Critical Systems Labs
Strategic Insight for Safety

Review of the
CERN LHC Beam Loss Monitoring System

7 February 2011

Submitted by:

Critical Systems Labs, Inc.
#618 - 475 Howe Street
Vancouver, B.C.
Canada V6C 2B3

Submitted to:

CERN
European Organization for Nuclear Research
CH-1211
Genève 23 Switzerland

Executive Summary

The CERN LHC Beam Loss Monitoring System (BLMS) is an integral element of CERN's approach to the protection of the LHC machine. The primary purpose of this system is to request a beam dump in response to a dangerous beam loss. The BLMS is a single point of failure for protection of the LHC, i.e., a failure of the BLMS to request a beam dump in response to a dangerous beam loss could be the sole reason why the beam is not dumped quickly enough to prevent catastrophic damage to the LHC.

The BLMS is a complex system whose design involves a variety of engineering challenges including extremely stringent, hard real-time constraints, distribution over a large area, harsh environment (e.g., radioactivity), intensive real-time data processing, sensing of physical phenomena (i.e., sub-atomic particles) and a planned lifetime that spans multiple decades. The design incorporates a variety of technologies ranging from analogue electronics to software written in a high-level programming language. Additionally, the designers of the BLMS have faced the challenge of minimizing unnecessary beam dump requests without compromising the safety of the BLMS, i.e., availability vs. safety.

We have reviewed the BLMS at several different levels of abstraction ranging from details of the VHDL code used to synthesize circuitry operating at the nanosecond timescale to matters of human performance in regard to management of threshold settings used by the BLMS.

It is clear to us that the BLMS has been developed by a team of skilful and conscientious experts who are deeply committed to the overall goals of LHC machine protection under very experienced and capable leadership.

With one minor reservation, we have found no reason to be concerned that the current configuration of the BLMS might fail to request a beam dump in response to a dangerous beam loss. This conclusion is based on the following assumptions: (1) appropriate threshold settings are used; (2) the current operating procedures, including regular execution of the "connectivity test", are maintained; (3) no element of the BLMS, such as an individual detector, is disabled without a decision by the appropriate authority that the safety risk is acceptable; (4) known limitations of the BLMS are adequately addressed before the machine is used at higher energies; and (5) the risk associated with the dependency of the BLMS on other systems for an accurate indication of bean energy is deemed to be acceptable by the appropriate authority. Changes to any aspect of the design may invalidate this conclusion.

The one minor reservation associated with the above conclusion is the fact that the critical real-time datapath responsible for generation of a beam dump request is not fully redundant. Hence, there are single points of failure in the design of this critical portion of the BLMS. In this regard, we are concerned mostly about non-redundant elements of this datapath within the Threshold Comparator Field Programmable Gate Array (FPGA).

This report includes a number of specific recommendations for the consideration of the BLMS developers and other stakeholders. More generally, some of these recommendations may also be of interest to CERN in regard to the development of machine protection technology in future accelerator projects.

Revision History

Version	Description	Status	Date
-	Document created.	DRAFT	31 January 2011
1.0	Document updated to incorporate revisions from internal CSL review.	FINAL	7 February 2011

Contents

1. Introduction

The CERN LHC Beam Loss Monitoring System (BLMS) is an integral element of CERN's approach to the protection of the LHC machine. The primary purpose of this system is to request a beam dump, by means of input to the Beam Interlock System (BIS), in response to a dangerous beam loss. The BLMS is a single point of failure for protection of the LHC, i.e., a failure of the BLMS to request a beam dump in response to a dangerous beam loss could be the sole reason why the beam is not dumped quickly enough to prevent catastrophic damage to the LHC. Although other kinds of "user system" inputs to the BIS might also trigger a beam dump request, the BLMS is regarded as the primary input to the BIS with respect to the protection of the LHC when there is circulating beam in the machine.

In the spring of 2010, CERN engaged Critical Systems Labs Inc. (CSL) to review the design of the BLMS. A team of CSL reviewers visited CERN in November 2010 for intensive briefings and discussions on the design of the BLMS. The CSL reviewers have also had access to a variety of documents and artifacts (e.g., VHDL code) for the purposes of this review.

Under this engagement, CSL has not performed, nor was expected to perform, a complete safety analysis of the BLMS. For this reason, this report does not offer a direct conclusion about the safety of the BLMS. However CSL has examined the design and implementation of the BLMS for the purpose of assessing the extent to which the BLMS has been developed with level of rigor and thoroughness appropriate for a system of this criticality. CSL has also examined many details of the BLMS design with respect to potential failure modes based on our experience with safety-critical systems across a wide range of industry domains.

The BLMS is a complex system for a variety of reasons including:

1. It has very stringent, hard real-time constraints, i.e., inputs from detectors must be compared to threshold settings every 40 microseconds.

