

THE LHC BEAM LOSS MONITORING SYSTEM COMMISSIONING FOR 2010

C. Zamantzas, B. Dehning, C. Chery, E. Effinger, J. Emery, S. Grishin, C. F. Hajdu, E. B. Holzer, S. Jackson, C. Kurfuerst, A. Marsili, A. Nordt, M. Sapinski, R. Tissier, G. G. Venturini, CERN, Geneva, Switzerland.

Abstract

The LHC Beam Loss Monitoring (BLM) system is one of the most complex instrumentation systems deployed in the LHC. In addition to protecting the collider, the system also needs to provide a means of diagnosing machine faults and deliver feedback of the losses to the control room as well as to several systems for their setup and analysis. It has to transmit and process signals from approximately 4'000 monitors, and has nearly 3 million configurable parameters. This paper will discuss its performance and ability to provide the expected measurements, the problems encountered and necessary improvements, the adequacy of related software and databases, and in general its readiness and suitability for 3.5 TeV operation.

INTRODUCTION

Based on the experience gained during the 2009 LHC operation [1] several corrections and improvements have been proposed to augment the functionality of the BLM system.

In addition, since the strategy for this year will require beams with higher energy and destruction potential, the full protection functionality and safeguards, expected from the BLM system, will be deployed and thus enforced by the system itself.

ACQUISITION

The acquisition part of the system resides in the least accessible part of the installation and is required to withstand a certain amount of radiation. Any degradation will have a direct impact on the minimum detectable signal, i.e. on the sensitivity of the particle flux measurements the system can achieve in total.

Offset Level

The offset current, which is a product of a protection mechanism to avoid lockups due to noise and radiation deposited in the electronics, is inversely proportional to the dynamic range that the system can achieve. For this reason, a regular investigation is being followed to discover the detectors and channels that exhibit levels exceeding tolerance.

On figures 1 and 2, an example of the mean offset level is shown in Gy/sec as measured by the 1.3 s (RS09) integration time window of each monitor in the L6 tunnel installation part. The average offset level is taken from a one hour run on the data recorded on the 18 December 2009. One can see that the offset level is less constant in the LSS and DS regions than in the ARC region.




Figure 1: Example of the mean offset level in Gy/sec from RS09 of each monitor in L6 as recorded on the 18/12/2009.

On figure 2, a similar example is shown for the evaluation of the mean offset level value - this time for installation point R3. The smaller plot shows the offset vs. time (i.e. 1 hour) for one monitor with a high offset level. It is visible that several sets of monitors exhibit a level higher than the expected due to the particularities of this region.

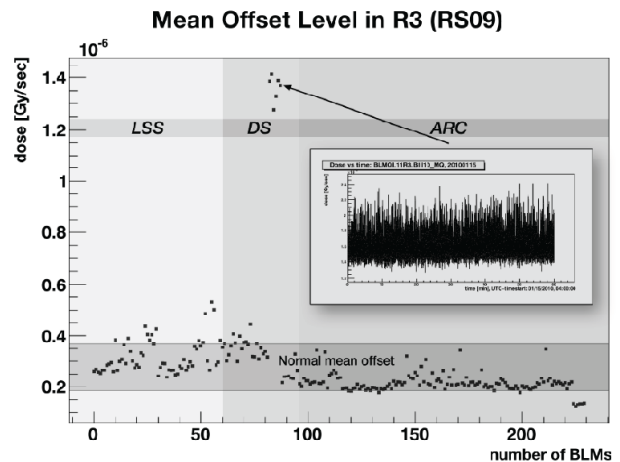


Figure 2: Example of the mean offset level in Gy/sec from RS09 of each monitor in R3 as recorded on the 18/12/2009.

On figure 3, the evolution of the mean offset level of all monitors in R3 is shown. In this example, the mean offset level was determined for four equally distributed time periods each day for each monitor during 14 consecutive days. The data used was recorded during the non-operational period between the 18/12/2009 and 1/1/2010.

