

THE LHC BEAM LOSS MONITORING SYSTEM'S REAL-TIME DATA ANALYSIS CARD

C. Zamantzas, B. Dehning, E. Effinger, G. Ferioli, G. Guaglio, R. Leitner
CERN, Geneva, Switzerland

Abstract

The BLM (Beam Loss Monitoring) system has to prevent the superconducting magnets from being quenched and protect the machine components against damages making it one of the most critical elements for the protection of the LHC. The complete system consists of 3600 detectors, placed at various locations around the ring, tunnel electronics, which are responsible for acquiring, digitising, and transmitting the data, and surface electronics, which receive the data via 2km optical data links, process, analyze, store, and issue

ANALYSIS CARD (BLMTC)

The data analysis card is based on the general purpose PCB that was implemented for the whole instrumentation group, named DAB64x [3]. It is comprised from an Altera Stratix™ [4] FPGA, an Altera MAX™ CPLD [5] for power-on configuration and VME functionality, and three SRAM memories. The functionality of the system is realised by using different firmware on the FPGA and CPLD devices, and different mezzanine cards that can be placed on either of the six general-purpose connectors.

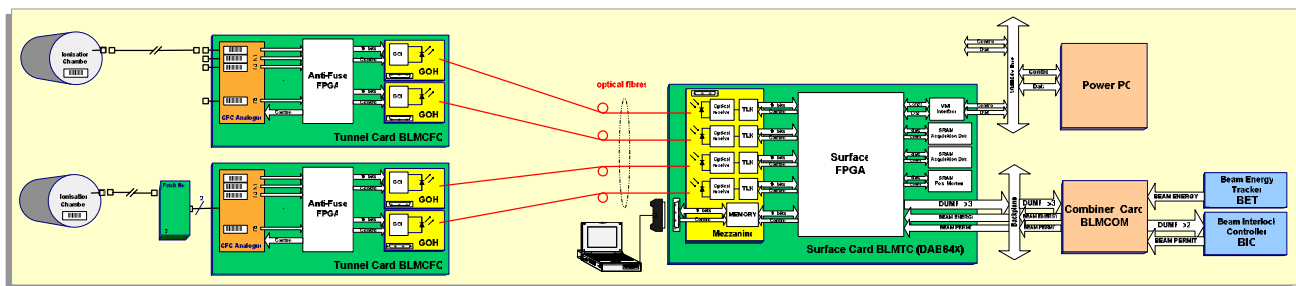


Figure 1: Overview of the complete BLM System for the LHC.

warning and abort triggers. At those surface units, named BLMTCs, the backbone on each of them is an FPGA (field programmable gate array) which treats the loss signals collected from 16 detectors. It takes into account the beam energy and keeps 192 running sums giving loss durations of up to the last 84 seconds before it compares them with thresholds uniquely programmable for each detector. In this paper, the BLMTC's design is explored giving emphasis to the strategies followed in combining the data from the integrator and the ADC, and in keeping the running sums updated in a way that gives the best compromise between memory needs, computation, and approximation error.

INTRODUCTION

Around 3600 *Ionization Chambers* are the detectors of the system. A set of up to 8 of them can be connected to each of the tunnel cards. In those tunnel cards, called *BLMFCs*, the digitisation of the detector signal is done by using *CFCs* (current-to-frequency converters) and *ADCs*. An *Anti-Fuse FPGA* [1] acquires the digitised data and transmits them at the surface using the *GOLs* (*Gigabit Optical Links*) [2]. There, the data analysis cards, named *BLMTCs*, receive those data. The surface *FPGA* will analyse by keeping a history of those data and decide whether or not a dump request should be initiated. Each surface card receives data from 2 tunnel cards, which means that it can treat up to 16 channels (see Fig.1).

The BLM mezzanine card [6] includes all necessary components for the four gigabit optical receivers and a Mbit of Flash memory for system specific data.

PROCESSES RUNNING IN THE FPGA

The surface FPGA has 41,000 LE (logic elements), 615 user I/O pins, and 400KBytes of internal memory as available resources for the system to be built from.