2. The BLMS is a highly distributed system. It includes approximately 4,000 detectors distributed around the 27 km length of the tunnel, together with electronics at various locations in the tunnel and on the surface. Analogue signal transmission cables from the detectors to the tunnel electronics may be as long as 2 km in the long straight sections of the tunnel.

3. The detectors and tunnel electronics operate in an environment subjected to radioactivity. The surface electronics operate in a physical environment comparable to a factory or workshop for light/medium industry, e.g., dust, seasonal fluctuations in ambient temperature and humidity.

4. The surface electronics process a very large amount of data in real-time, e.g., for each of the approximately 4,000 detectors, the BLMS calculates the sum of the previous 2,097,152 inputs from the detector every 1.3 seconds approximately. A number of other sums are also calculated for each detector over shorter intervals.

5. The BLMS is a mixture of analogue and digital electronics, including the use of very specialized sensors. While the detectors are based on previous CERN experience, they are not commercial mass-produced “off-the-shelf” devices with quantified reliability data.

6. The BLMS design incorporates a wide variety of technologies ranging from analogue electronics to software written in a high-level programming language, thus requiring a very extensive range of expertise and knowledge from the BLMS developers.

7. While other accelerators at CERN and elsewhere also include systems to monitor beam loss, the LHC BLMS is required to operate at a higher energy level and with a larger dynamic range.

8. The planned lifetime of the BLMS spans multiple decades.

It is a remarkable accomplishment for CERN, and especially the BLMS developers, to have created a solution that addresses all of the above challenges. It is easy to imagine that the development of a comparable system by industry would have required a much larger team of experts to produce a satisfactory solution.

The essential and foremost question that has driven the technical review of the BLMS by CSL is “Are the digital and programmable parts of the BLMS going to perform as they are intended to in the context of the BLMS?”

More specifically, CSL was asked to:

- assess the adequacy of the overall BLMS design with a focus on the programmable parts

- identify possible weaknesses in the programmable parts of the mission-critical BLMS

- suggest activities that could increase the level of confidence that the programmable parts of BLMS perform as intended

- suggest potential improvements of the BLMS

- provide a general comparison of the BLMS with approaches in industrial systems.

For the purpose of this project, components of the BLM system can be itemized as follows:

1. Beam loss detectors – The passage of particles through the detector results in charge. This part is purely analogue.

2. Acquisition card (BLECF) – Performs some analogue processing and then produces a digital output.

3. Threshold Comparator card (BLETC) – This is entirely digital and likely the most complex component in this list. The real-time data processing is performed by an FPGA that makes uses of preset thresholds and masking information.

4. Combiner and Survey card (BLECS) – This combines outputs from multiple instances of threshold comparators to produce one signal. This card also receives and distributes an indication of the current beam energy.

5. Settings Management – This is an oracle database of threshold settings which are accessed, distributed and then read-back to verify that the settings in use are valid. This involves software implementation (in C and Java).

6. Display of beam losses monitoring data for post mortem analysis.

The scope of this review, as agreed with CERN, has been limited to Component #2 (digital part only) and components #3 to #5. Component #6 of the above list is outside the scope of this project.

The task of reviewing the "Settings Management" component was not expected to include a review of the framework code, which is understood by CSL to be previously developed code used by a variety of projects in CERN. However, the review was expected to consider how well the application code conforms to rules of use for the framework, as given in documents provided in advance to CSL in preparation for the site visit.

The scope of this review is also limited to a consideration of:

1. Potential sources of un-safety within the BLMS, where the detection of an amount of particle losses that has the potential to quench the magnet is not relayed to the Beam Interlock System, resulting in a 'missed generation of beam dump trigger' and potential machine damage.

2. Potential sources of unavailability, where failure of the BLMS leads to a request to the Beam Interlock System to dump the beam, resulting in a 'false dump trigger' and some machine downtime.

While motivated by CERN's interest in the protection of the LHC from damage, this review is not a safety analysis of the BLM. In particular, the scope of this review does not include the task of identifying additional hazards. For the purpose of this review, the only hazard of interest is a 'missed generation of beam dump trigger', i.e., when the BLMS fails to generate a dump request in response to a dangerous beam loss.

A quantitative assessment of residual risk or reliability is also outside the scope of this review.

Section 4 of this report includes a number of specific recommendations. We suggest that a responsible authority within CERN should ensure each recommendation in this report is addressed within six months of the delivery of this report to CERN.

2. References

Documents

[1] Functional Specification, LHC-BLM-ES-0001 Rev 2.0.

[2] B. Dehning, E. Effinger, J. Emery, G. Ferioli, C. Zamantzas, Single Gain Radiation Tolerant LHC Beam Loss Acquisition Card, presented at Presented at DIPAC'07 – 20-23 May 2007 – Venice-Mestre/IT, CERN-AB-2007-028 BI.

[3] B. Dehning, E.Effinger, 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.

[4] B. Dehning, E. Effinger, J. Emery, R. Leitner, C. Zamantzas, Functional and Linearity Test System for the LHC Beam Loss Monitoring Data Acquisition Card, CERN-AB-2007-063 BI.