One can see that fluctuations over time are still higher in the LSS and DS than in ARC regions, but the drift over time is in the acceptable region at the moment.

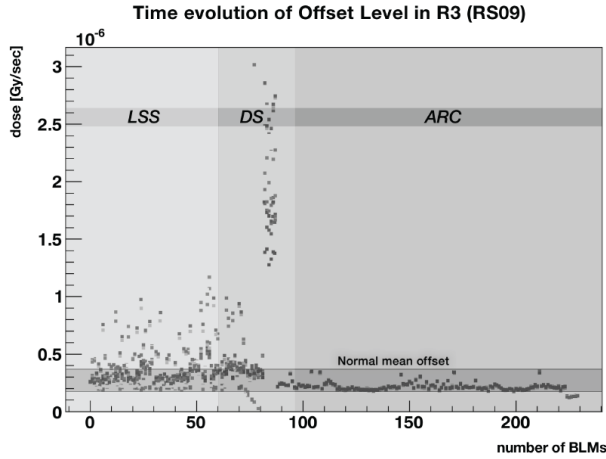


Figure 3: Example of the evolution of the mean offset level in Gy/sec from RS09 for each monitor in R3 as recorded during a period of two weeks.

Finally, in order to provide a complete overview of the installation, table 1 summarises the state of the offset level in all installed ionisation chamber monitors. All monitor channels outside the 10 to 80 pA range, i.e. the acceptance region set for this year's operation, need to be repaired.

Table 1: Mean offset values for all ionisation chamber monitors. Calculated using 1 hour of 1 Hz data recorded on 18/12/2009.

Offset [pA]	% of total	Comments
0	0.28%	no data - faulty cards/connections
1-9	6.91%	too low - problem has been identified
10-19	77.71%	very good
20-29	6.83%	good
30-80	3.28%	high - DAC reset needed
>80	5.00%	problematic

Noise Level

At the moment, the main contributor in the change of the offset level of a channel is the noise introduced in the acquisition input. The charge balance integrator used to construct the Current to Frequency Converter (CFC) [2] can enter a locked-up state if the current in the input flows in the opposite direction. For this reason, a protection circuit injects a small but constant current and monitors the output of the CFC. If it detects that in the last 20 seconds the output is lower than the expected it increases the injected current by steps of one pA until the CFC exits the locked-up state.

In general, there is less noise in the ARC region than in LSS and DS regions. This is mainly due to the cable

length and to the quality of the cabling. The effect can be seen graphically by comparing figures 4 and 5. They show examples of the frequency distribution of the noise on two channels connected using significantly different cable lengths. It should also be noted that the noise level of the secondary emission monitors is 5 orders of magnitude higher than for the ionisation chamber monitors.




Figure 4: Plot of a single channel frequency distribution of the noise on a short cable as seen by the 40 μ s (RS01) integration time.




Figure 5: Plot of a single channel frequency distribution of the noise on a long cable as seen by the 40 μ s (RS01) integration time.

Table 2: Noise in the 40 μ s (RS01) integration time for all ionisation chamber monitors. Maximum values from 9 hours of 1 Hz data recorded on 13/01/2010.

Noise [BITS]	# of ICs	% of total	Comments
0	18	0.5%	no data were available
1-30	1083	30.1%	very good
30-100	1646	45.7%	Good
100-200	600	16.7%	Ok
200-300	201	5.6%	Candidates for problematic channels
> 300	54	1.5%	critical noise

In order to avoid false triggers it has been decided for last year's operation to mask all channels that exhibit a noise level higher than 0,027 Gy/s (or 300 bits) at the RS01 integration time. Table 2 summarises the current status of the noise level in all ionisation chambers installed and shows the number of the monitors that fall outside this acceptance region and thus need to be corrected for this year's operation.