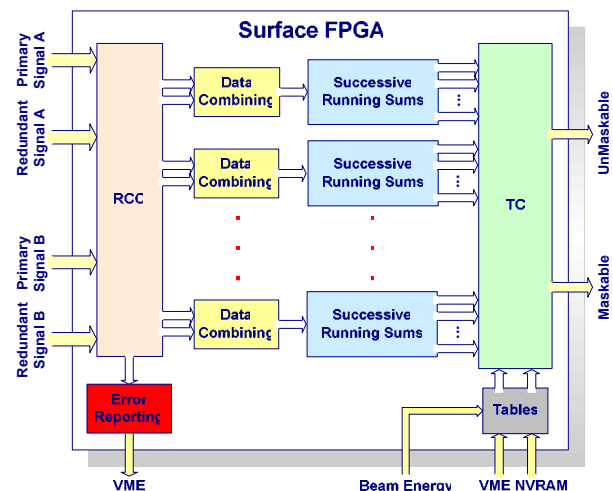


Figure 2: Block diagram of processes related to data analysis running in the surface FPGA.

A description of each of the system blocks related to the data analysis process follows (see Fig.2).

Transmission Check (RCC Block)

In the *BLM* system, for reliability reasons, many parts were decided to have redundancy. The communication link was one of them. A comparison of the two could trigger a dump of the beam in the case of a difference. By including, some digital techniques, like the 8b/10bit and the CRC (Cyclic Redundancy Check), a more sophisticated receiver part was produced that achieves even higher data reliability and system availability.

Table 1: Decision Table for the Signal Select Block.

| CRC Check | | Compare | Output | Remarks |
|-----------|-------|---------|----------|----------------------------------|
| A | B | A to B | | |
| Error | Error | Error | Dump | Both signals have error |
| Error | Error | OK | Dump | CRC generation or check is wrong |
| Error | OK | Error | Signal B | Error at CRC part |
| Error | OK | OK | Signal B | Error at data part |
| OK | Error | Error | Signal A | Error at CRC part |
| OK | Error | OK | Signal A | Error at data part |
| OK | OK | Error | Dump | One of the counters has error |
| OK | OK | OK | Signal A | Both signals are correct |

With the 8b/10b encoding before transmission, the complete frame is split into 8-bit blocks and each of them is encoded into a 10-bit block. On receipt, the reverse procedure (i.e. decoding) is used. It serves two purposes. First, it makes sure there are enough transitions in the serial data stream so the clock can be recovered easily from the embedded data. Second, because it transmits the same number of ones as zeros, it maintains a DC balance.

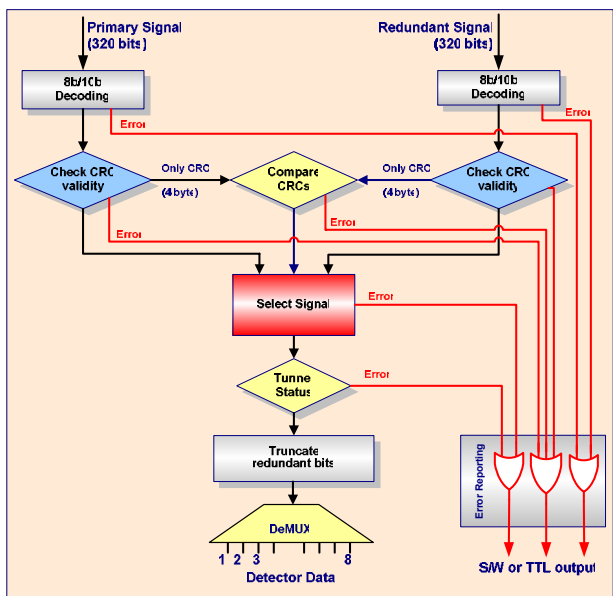


Figure 3: Transmission Check (RCC) Block Diagram.