[5] C. Zamantzas, The Real-Time Data Analysis and Decision System for Particle Flux Detection in the LHC Accelerator at CERN., Brunel University, CERN-THESIS-2006-037.

[6] MPS Commissioning Procedure: MPS Aspects Of The Beam Loss Monitor System Commissioning.

[7] The Commissioning Of Beam Instrumentation on the LHC Sectors: Individual System Tests Of the LHC Beam Loss [BLM] Monitors.

Presentations

[8] B. Dehning, The LHC Beam Loss Monitoring: presentation, 8 November 2010.

[9] E.Effinger, BLM tunnel installation and data acquisition card (BLECF), 8 November 2010.

[10] Christos Zamantzas, The Beam Loss Threshold Comparators card (BLETC), 8 November 2010.

[11] J. Emery, The Beam Loss Combiner and Survey card (BLECS), 8 November 2010.

[12] B. Dehning, System Issues observed, 8 November 2010.

3. Method

The CSL review team is highly experienced in the development, analysis and verification of safety-critical systems across a wide variety of technical domains including aerospace, rail signaling, defence, medical technology and advanced automotive electronics. Although most of our experience is from outside the accelerator community, CSL performed a similar review of the LHC BIS in 2009.

This review was performed in three phases. The first phase involved a study of various documents and related artifacts including CERN technical documents, journal articles, conference papers, VHDL code and a doctoral dissertation by one of the BLMS developers, Dr. C. Zamantzas. Next, three members of the CSL review team (L. Fabre, N. Ghafari and J. Joyce) visited CERN for four days, 8-11 November 2010 for intensive discussions with the BLMS developers. On the final day of this site visit, the CSL team presented a number of preliminary findings. Following this site visit, the third phase of this review involved further study of the above-mentioned BLMS documentation and related artifacts, as well as additional material collected during the site visit to CERN. This third phrase has also included several informal email discussions between the review team and the BLMS developers.

The BLMS development team prepared an extraordinarily rich set of presentations for the site visit. These presentations, the doctoral dissertation written by Dr. Zamantzas and the VHDL code have been the most important sources of information about the BLMS.

The CSL team encountered two significant obstacles during this review. As explained in Section 4.8 of this report, the absence of a comprehensive specification of the BLMS requirements made it difficult to assess certain aspects of the system. The absence of a comprehensive specification also made it more difficult to assess the adequacy of the efforts to verify the design of the BLMS. The other main obstacle encountered by the CSL team was the difficulty of understanding some aspects of the VHDL code. This difficulty was largely due to the lack of meaningful comments in the VHDL code and related problems described in Section 4.9.

Although a review of the Settings Management application code is within the planned scope of this review, this code has not been reviewed due to time constraints. In particular, the time allocated for the review of this code was instead spent on the review of the VHDL code which was more time-consuming than anticipated.

4. Findings

This section of the report presents findings that have resulted from our review of the BLMS. A number of recommendations are also stated in the context of specific findings. These recommendations are enumerated in Appendix A.

4.1 Novelty of Design Solutions

While most aspects of the BLMS design are based on pre-existing technology, it also incorporates some novel solutions. One example of novelty in the design of the BLMS is the use of the Analogue-to-Digital converters (ADC) to increase the dynamic range of the detectors. As a general principle, it is preferable to avoid novelty in a critical system especially when the requirements can be satisfied using only proven technology. However, we believe that the novel aspects of the BLMS design have been introduced to overcome limitations of pre-existing technology. In other words, this novelty has been introduced out of necessity than merely for the sake of novelty. Novelty usually entails some uncertainty which, in turn, contributes to risk. However, the additional risk, if any, associated with novelty in the BLMS design appears to be warranted by increased fidelity of the system which, in turn, improves the capability of the BLMS to detect dangerous beam losses without adversely affecting availability.

4.2 Error Detection and Fault Tolerance

The BLMS includes very substantial provisions for error detection, especially the detection of errors that may occur while data is acquired by the tunnel electronics and while data is transmitted from the tunnel electronics to the surface electronics.

For example, a combination of CRC checks and the 8b/10 protocol is used for the communication of detector data by optical links from the tunnel electronics to the surface electronics. There are also redundant optical links for communication of this data from the tunnel electronics to the surface electronics. Other sections of this report specifically address the use of physical redundancy in the BLMS, which also contributes to the ability of the BLMS to detect errors.

The BLMS includes a limited measure of fault tolerance. One example is the triplication of the the tunnel electronics with 2 out of 3 voting. Another example is dynamic selection by the surface electronics of the error-free data received from the tunnel electronics when an error is detected in the data on one of the two redundant optical connections. We find this use of fault tolerance in the BLMS appropriate because it is very unlikely to have an adverse effect on the ability of the BLMS to request a beam dump in response to a dangerous loss. However, it will avoid some unnecessary interruptions to the operation of the LHC.