On figures 6 and 7, one can see the frequency distribution of the maximum noise values for all 3602 installed ionisation chambers and for the 302 secondary emission monitors, respectively, in Gy/sec, as recorded by the 40 μ s (RS01) integration time window.




Figure 6: Maximum noise value frequency distribution for the Ionisation Chambers (3602)




Figure 7: Maximum noise value frequency distribution for the Secondary Emission Monitors (302)

In an effort to remedy the issue, several cables have already been exchanged, some of them 800 m long. The expected noise reduction with any exchange using standard cables is a factor of 2, but there is the possibility to use a new type of cable, i.e. shielded single pair. The

second can potentially give a reduction of more than a factor 5. The mid-term plan is to use the new type of cable wherever it is found to be necessary.

Finally, at the moment the installation consists of 3 batches from different manufacturers and an effort has started to document each cable by manufacturer and understand if there are any correlations.

Over-Injection

During the 2009 operation, a twofold problem has been observed in the injection region. It consists of saturation of the monitors at the LHC TDI elements and exceeding threshold values for transient losses at the monitors protecting the MQXA magnets during injections.

Figures 8 and 9 show plots of the measurements taken while injecting on Beam 1 and 2 respectively as recorded by the 40 μ s (RS01) integration time window. The signal pattern corresponds to a particle shower initiated on the TDI and has been normalised to Gy/s.




Figure 8: Plot of the maximum measured values as recorded by the 40 μ s integration time window (normalised to Gy/s) during an injection in the beam 1 injection region.




Figure 9: Plot of the maximum measured values as recorded by the 40 μ s integration time window (normalised to Gy/s) during an injection in the beam 2 injection region.

Close observation can reveal some difference between the two injection points. For this reason, the radiation source measurements (performed during the summer of

2009) have been reviewed. The installation has also been checked and proved correct.

Moreover, during operation on the 08/12/2009, extensive tests have been done using the second beam to get a better understanding of the situation. The configurations that have been tested:

- Over-injection with the TDI open, 2e9 proton
- Injection on closed TDI (+- 3 mm) and the kicker off
- Injection on closed TDI (+- 3 mm) and TCTV (upper jaw -6 mm, lower jaw -12 mm)

As it can be seen on figure 10, the monitor signals recorded during the consecutive injections with those different configurations exhibit no significant differences. One possible explanation could be that losses come from outside the cryostat.




Figure 10: Plot of three measurements recorded in the beam 2 injection region during consecutive injections using different configurations; overshoot on TDI, injection shot on closed TDI and with, in addition, closed TCTV.

Since the strategy for beam operation for this year will be based on over-injection, there is an immediate need for a solution to this problem. It is not possible to remove those monitors from the protection scheme or connect them to separate power supply lines. In addition, the solution adopted for last year's test, with threshold values for the transient losses set over the maximum allowed limit, is not safe for higher intensities, since it would require modification of the constraints set in the database for allowed values, which have been set as a safeguard from human errors.

The most favourable solution would be to add shielding on the monitors that are foreseen to protect against losses on the opposite beam in those specific areas. Unfortunately, in order to do this correctly, an estimation of its efficiency will be necessary, before installation. The required shower simulation of the TDI is a difficult task to complete in the given timeframe.

To the second part of the problem, i.e. the saturation of the monitors in the TDI, one possible solution is to stretch out the signal in time by hardware means. That is, to introduce an additional capacitor and resistor in the installation, increasing the time constant of the input. This modification will not require changing the threshold

values and the elements those monitors protect will equally be protected from steady-state losses.

Other ongoing modifications

Approximately 7% of the acquisition cards will need to be exchanged due to a non-conformity that arose for a specific production lot. At the moment, those cards show lower than expected offset current.

It has been measured that the ionisation of the air inside the signal interconnection boxes is a significant source of noise, especially for the secondary emission monitors. Actions have been initiated to add better isolation to all those boxes.

