as data are being collected, in order to spot small (unexpected) deviations in rate. As well as this real-time component, the software provides a variety of additional features that are used in offline analysis.

2. Fits
First, fits are made to the trigger rate in previous runs as a function of the average number of collisions in an LHC bunch crossing, $<\mu>$ (also known as average in-time pile-up). Runs to be

fit are selected from a list of known good runs, and the fits are performed using the ROOT [2] software.

The rate information is obtained from a central SQL database at CMS. Before fitting, the raw trigger rate is first corrected for deadtime, Level-1 (L1) and High Level Trigger (HLT) prescales, as well as the number of colliding bunches in the LHC. Performing this correction facilitates comparisons and extrapolations between runs with different conditions, and allows a smooth function to be fit between runs. For each fit, several fit functions are attempted, and the final function is selected based on a χ^2 minimization.

Ideally, trigger rates should depend linearly on instantaneous luminosity. However, some have nonlinear behavior, mainly due to pile-up effects. Candidates for multiple object triggers may be found in independent collisions, and triggers that rely on measuring calorimeter deposits (such as jets and jet sums) increase in rate as the number of contributions from independent collisions increases. After exploring several options, we found that the behavior of most triggers is well-described by either linear, quadratic, or exponential (sinh) functions. The triggers with linear and quadratic behavior are primarily those that trigger on single and double objects, respectively. The triggers with the highest sensitivity to pile-up are the jet sum triggers, and they are described best by exponential functions. Examples of fits are shown in Figure 2.

3. Online Monitoring

3.1. Real-time Component

Using fits from previous runs, the trigger rates can be monitored during data-taking by comparing the instantaneous rate to the rate predicted by the fit. A list of approximately 20 L1 and HLT triggers is used for real-time online monitoring in the CMS control room. This list is selected such that all CMS subdetectors are monitored and all physics objects are represented. An automated script running continuously during data-taking checks current trigger rates against the prediction from the fit for each trigger in the list. Rate information is displayed in the trigger shifter terminal. The rate information

Figure 1. Flowchart for the real-time component of the rate-monitoring software.

is updated every two minutes— updating more frequently can cause errors as there is a lag associated with the time it takes to propagate new rate information to the database (this is an aspect of the software we are working to improve). When the rate of a given trigger is at least five times higher than the uncertainty on the fit for a significant period of time, the script prints a warning, and an email summarizing the problematic triggers is automatically sent to on-call experts. In order to decrease false alarms due to noise, the warning and email are only activated when three consecutive queries of the database indicate a higher than expected rate. The script also activates audible alarms in the control room in the event that rates exceed more extreme, fixed thresholds; in this case, warnings are tailored to help diagnose problems with individual subsystems. For example, if the rate of a certain L1 e/gamma trigger is higher than 25kHz, the warning "critical L1 e/gamma trigger rate: please check that ECAL and calorimetric triggers are behaving correctly" is given.

3.2. Summary Plots

In addition to the real-time component, the software produces plots that compare the trigger rate versus $<\mu>$ to the expected rate for each fill. These plots have been integrated into the