4.3 Redundancy of the Critical Datapath

In the design of critical systems, redundancy is a very common strategy for mitigating risks that are associated with component failures. In particular, redundancy may be used to eliminate the possibility of a single point of failure in the implementation of a critical system.

For a considerable portion of the critical datapath that contributes to the generation of the beam permit, redundancy is achieved by means of physical replication. For example, there are two separate optical fibre connections from the tunnel electronics to the surface electronics for each channel. In addition to replication, other forms of redundancy may be found in the BLMS

design. For example, the use of CRC checking is a form of redundancy known as informational redundancy.

We note that some elements of the critical datapath for beam permit generation are not physically replicated.

One obvious example is the fact that every single detector is uniquely positioned to detect a loss in a particular zone. There is a possibility that a dangerous loss could be very localized and only visible to a single detector. For this reason, it is not possible to claim that redundancy fully exists in the form of physical replication of the detectors. Concern that the detectors are single points of failure is somewhat dimished by the fact that these detectors are relatively simple devices based on a proven design re-used from earlier particle accelerators. As well, certain kinds of failures are likely to be detected by the BLMS connectivity test which, according to current operating procedures, must be performed at least once within the 24 hour period prior to an injection of beam into the LHC.

Another example is the fact that parts of the BLMTC design that implement the calculation of the running sum and comparison of the running sums to threshold values are not physically replicated. Clearly, the decision not to replicate these parts of the design is partially due to the size limitations, i.e., physical replication would consume too much of the limited capacity of the FPGAs selected for the implementation of the BLMS.

It might be argued that some measure of redundancy is achieved for these particular aspects of the BLMS by means other than physical replication. For example, there is a reasonable expectation that an occurrence of a dangerous loss will likely be detected by more than one of the twelve different running sums calculated by the BLMTC FPGA. Hence, the dangerous loss might still be detected even if there is a failure in the circuitry responsible for threshold comparison of a particular running sum. However, this argument is weakened by several details. For instance, the "shorter" running sums would not necessarily yield indications of a "slow loss" that exceeds the tolerable threshold for slow losses, and hence, they would not provide full redundancy for a failure of one of the "longer" running sums. Furthermore, the circuitry for calculating each of the running sums is not independent, as the output of some of the "shorter" running sums are used as input to the calculation of some of the "longer" running sums (as explained in the remarks at the bottom of Table 6.2 in Zamantzas's doctoral dissertation).

It might also be argued that some protection against single points of failure exists in the form of both spatial and temporal redundancy. For instance, most occurrences of a dangerous loss are likely to be observed by multiple detectors - although exceptions are known to be possible. This may be regarded as a form of spatial redundancy. However, losses are generally expected to be localized, i.e., only seen by detectors in close physical proximity that are processed by the same electronics and may be susceptible to the same single point of failure. Hence, this form of spatial redundancy has limited value towards the mitigation of risk due to single points of failure. Some degree of temporal redundancy also exists due to the fact that dangerous losses

are likely to be observed over multiple turns of the beam. However, temporal redundancy has limited value in the case of fast losses that must be detected within a small number of turns.

Overall, we conclude that the design of the BLMS contains some single points of failure. The existence of these single points of failure is not an oversight or a mistake in the design. Rather, the existence of single points of failure is a consequence of practical limitations of resources. The fact that the detectors are single points of failures does not greatly concern us. However, we are less comfortable with the fact that some elements of the BLMTC FPGA datapath, such as the running sums structure, are not physically replicated. It might have been possible to avoid this situation by selecting a different type of FPGA for the BLMTC with greater capacity. Alternatively, it might also have been possible to avoid this situation by allocating non-critical functionality to a separate FPGA (as discussed elsewhere in this report).

4.4 Mixed Purposes of the BLMS

In addition to requesting beam dumps, a considerable portion of the BLMS functionality exists exclusively for the separate purpose of generating data about beam loss. This data used for purposes outside the scope of machine protection, or at least, for purposes that are not part of the critical real-time datapath that triggers a beam dump. Some of this data might be used to adapt the BLMS (e.g., adjusting threshold settings). This data may also be used for other purposes entirely separate from machine protection such as improving performance of the LHC – for example, optimization of the beam collimation system. For this reason, we regard the BLMS as a system with mixed purposes. On one hand, it is an integral part of the machine protection strategy for LHC. On the other hand, it is a measurement tool that generates data which is used for a variety of purposes. As part of our review, we considered the possibility that the obligation to satisfy the needs of these other purposes might contribute to a situation where the integrity of the safety function of the BLMS is compromised.

In this regard, we also note that a considerable portion of the BLMS functionality exists for the purpose of proof testing the BLMS. While this test functionality is very relevant to the goal of ensuring that the safety function of the BLMS is achieved, it is not part of the critical real-time datapath that triggers a beam dump. For this reason, we also considered the possibility that the obligation to include this test functionality might similarly contribute to a situation where the integrity of the safety function of the BLMS is compromised.