The global reset is not reliable on 5% of the acquisition cards due to the accuracy of the resistors used to create the voltage dividers on the backplanes. The strategy for this year was to replace this type of reset with a recently implemented feature that would require the usage of the WorldFIP connections. Nevertheless, latest results have shown some uncertainty about the reliability of WorldFIP under radiation and it has been decided to enable this type of connection only on the cards residing at point 6 to allow a better understanding of the implications.

Finally, an effort is being made to cover the gap in the dynamic range between the ionisation chambers and the secondary emission monitors. A production of prototypes of a shorter ionisation chamber type has started and a few of them will be installed in the LHC tunnel for evaluation during next year's run.

THRESHOLD VALUES

By increasing the beam energy, the level of potential damage to the LHC also increases. The threshold values will have to provide the safe limits for all different elements monitored accordingly. Apart from damage protection, the threshold levels also have to be precise enough not to compromise the operation efficiency by false dumps or magnet quenches.

Thresholds and Noise

It is important to avoid false beam dumps due to noisy monitor channels for all beam energy levels. The analysis examples provided in this section are given grouped into four main beam energies, i.e. 450 GeV, 3.5 TeV, 5.0 TeV and 7. TeV, respectively.

The noise levels have been extracted from the integration time windows of 40 μ s (RS01) and 1.3 s (RS09) as accumulated during 9 hours on 13/01/2010 with the 1 Hz logging, when there was no beam in the machine and no significant disruption of the BLM system occurred. For each ionisation chamber and each RS a total of 32400 values were taken into account for this analysis and only the maximum noise value for each ionisation chamber is plotted in the histogram.

As a result, the comparison between the maximum noise level for each of the 3432 ionisation chambers and their corresponding beam abort thresholds, as set for the 2009 run, is shown in the histograms of figures 11 and 12

for the 40 μ s and 1.3 s integration time windows, respectively.

It should also be noted that:

- The noise data were taken at a time when the system was not yet fully repaired and prepared for this year and thus the actual maximum noise level is expected to become smaller than what is shown.
- The beam abort threshold values for a given energy are different according to the damage and protection level of the single LHC ring elements, e.g. the MB magnets are among the most sensitive elements and therefore have smaller thresholds than those set for MQ magnets.

IC Thresholds and Max. Noise (RS01)




Figure 11: Histogram showing a comparison between the maximum noise levels in the 40 μ s (RS01) integration time window for each of the 3432 ionisation chambers and their corresponding beam abort thresholds, as set for the 2009 operations.

IC Thresholds and Max. Noise (RS09)




Figure 12: Histogram showing a comparison between the maximum noise levels in the 1.3 s (RS09) integration time window for each of the 3432 ionisation chambers and their corresponding beam abort thresholds, as set for the 2009 operations.

The dotted threshold lines indicate the smallest beam abort threshold values for each of the aforementioned beam energies. In order to make sure the system responds to beam losses but not to noisy monitor channels, a signal to noise ratio of a factor of 10 is the minimum required to avoid false beam dump requests. At present, this is ensured up to energies of 5.0 TeV.

IC Collimator Thresholds and Max. Noise (RS01)




Figure 13: Histogram showing a comparison between the maximum noise levels in the 40 μ s (RS01) integration time window for each of the 160 ionisation chambers installed on the collimation elements and their corresponding beam abort thresholds, as set for the 2009 operations.

IC Collimator Thresholds and Max. Noise (RS09)




Figure 14: Histogram showing a comparison between the maximum noise levels in the 1.3 s (RS09) integration time window for each of the 160 ionisation chambers installed on the collimation elements and their corresponding beam abort thresholds, as set for the 2009 operations.

A similar analysis is shown in figures 13 and 14, but here the thresholds and the noise histograms are plotted only for the 160 ionisation monitors installed on the

collimation elements of the LHC. Since they are one of the most robust elements, one would expect their thresholds to be higher in general than for the other elements.