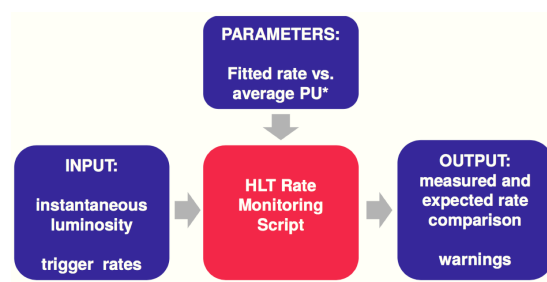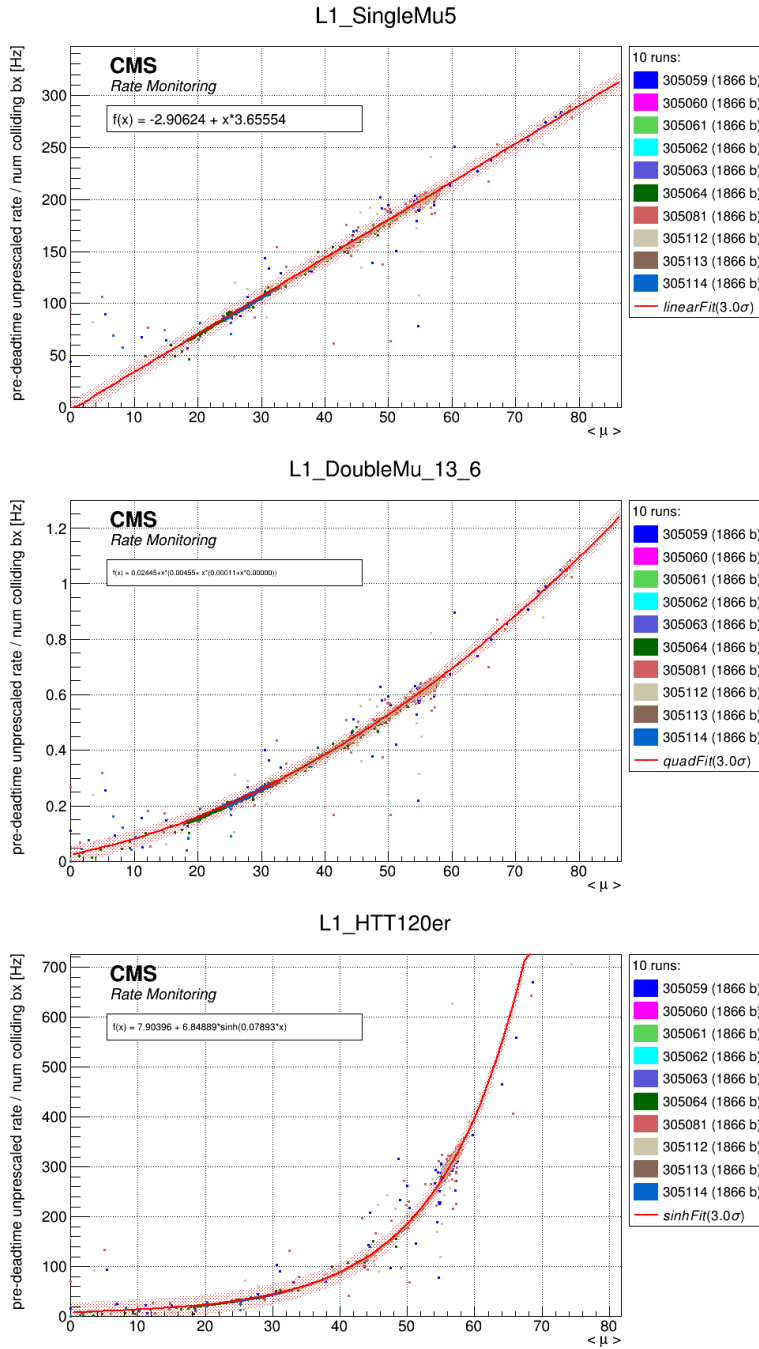**Figure 2.** Examples of adjusted trigger rate versus $<\mu>$. The points correspond to the per-lumisection rate, color-coded for different runs, and the pink band is the 3σ error band from the fit. From top to bottom: examples of triggers whose behavior as a function of $<\mu>$ is best described by linear, quadratic and hyperbolic sine functions, respectively.

central CMS Web Based Monitoring service (WBM); a dedicated page is linked from each Fill Report page on WBM, containing plots with rates for that fill (Figure 3). The page contains rate versus $<\mu>$ plots for all the triggers in the monitored trigger list. Plots for all HLT and L1 triggers, as well as stream and dataset plots, are available via links from this main page.
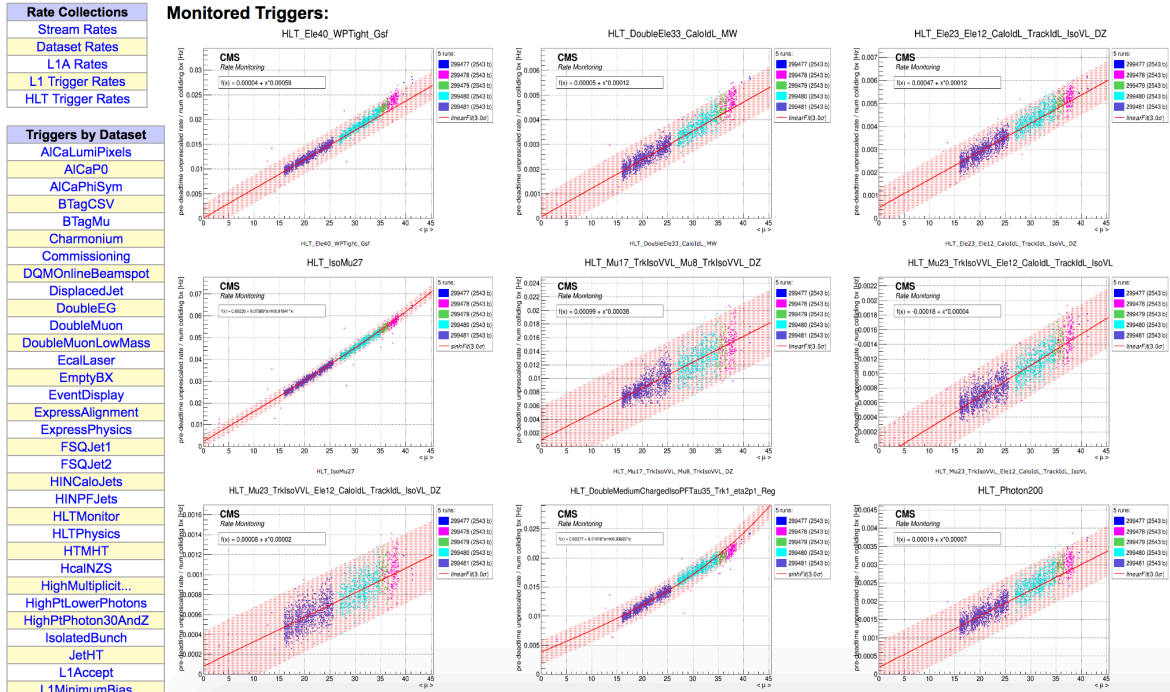
Figure 3. Screenshot of WBM page produced by the rate-monitoring software, showing plots of rate versus $<\mu>$ for LHC fill 5976. Selected triggers are shown on the main page, with links on the left-hand column to further pages containing plots for every L1 and HLT trigger in the trigger menu, as well as dataset and stream rates.