A concrete manifestation of the mixed purposes of the BLMS is the fact that a substantial portion of the circuitry on the BLECS FPGA (and to a lesser degree, the BLMTC FPGA) is not part of the critical real-time datapath that triggers a beam dump. The safety function of the BLMS would be unchanged if this additional circuitry did not reside on the same FPGAs as the circuitry that triggers a beam dump – for example, if this circuitry had been allocated to separate FPGAs.

We assume that the functional concept of the LHC BLMS was carried over from past experience with other accelerators, which operate at lower energy levels. We suspect that beam loss measurement has historically been viewed primarily as a matter of calibration as much or more than a matter of machine protection. It is possible that the genesis of the BLMS may have

influenced some decisions in its design, such as allowing additional circuitry to be co-located with circuitry that triggers a beam dump.

We are concerned that future decisions about the maintenance and use of the BLMS may be adversely affected if stakeholders who make or influence these decisions give priority to non-safety purposes of the BLMS. For example, a decision could be made to add another function to the BLMS for processing of beam loss data by adding more circuitry to the BLMTC FPGA. In light of our remarks in Section 4.6 about the very high utilization of logic elements in the BLECF and BLMTC FPGAs, a decision to add more circuitry to the BLMTC FPGA could jeopardize the dependability of these FPGA as integral elements of the LHC approach to machine protection. Another example would be a decision at some future point in time to relax the current practice of requiring, as a condition for injection to the LHC, that the BLMS "connectivity test" has been run within the previous 24 hours. If stakeholders regard the BLMS as primarily as a source of measurement data, then they might accept an argument that operational requirements for ensuring the integrity of the BLMS can be relaxed for runs in which the measurement data might less important.

The fact that the BLMS has mixed purposes is not a mistake. Indeed, the mixed purposes of this system are explicitly declared in [1]: "Its use, both for machine protection and for machine operations and studies is considered." However, care must be taken when changes are made to the BLMS over its lifetime that changes motivated by these other purposes are not given priority over the goal of preserving the integrity of the BLMS safety function. We are especially concerned about pressure being applied by users of BLMS data in the LHC community for changes to the BLMS that would increase the size or complexity of the BLMS functionality for purposes other than machine protection.

Recommendation #1: There should be a clear policy statement by the Director for Accelerators and Technology that changes to the BLMS design for the purpose of enhancing or modifying the capability of the BLMS to provide measurements of beam loss must not compromise the safety of the BLMS.

4.5 Separation of Critical and Non-critical Parts

With respect to the protection of the LHC machine, the generation of the user permit is the only critical function of the BLMS. We considered other functionality of the BLMS, as described above in Section 4.4, to be "non-critical" in the sense that it is not part of the real-time datapath from the threshold comparators in the BLMTC to the connection point with the BIS interface. When a system contains a combination of critical and non-critical functionality, the critical functionality should be isolated from the non-critical functionality as a general principle for the design of highly dependable systems. The aim of this general principle is to minimize the possibility that the implementation of the non-critical functionality could adversely affect the implementation of the critical functionality. As well, the separation of critical and non-criticality functionality simplifies the task of demonstrating the integrity of the safety function.

We are concerned that the design of the BLMS does not sufficiently isolate the generation of the user permit from the generation of measurement data. One aspect of this concern is due to

the potential consequences of the very high utilization of logic elements on the BMLTC FPGA, as discussed in Section 4.6. Another aspect of our concern arises from our examination of the VHDL code for instances of signals that connect parts of the design that appear to be non-critical to parts of the design responsible for generating the user permit.

At a high level of abstraction, the critical function of the BLMS may be viewed as a datapath that extends from the detectors to the connection of the user permits to the BIS. At various points in this datapath, there are branches off to elements of the design that implement non-critical functionality. However, there should not be any signal that flows from one of these non-critical branches back into the datapath for generation of the user permit.

We have spent a significant portion of our review effort on an examination of the signal connections between portions of the BLMS circuitry that we believe to be critical and portions that we believe to be non-critical. In spite of this effort, we are unable to decide conclusively whether there are signals that connect non-critical portions to critical portions. To the best of our understanding, there is no instance of a signal from a non-critical portion of the design that could mask, delay or otherwise interfere with a beam dump request. However, our understanding falls short of absolute certainty. One of the difficulties in making a conclusion about adequate separation of critical and non-critical portions is that we must "reverse engineer" the design from VHDL code and schematics to distinguish between critical and non-critical portions. With more documentation, such as a requirements specification, more time and access to tools used by the developers, it would be possible to decide this question. However, the more fundamental issue is the fact that this separation between critical and non-critical portion is not an explicit feature of the design. The separation should not only exist, but exist in a way that is easily traceable to a differentiation between critical and non-critical functions at the requirements level.