The *CRC*, a polynomial arithmetic based algorithm, is widely used as error detection scheme. In this system, the redundant bits it introduces are not only used to detect transmission errors but also to find any difference between the two signals.

The “*Signal Select*” block receives the outputs of the checks and the comparison and gets the responsibility of selecting the error free signal to convey on the stages below. By using efficiently all information it is expected to increase the availability of the system (see Table 1, cases 3 to 6). Additionally, its error reporting features can also indicate problematic areas and failing components in the tunnel installation (see Fig.3).

CFC & ADC Data Combining

Based on the fact that the two types of data acquired for each detector are different, a pre-processing is needed for them to be combined seamlessly. The measurement with a counter of the frequency produced by the *Current-to-Frequency Converter* relates to the average current between the two last acquisitions. On the other hand, the voltage measured by the *ADC* is the voltage of the *CFC* integrator at the time of the counter readout. It represents then the fraction remained between the last count and the first from the next acquisition.

In order to combine those data the difference of the last two *ADC* measurements is needed. It will then correspond to the counter fraction of the last 40 μs and thus could be added to the counter value. Of course, since the difference could be a negative number, signed number arithmetic is used for the addition (see Fig.4).

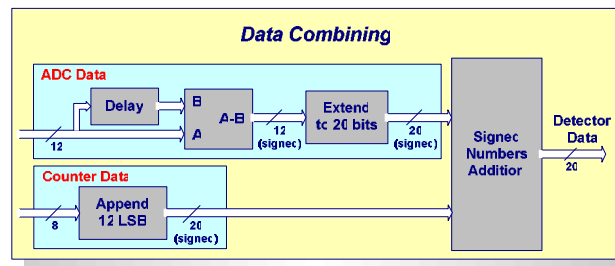


Figure 4: Data Combining Block Diagram.

Using this scheme, the usage of resources and computation effort is minimised since it is concentrated in this beginning stage. All later stages are just exploiting those combined values.

Successive Running Sums

The *Running Sums* can be produced simply by adding the newly acquired value to a register and subsequently subtracting from it a value coming delayed by a number of cycles in a shift register. A similar configuration is used in the *BLMTC* where, in order to increase the efficiency in resources, it is making use of the multipoint shift registers available on the Stratix devices. It is simply adding the difference of those two values, the new and the delayed, to an Accumulator.

Another scheme of resource sharing is employed for reaching higher integration time with relatively small in

length shift registers. It makes use of the already calculated sums in order to calculate bigger in length running sums with the only expense of some additional latency (see Fig.5).




Figure 5: Production of Successive Running Sums.

The latency introduced has little effect to the optimal approximation accuracy though since it varies between them. More specifically, the running sums that span to the low range (fast losses) have zero or very small additional latency and gradually increases reaching up to 1.3s for the 21 and 83s time range (see Table 2).

Table 2: Integration Time-Windows Configuration.

The red coloured RSs (running sums) outputs, i.e. RS1, RS4, RS6, and RS8, represent their additional utilisation as inputs for the adjacent SRs (shift registers), i.e. SR2, SR3, SR4, and SR5.

| Range | | Refreshing | | Shift Reg. Name | Signal Name |
|------------------|----------|------------------|---------|-----------------|-------------|
| 40 μ s steps | ms | 40 μ s steps | ms | | |
| 1 | 0.04 | 1 | 0.04 | | RS0 |
| 2 | 0.08 | 1 | 0.04 | | RS1 |
| 8 | 0.32 | 1 | 0.04 | SR1 | RS2 |
| 16 | 0.64 | 1 | 0.04 | | RS3 |
| 64 | 2.56 | 2 | 0.08 | SR2 | RS4 |
| 256 | 10.24 | 2 | 0.08 | | RS5 |
| 2048 | 81.92 | 64 | 2.56 | SR3 | RS6 |
| 8192 | 327.68 | 64 | 2.56 | | RS7 |
| 32768 | 1310.72 | 2048 | 81.92 | SR4 | RS8 |
| 131072 | 5242.88 | 2048 | 81.92 | | RS9 |
| 524288 | 20971.52 | 32768 | 1310.72 | SR5 | RS10 |
| 2097152 | 83886.08 | 32768 | 1310.72 | | RS11 |

This gained efficiency was necessary for this system to be applicable in a configuration with relatively very low memory available. In a “normal” configuration of this system, the running sums would need to hold approximately 3 million values for each of the 16 detectors to achieve the same approximation error. That means a total of approx. 150MBytes. Instead, by using the successive running sums the system is using only some of the FPGA internal memory since it does not need more than 100KBytes.