Nevertheless, in order to provide the needed one order of magnitude of signal to noise ratio, it was necessary to set the threshold monitor coefficient to the maximum allowed value for all the monitors at the collimation elements during last year's operation. This means that, contrary to the rest of the monitors that have been set to 10% of the maximum, the threshold values set for the monitors at the collimation elements will not have a margin of increase based on their current calculation, even if it were required.

Validity of threshold values

At the moment most of the threshold values are based on simulations and need to be fine-tuned in the future based on a detailed analysis of the data to be recorded during beam operation.

Even though no threshold values are missing, there are some precise simulations pending as well as some weak points that have been observed during the periods of beam operation.

In order to verify the currently used threshold values and adjust to more precise values, a set of beam tests is being prepared. The transient losses will be cross-checked with dedicated nQPS buffers; the steady-state losses with temperature sensors; fine-tuning will be performed with collimator tests. Especially the latter will need to be done for each collimator type (i.e. primary, secondary, tertiary), but several tests to determine the level of cross-talk and back scattering also have to be foreseen.

For more information see also [3].

REAL-TIME DATA PROCESSING

The backbone of the system, which is responsible for processing the acquired data in real-time, creating integration histories and deciding if operation respects the predefined limits [4], resides in the SR buildings of each IP.

System Response Latency

From last year's operation, several beam dump requests generated by the system have been analysed in detail. In order to independently verify the system latency in acquiring the measurements, processing the data and taking the decision to request a beam dump, timestamps from other systems have been used.

The difference in time between the bunch at the MKI and the break of the beam permit loop (by the BLMS) recorded at the BIC level was always between 100 and 130 μ s. From that value, one should subtract the time of flight (t.o.f.) of the bunch from the MKI to the monitor to get the latency of the system. In the specific tests, the monitor was placed approximately 3 km away from the injection point, which corresponds to a t.o.f. on the order of 10 μ s.

Other delays that are incurred in the system are: the t.o.f. of the signal from the monitor to the acquisition electronics – on average 0.5 km of cables or 3 μ s; the t.o.f. of the digitised signal from the acquisition to the processing electronics – on average 1 km of fibre or 3 μ s; the detection of the change in frequency in the daisy-chain for the beam dump requests in the VME crate – approximately 5 μ s; the integration window in the acquisition electronics – once every 40 μ s. The processing of the data and the decision at the BLETC (for transient losses exceeding the threshold) is in the sub- μ s level.

Even though the system latency is well below the allowed maximum, it is clear from the above that an unexpected delay has been added to the system. It has recently been identified as an extra 40 μ s cycle in the merging of the ADC and CFC data. It will be investigated if this can be safely removed in a future release.

System Self-Monitoring

In order to provide the protection level required for the next state of the LHC operation with higher intensities, the activation of all of the system's continuous self-monitoring processes will be required.

After their activation on the firmware level, any of the following cases will be able to set the beam permit state to FALSE and initiate an emergency dump:

- Configuration cross-check failure - one case, where the embedded serial numbers in the packets are checked to be coming from the same acquisition card.
- Data transmission between the tunnel and the surface installation provides non-valid data - 22 cases, where checks on the continuity of packets, the validity and correctness of the received data packets (by embedding frame numbers in the data stream), card identification and strong checksums are done every 40 μ s.
- Any update of the operational parameters, i.e. of the threshold values or configuration settings.
- Whenever the system enters any of the regular self-check modes (described later in the document).

Nevertheless, it should be noted that exceptional effort has been put into designing those checks, not only to be rigorous and strict but also to avoid unnecessary beam dump requests as much as possible and, in some cases, even provide an increase in the availability of the system.

By smartly exploiting the redundancy of the optical links, the function will allow the continuation of beam operation if a certain amount of checks validate for at least one of the redundant data packets received. The non-conforming links will be repaired or reconfigured using the spare fibres in the next technical stop.

To avoid unnecessary stops, beam dump requests generated by checks in cards which have none of their channels connected to the BIS will be automatically inhibited.