A number of important industry standards for safety critical functionality require explicit evidence that such separation exists. In the case of the functionality implemented by the FPGAs, this separation would be most easily demonstrated by allocating all of non-critical functionality to a separate FPGA. (This would also address another concern described in Section 4.6 regarding the high utilization of logic elements in the BLECF and BLMTC FPGAs.) Even though this would entail some duplication of the circuitry also needed for the critical functionality, the cost of this duplication is insignificant relative to the potential cost of catastrophic damage to the LHC.

Recommendation #2: The engineering documentation for the BLMS design should be enhanced to distinguish critical from non-critical elements of the design, and to record an evidence-based argument that the critical functionality is sufficiently isolated from the non-critical functionality (if such an argument can be made).

The above recommendation is intended to serve as a means of increasing confidence in the design of the BLMS with respect to its role in the protection of the LHC machine. This recommendation is also intended to support the maintenance of the BLMS system, i.e., to make

system maintainers aware of changes to the design that could impact the dependability of the BLMS with respect to its role as a key element of the LHC machine protection strategy.

4.6 Utilization of FPGA Elements

The functionality of the BLECF, BLMTC and BLECS FPGAs, as synthesized from a combination of schematic diagrams, pre-defined cores and VHDL code, consumes a surprisingly large portion of the discrete elements available on these FPGAs. In theory, all of these discrete elements are available for use. However, there are several factors that effectively limit what portion of these elements may be used without being exposed to certain risks that could conceivably affect the reliability of the BLMS or limit its maintainability.

It is understood that the utilization of the BLECF FPGA is 85% and for the BLMTC FPGA this value exceeds 90%. We are not certain if the manufacturer of the FPGAs officially provides guidance, in terms of a percentage, on how much of the FPGA can be used. However, the very high utilization of logic elements on these FPGAs exceeds what we believe is considered good practice for a critical system that incorporates FPGAs – or at least, it is at a level of utilization that warrants careful scrutiny and consideration for the future evolution of the BLMS.

In general, very high utilization of the discrete elements on a FPGA might increase exposure to certain risks including problems with heat dissipation and an inability to achieve the required operating frequency of the FPGA. We are not aware of the extent to which the particular FPGA used in the BLMS are susceptible to these specific risks – and it is possible that the manufacturers of these FPGAs might not even explicitly provide such information.

As remarked earlier in this report, a significant portion of the BLMS functionality exists exclusively for the purpose of generating measurement data. If this extra functionality had been allocated to separate FPGAs instead of being co-located with the functionality needed for the generation of the user permit, then the utilization of the FPGA would almost certainly be well below any reasonable threshold for concern. Our concern in this regard could have been avoided by a decision to allocate separate hardware for the surface electronics for the purpose of implementing non-critical functionality. This would have also isolated the non-critical functionality from the critical functionality which, as previously mentioned, is desirable as a means of minimizing the possibility that non-critical functionality might interfere with critical functionality. Obviously, an even simpler solution might have been to use a larger FPGA – although this may have been seen as wasteful at the time when the decision to select the FPGAs for the BLMS was made.

It is very clear from our discussions with the BLMS developers that they are fully aware of the possible implications of very high FPGA utilization. Moreover, we are aware of refinements to the design of the BLMTC FPGA that are intended to mitigate some of the risks generally associated with high utilization of the FPGA. For example, the manner in which elements of the "running sums" function in the BLMTC FPGA are updated reduces the number of logic elements that may switch logic state in a particular clock cycle.

We also understand that the synthesis tools provided by the FPGA manufacturer include verification functions that check whether the operating frequency specified by the BLMS designers are achievable – and we have been assured that these tools have been used by the BLMS developers to verify that the FPGAs can operate at a frequency well above the frequency actually used in the system.

Thus, our initial concern about the immediate implications of the high level of FPGA utilization in the BLMS has been addressed by a combination of engineering ingenuity and thoroughness. Nevertheless, this high level of utilization has very significant implications for the evolution of the BLMS over its lifetime. Obviously, this high level of utilization severely limits the possibility of adding more functionality to these FPGAs if, for example, stakeholders decide that another kind of measurement function is needed. Given the already high level of utilization, the routing of signals to accommodate additional functionality will likely be difficult and the maximum operating frequency may be reduced. As well, phenomena such as Negative Bias Temperature Instability (NBTI) may change the device physics over the expected lifetime of the BLMS and this could also reduce the maximum operating frequency.

Recommendation #3: Any proposal to modify the BLMS in a manner that would increase the amount of utilization of the FPGAs should be strongly resisted unless some previous change to the BLMS has reduced the utilization to a level that easily accommodates additional functionality without increasing the utilization back to a very high level.

Recommendation #4: Options to reduce the utilization of the FPGAs should be identified and evaluated.

One obvious possibility is to move non-critical functionality to a separate FPGA. If this option is selected, it should be implemented while the developers of the BLMS are available to implement a change that involves de-coupling critical from non-critical functionality since this task might depend on important details that are not documented. Moving non-critical functionality to separate hardware would also address concerns about insufficient isolation of the critical functionality from the non-critical functionality, as stated in Section 4.5.