Threshold and Masking Tables (TC Block)

The quench and damage levels are time and energy dependent [7, 8]. For this reason, from each detector’s data, as it was shown, this system calculates 11 *Running Sums*. Every *Running Sum*, after every new calculation, is compared with its corresponding *Threshold* value that was chosen by the beam energy reading at that moment. If found to be higher, the *Comparator* will initiate the necessary dump request (see Fig. 6).

All requests are being gathered by a *Masking* block with the purpose of distinguishing between “Maskable” and “Unmaskable” channels. The tables are stored on the BLM Mezzanine’s non-volatile memory and can be unique for each detector. For added functionality both of them can be loaded locally either by the frond-panel’s connector or by the Power-PC.

Finally, all outputs will be summed and collected by the “Combiner Card” which will forward them to the “Beam Interlock System”.

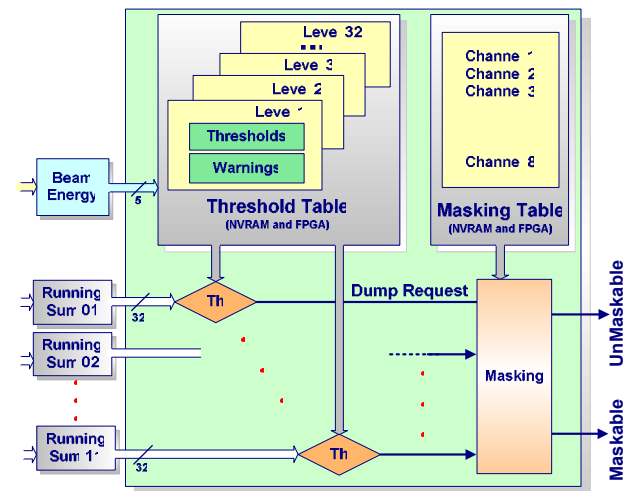


Figure 6: Threshold Comparator (TC) Block Diagram.

REFERENCES

- [1] Actel Antifuse Devices SX-A/ SX, <http://actel.com/>
- [2] P. Moreira, G. Cervelli, J. Christiansen, F. Faccio, A. Kluge, A. Marchioro, T. Toifl, J. P. Cachemiche and M. Menouni “A Radiation Tolerant Gigabit Serializer for LHC Data Transmission”, *Proceedings of the Seventh Workshop on Electronics for LHC Experiments*, Stockholm, Sweden, 10-14 September 2001
- [3] R. Jones, “VME64x Digital Acquisition Board for the LHC Trajectory and Closed Orbit System”, *LHC Engineering Specification, LHC-BP-ES-0002*.
- [4] Altera Stratix Device Family Overview, www.altera.com/products/devices/stratix/overview/stx-overview.html
- [5] MAX 3000A CPLD Family Overview, www.altera.com/products/devices/max3k/overview/m3k-overview.html
- [6] Schematic files of the BLM Mezzanine Card (version 2), <https://edms.cern.ch/item/EDA-00780>
- [7] J.B. Jeanneret, D. Leroy, L. Oberli and T. Trenkler, LHC Project, “Quench levels and transient beam losses in LHC magnets”, *LHC Report 44, CERN, July 1996*.
- [8] A. Arauzo-Garcia et al., “LHC Beam Loss Monitors”, *5th European Workshop on Diagnostics and Beam Instrumentation DIPAC 2001, Grenoble, France*.