Additional links directly to the plots of the individual triggers are available from the HLT and L1 summary pages. The plots are produced by a cron job run on a dedicated machine that updates the plots on an hourly basis (for the current fill only). Work is ongoing to develop a better implementation which would allow the plots to be updated in real time.

4. Offline Use
The software is also used for offline data certification. For a given list of runs and list of triggers, rate-versus-lumisection plots can be produced for each run and trigger (one lumisection equals 2^{18} orbits of the LHC beam, or about 23 seconds). An overlay of the rate prediction from the fit enables easy comparison by offline validators (Figure 4), and a text summary is also produced highlighting the runs and lumisections where triggers deviated significantly from expectation. The software has a modular organization, and can be extended to perform additional functions by other collaborators.

5. Summary
We have described a set of software tools currently in use at CMS to monitor trigger rates. This software is used for real-time monitoring of trigger rates by shifters and on-call experts, as well as offline for data certification and other studies. In particular, the trigger shifter script together with the plots on WBM have allowed fast identification and diagnosis of detector issues. The sofware has proven to be a critical tool to the successful operation of the trigger and to successful data-taking at CMS.

For further information about the rate-monitoring software, see the Github repository here [3].
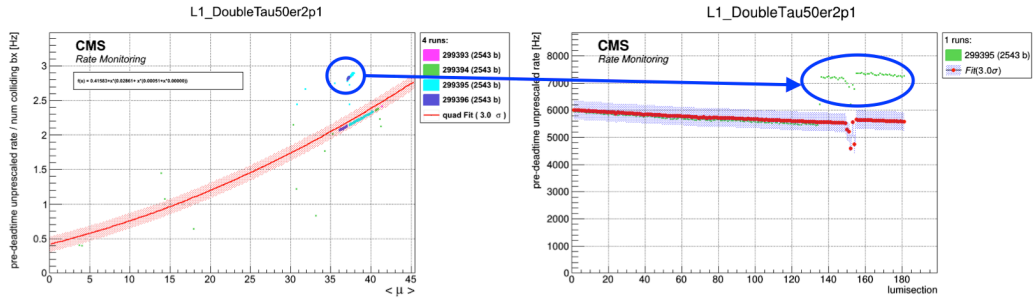
Figure 4. Examples of plots used for run certification. Runs where the trigger rate deviated significantly from the fit are first identified in rate-versus-$<\mu>$ plots (left). Rate-versus-lumisection plots are then used to identify the affected lumisections (right).

6. References

[1] CMS Collaboration 2008 *The CMS experiment at the CERN LHC* JINST 3 S08004
[2] Rene Brun and Fons Rademakers, *ROOT - An Object Oriented Data Analysis Framework*, Proceedings AIHENP'96 Workshop, Lausanne, Sep. 1996, Nucl. Inst. & Meth. in Phys. Res. A 389 (1997) 81-86. See also http://root.cern.ch/.
[3] A. Wightman et. al, *Repository with various tools to monitor HLT and L1 rates*. https://github.com/cms-tsg-fog/RateMon.