4.7 Successive Running Sums Calculation

A significant portion of our review effort focused on the calculation of the successive running sums for each detector by the BMLTC FPGA. In our opinion, this calculation is the most complex aspect of the data processing performed by the BLMS. While this calculation is quite simple from a functional point of view, implementation of this function in the BLMTC is nothing short of *a tour de force* in algorithm and circuit design.

We suspect that the complexity of this part of the BLMTC is partially due to concerns about the high utilization of logic elements on the BLMTC FPGA. In other words, some of this complexity might have been avoided if earlier decisions about the design of the BLMS had not resulted in such high levels of utilization. For example, it is possible that some of this complexity could have been avoided if a different kind of FPGA with a larger capacity had been selected.

The running sum calculation is a single point of failure, as there is no physical replication of the circuitry that implements this calculation for each detector. Moreover, the different running sums for a single detector are not completely independent. In particular, an error in the running sum calculation for one of the "shorter" running sums would almost certainly cause an error in the running sum calculation for one of the "longer" running sums for the same detector.

We are satisfied that the calculation of the successive running sums has been verified very extensively by means of simulation and testing. However, the complexity and criticality of this calculation beckons for an analytical argument about the correctness of this calculation in terms of the register-transfer level structures used to implement this design and the non-trivial timing dependencies between these structures.

4.8 Requirements Specification

While there is a substantial volume of documentation about the BLMS in the form of doctoral dissertations, journal articles and conference papers/presentations, there is no comprehensive and current specification of the functional requirements of the BLMS. The developers of the BLMS and perhaps also a few other member of the CERN technical staff with a keen interest in machine protection have a detailed understanding of the current functionality of the BLMS. However, this understanding cannot be relied upon for the long term maintenance of the BLMS as members of the development team will inevitably retire or move to other projects and responsibilities over the lifetime of the system.

There is a project document titled "Beam Loss Monitor Specification" [1] which is self-described as a functional specification of the BLMS. However, the version of this document provided to us for the purposes of this review was last updated in 2004 which is clearly well before the functionality of the BLMS was finalized. This document has more of the flavor of what is typically known in industry as a "system concept" or "white paper", in contrast to an enumeration of functional requirements that specify the response of the BLMS to specific stimuli and other events. This project document is undoubtedly a key artifact of the development process, especially as a source of background information about the BLMS. It exhibits the same "deep thinking" that made a strong impression on us during the review of the BIS. Echoing remarks contained in our October 2009 report on the BIS, we appreciate that the research-oriented culture of CERN is different from the culture of our industry-oriented clients. We are also impressed by the richness of the various technical documents that have been produced in the course of developing the BLMS. Nevertheless, we are concerned that a comprehensive record of the current behaviour of the BLMS does not exist except in the minds of its developers.

For complex digital systems that implement critical functionality such as the BLMS, a comprehensive specification of the requirements should be considerably more detailed than a high level statement of the system concept or purpose. For instance, it should provide explicit details about the required behaviour of the system in response to anomalous conditions, e.g., it should explicitly specify the required response of the BLMS when it does not have a valid indication of the current beam energy (which is used to select the appropriate thresholds for

detecting dangerous beam loss). One possibility is a fail-safe response to this anomalous condition, i.e., a beam dump request is generated. Another possibility is a fail-operational response in which a default value is used for the beam energy such that the most conservative threshold settings are used. This particular example of an invalid indication of beam energy is one of many anomalous conditions that we believe are detectable by the functionality of the BLMS. But in the absence of a comprehensive specification of its functionality, we do not have a means of determining exactly what conditions are detectable and what response is required for each of these conditions, i.e., fail-safe vs. fail-operational. Although we could resolve such uncertainties by analysis of the VHDL code, reverse-engineering of the design would defeat the purpose of documenting the required behaviour separately from the design. Alternatively, we could ask the BLMS developers about the behaviour of the BLMS in such situations. However, this approach to clarifying the behaviour of the BLMS is not sustainable over the lifetime of the BLMS.

As explained in other sections of this report, the absence of a comprehensive specification of the BLMS requirements has a number of undesired implications with regard to the verification of the BLMS design and the long-term maintainability of the BLMS. Even a modest effort could produce a comprehensive specification of the functional requirements for the BLMS that would be sufficient for the maintenance of the system.

Recommendation #5: A comprehensive specification of the required functional behaviour of the BLMS should be created and maintained using a style and format that is amenable to modification as changes are made over its lifetime.

4.9 VHDL Code

A very large portion of the functionality of the BLMS is determined by the VHDL code which is synthesized to configure the FPGAs used in the tunnel and surface electronics. CSL has examined a substantial portion of the VHDL code developed by CERN for the BLMS, with particularly emphasis on the critical datapath that propagates a beam dump request from the threshold comparator to the connection with the BIS. We examined this code in regard to the both correctness and understandability. As previously explained in Section 4.5, we also examined the VHDL code in regard to the separation of critical functionality from non-critical functionality.