Similarly, the system will ignore any requests to update its threshold tables or to enter any of the externally

initiated self-check modes when the BEAM INFO flag (coming from the BIS) is not set to the TRUE state.

APPLICATIONS AND DATABASES

The system requires vast amount of parameters for operation and at the same time produces an abundance of data for display in the control room, to assist in the configuration and operation of several other systems and even after heavy reduction a large amount of measurements is sent for long term storage in the Logging database [5].

Storage of Parameters

For a highly complex system like the LHC BLMS, several millions of database entries are required in order to document, generate and manage its parameters, i.e. the configuration data, threshold values and system settings, and to provide a fully functional system under all operational conditions.

At this moment in time, where the uploading and verification of all this vast amount of data has been completed, the requirements and expectations are changing. For this reason, a restructuring of all databases used to enter and store parameters has been initiated to provide simpler views, easier manipulation and more features in the years to come.

Those actions, summarised for each database employed, include:

- *MTF* – removal of the BLMS architecture; only the reference measurements and the track history of components will be kept.
- *Layout* – will become the new entry point for updated or new connections and installations.
- *LSA* – its tables holding either system architectural data transferred from Layout db or threshold values are being reorganised to become more relational.

The LSA reorganisation has already been completed and verified to be correct. This will allow continuing the rest of the restructuring throughout the year without affecting the operation and safety of the system.

Moreover, numerous constraints on the type and the range of the values one can enter for each field have been configured internally in the LSA database. Those constraints will be reviewed and augmented, where necessary, to provide a maximum of protection from human errors.

Finally, one major additional protection feature will be included this year. Before a commit action to the DB is executed, a check will be made for the currently disabled channels. A set of rules have been created for each monitor based on criticality and the maximum number of allowed adjacent disabled channels. Each of the installed monitors is already being tagged for its criticality and a first version is under test in the development database. It is expected to be operational before the start-up and any rule violation will block the commit of a new configuration.

Data Concentration

Logging the measurement and status data from the system at 1 Hz has undergone a major improvement in the past year as a result of the efforts of the CO/AP and BI/SW sections. Several modifications of the CMW settings and a rebuild of the Concentrator processes have improved the situation and also caused the positive side effect of significantly increasing the stability of the front end computers (FECs). Unfortunately, the reasons for randomly losing packets are still uncertain and the investigation will continue.

The problem with the ability to deliver the Capture data has been solved with an unorthodox solution of delivering the notification of new data and the actual data through separate properties, which is now called a “synchronised GET”, and by reducing the amount of data requested by three quarters. An effort will be made to return to the nominal 2048 samples per channel.

The Java Messaging System’s broker had only one period of unavailability last year. Due to the criticality of its service, a new configuration is being investigated with the usage of a dedicated broker for the BLM system and a failover cluster configuration. Several tests are foreseen during the dry runs.

Warnings (ALARMS)

In order to provide the necessary information to the LHC operators, it has been agreed a set of warnings to propagate in the ALARMS consoles, which is now under implementation.

Those warnings will be generated and shown whenever:

- A system check is ongoing. It will be a concentrated value from all checks needed to be performed on the system. During this period it will be not possible to arm the BIS and thus inject beam into the LHC.
- A timeout from the last system check has occurred. In that case, execution of the automated regular checks will be mandatory before the next injection.
- One of the system checks has failed. In that case, the console should give some details on actions.
- Transfer of measurement data is failing to reach the applications or the database. It will provide separate warnings for each of the main elements in the infrastructure, i.e. the devices, the concentration and logger processes, the transportation layer or the unavailability of the database.
- A beam dump request has been made.

Data Displays and Expert Applications

In general, the main set of applications, like the fixed display, capture data display, monitor factor editor, expert threshold editor, etc., needed for the setup and operation of the system or to display its data are ready or will require minimal changes. Work will concentrate mainly on enhancing the following:

- Post-Mortem data and analysis modules - adding some more variables, and more functions.

- Status application – some of the main parts have been created but several parts are still missing; needs usability update for OP.
- Management of parameters – provide detailed log of each user and their actions, dedicated DB API, configuration of mobile monitors, and simplification of the enabling/disabling of channels.
- Diagnostics application for misbehaving channels or when checks fail - expert tool initially but later will be handed over to the operations group.

RELIABILITY AND AVAILABILITY

The BLM system's high reliability and availability have been major concerns since the early stages of its design cycle. Several features have been embedded that continuously monitor and check the system during operation. Aside of those, a whole suite of additional checks is being deployed, some of them to be run after each design change, while others after any maintenance interventions or before each injection.

Automated Regular Checks

In order to guarantee the requested SIL level a set of automatic checks that need to be performed on a regular basis has been included in the system's design.

Those are:

- The MCS Online check, which verifies the consistency of the parameters stored onboard, i.e. settings and threshold values, with those at the database level.
- The Connectivity (HV Modulation) check, which verifies that all monitor connections, and only those, are present, that the monitor connected is of the expected type, and that those monitors' responses are inside predefined boundaries (see also the connectivity check chapter).
- The Internal Beam Permit check, which verifies the ability of each of the threshold comparator (BLETC) modules in the daisy-chain to trigger a dump request and that the Combiner & Survey (BLECS) module is able to receive those requests.

Since performing any of those checks removes the Beam Permit signal sent to the BIS and compromises the measurement ability of the system, the control and initiation of those tasks has been given to the LHC Sequencer. The development and testing cycle of those specific sequencer tasks is now in its final stages. In turn, the BLMS will activate the enforcing of their agreed periodic run on the firmware level (using internal timers) and on failure or timeout will lock the Beam Permit in the false state on the next possible occasion.

Finally, one more check has been implemented, this time to verify the BLMS connection to the Beam Interlock System (BIS). This is known as the *External Beam Permit check*, and has been handed over to the responsibility of the BIS team.

Connectivity Check

A current on the monitors is generated through modulation of the monitor bias high voltage and measured by the normal acquisition system to determine the connection status of all the system's monitor inputs (see also [6]).

The BLECS module uses the integrated values, i.e. RS09, from the BLETC modules to determine the amplitude and phase of each monitor by employing digital FIR filters in its onboard FPGA device.

The results are sent to the Logging database and are also compared to the relevant thresholds to permit or block the next injection if a non-conformity is detected.

The detailed signals recorded for each monitor are stored in the cards' embedded memories (see figure 15). They can be read on demand and further analysed offline with the dedicated diagnostics application.




Figure 15: View of the SRAM data (in bits vs. sample) containing the original (stairs) and filtered RS09 data of each channel for one crate.

The minimum and maximum threshold values (usually referred to as internal parameters) defining the acceptance region for the phase and amplitude measured through this function, are unique for each monitor. They are calculated out of multiple measurements and the analysis of the logged values.

System Verification Tests

Acquisition card problems, cable degradation, channel noise, ground loops and in general many of the unwanted changes in the acquisition properties of the system can be identified by analysing the frequency distribution of the channel noise.

At the moment, the noise level is checked for all channels on a regular basis and especially before and after any intervention. This has been implemented as a semi-automatic task, where a large amount of data are extracted from the Measurement database, analysed and plotted using the ROOT framework for data processing and the results are presented for human inspection.

A similar task is being conceived to analyse and report regularly on the degradation of optical fibre links either from aging or radiation. Configuration changes have already been made to propagate all the necessary data to the long term storage of the Logging database.

The two tasks will not only improve the availability of the system but will also help to organise the needed regular maintenance. Therefore, an effort has been initiated to turn those into fully automatic analysis packages that will be executed inside the database on a regular basis to avoid the unnecessary extraction of significant amounts of data that are needed.

Vertical Slice Test system

In an effort to reduce the risk of impeding the ability of the BLM system to perform its critical functions, a verification environment based on custom hardware and software is under development (see also [7]).

This verification environment, named Vertical Slice Test system, residing next to the operational crates of point 2 can perform the following tests:

- Exhaustive measurement over threshold values tests - where all possible permutations of the threshold comparator and their ability to trigger a beam dump are checked.
- Correct reception and status tests - where all the system self-check mechanisms (described earlier in the document) are tested for their ability to detect the erroneous cases and trigger a beam dump.
- Direct injection of linearly increasing, impulse, and other predefined patterns of input signals - where the response of the system and its ability to process the data correctly is tested. This part also has the ability to 'playback' any signal patterns recorded by the post-mortem and capture buffers.

All the above test cases are already part of the standard 'protocol' followed before releasing a new firmware and have to be executed in the vertical slice test system, while some of them are used in the commissioning strategy for the operational system.

Other Known Issues:

Other issues that have been identified in the processing or the readout of the data stages of the system and will require solutions are:

- The PM and XPOC data are sometimes wrongly decoded – firmware issue.
- The channels that have been over threshold remain in history and are appended to the PM data every time – FESA issue.
- The MCS online check always fails the first time after a reboot – MCS issue.

ACKNOWLEDGEMENTS

The authors would like to express their gratitude for the support given by the Controls and Operations groups (special thanks go to Marek Misiowiec, Fabio Follin, Chris Roderick, Grzegorz Kruk, Pascal Le Roux, and Sonia Mallon Amerigo), to Brennan Goddard and Jorg Wenninger for their assistance in taking measurements with beam, and to the Protvino and FSU teams (especially Ion Savu and Romain Ruffieux) for their diligence in the installation and measurement of the tunnel installation.

REFERENCES

- [1] B. Dehning et al., "First Experience with the LHC Beam Loss Monitoring System", Particle Accelerator Conference 2009 (PAC 09), Vancouver, Canada, 2009.
- [2] E. Effinger, B. Dehning, J. Emery, G. Ferioli, G. Gauglio, C. Zamantzas, "The LHC Beam Loss Monitoring System's Data Acquisition Card", 12th Workshop on Electronics for LHC and future Experiments (LECC 06), Valencia, Spain, 2006.
- [3] E.B. Holzer, B. Dehning, E. Effinger, J. Emery, S. Grishin, C. Jr Hajdu, H. Ikeda, S. Jackson, C. Kurfuerst, A. Marsili, A. Nordt, J. Perez Messeri, V. Prieto, M. Sapinski, C. Zamantzas, "Lessons learnt from beam commissioning and early beam operation of the Beam Loss Monitors (including outlook to 5TeV)", LHC Performance Workshop - Chamonix 2010.
- [4] C. Zamantzas, B. Dehning, E. Effinger, J. Emery, G. Ferioli, "An FPGA Based Implementation for Real-Time Processing of the LHC Beam Loss Monitoring System's Data", Nuclear Science Symposium Conference Record, 2006. IEEE, vol.2, no., pp.950-954, 2006.
- [5] C. Roderick, R. Billen, RD Gaspar Aparicio, E. Grancher, A. Khodabandeh, N. Segura Chinchilla, "The LHC Logging Service : Handling terabytes of on-line data ", 12th International Conference On Accelerator And Large Experimental Physics Control Systems (ICALEPCS 09), KOBE, Japan, 2009.
- [6] J. Emery, B. Dehning, E. Effinger, G. Ferioli, H. Ikeda, E. Verhagen, C. Zamantzas, "LHC BLM Single Channel Connectivity Test using the Standard Installation", 9th European Workshop on Beam Diagnostics and Instrumentation for Particle Accelerators (DIPAC 09), Basel, Switzerland, 2009.
- [7] C. Zamantzas, B. Dehning, C. F. Hajdu, S. Jackson, "Reliability Tests of the LHC Beam Loss Monitoring FPGA Firmware", to be published 2010.