We have observed nothing that would make us doubt or question the correctness of the VHDL code. The fact that we do not have a comprehensive specification of the required behaviour of the BLMS has limited our ability to assess the correctness of some specific details of the VHDL code, especially details of the required behaviour that are adjunct to the primary objective of requesting a beam dump when a dangerous loss is detected, e.g., data logging.

In addition to the purely “logical” aspects of the code, we have examined more “circuit level” details such as protection from clock slew and reset mechanisms. Several questions and observations about these details have been informally discussed with members of the BLMS development team, who have consistently provided knowledgeable answers and feedback that support our confidence in the correctness of the VHDL code.

In addition to correctness, we examined the VHDL code in regard to its understandability. This is especially important in regard to the likelihood that changes will be made to the VHDL code over the lifetime of the BLMS and that some of these changes, especially with the passage of time, will be made without the involvement of the individuals who originally developed this code.

There are several opportunities to improve the understandability of the VHDL code. Many parts of the VHDL code for the BLMS would be significantly more understandable with the addition of more in-line comments, and also, by ensuring that the comments are truly meaningful. The dearth of meaningful comments about instantiations of Intellectual Property (IP) in the design (i.e., Altera Megafunctions) is also a serious limitation on the understandability of the VHDL code. The understandability of the VHDL code would also be improved by the adoption of a set of naming conventions, and conformance to these conventions. For example, the names of signals should clearly distinguish "active-low" signals from "active-high" signals.

Recommendation #6: The understandability of the VHDL code should be significantly improved with the addition of meaningful in-line comments.

4.10 Design Verification

An impressively diverse combination of test methods and strategies has been used to verify the design of the BLMS. The verification activities include both analysis and testing. It is very evident from presentations and discussions that the BLMS development team has made extensive efforts to analyze and test the design of the BLMS. However, it is difficult to objectively assess the adequacy of the verification results.

The task of verifying the BLMS is complicated by several factors, some of which were already mentioned in the introductory section of this report. The verification task is also complicated by the fact that the BLMS is not "purely" digital; the front end electronics include analogue elements whose behaviour cannot be adequately described in terms of "if-then-else" logic. As well, the correctness of the BLMS must ultimately be assessed in terms of how it responds to a phenomenon (i.e., a beam of protons or lead ions at unprecedented energy levels) whose behaviour is not completely predictable. In light of such complexities, the task of thoroughly verifying the BLMS cannot be limited to a strictly formal process that mechanically cycles through an itemized list of functional requirements. Instead, there is a need for some measure of exploratory testing combined with expert knowledge to reach a credible conclusion about the correctness of the design. Moreover, the results of verifying the design of the BLMS are not necessarily "binary" (i.e., "pass" or "fail"). Some of the results may be used to make improvements or adjustments to the BLMS (e.g., adjusting thresholds) even though the observed behaviour is considered "correct".

The lack of emphasis on formal verification (e.g., production of written test procedures with detailed "step by step" instructions accompanied by a written record of the "pass/fail" results of each step) limits our ability, as reviewers, to form an objective conclusion about the adequacy

of the effort to verify the design of the BLMS. Instead, our confidence in the adequacy of these verification results rests very much on our perception of the dedication and skill of the individuals who performed these verification activities, and our belief that these individuals possess a very detailed understanding of the required behaviour of the BLMS.

While recognizing that a completely formal approach to verification is not appropriate for the BLMS, we are obliged to consider how the approach taken by CERN differs from what commonly done by industry to verify a critical system. In particular, we focus on three aspects of verification that would be emphasized typically for a critical system, namely, (1) an objective measurement of coverage, (2) objective criteria for deciding whether a particular requirement has been correctly implemented and (3) independence.

While substantial and extraordinary diverse, it is difficult to be sure that the verification of the BLMS is comprehensive in terms of an objective measurement of coverage. It would be typical in an industry project of this level of criticality to see an explicit record of the traceability between requirements and verification results, e.g., a table that identifies the method of verification and the corresponding verification results for each requirement. But as discussed earlier in Section 4.8, a comprehensive specification of the functional requirements does not exist for the BLMS. Hence, it is not possible to record traceability or even to make an objective measurement about the coverage achieved. For example, the BLMTC and BLECS FPGAs include a number of error detection functions which must generate a beam dump request if an error is detected. Other than reverse-engineering of the VHDL code, we are not aware of any way to make a list of each of the error detection functions that should have been implemented by the code. Without such a list, it is not even meaningful to ask if all of the required error detection functions have been implemented in the VHDL code.

Since there is no comprehensive specification of the required behaviour, there appears to be no objective and repeatable basis for deciding what needs to be tested or what criteria should be used to decide that the result of executing the test case is correct, other than the knowledge and insights of the individuals who designed the system. We have noted that references to specific sections of design documents are embedded in at least one of the commissioning test procedures for the BLMS; these references are presumably intended to clarify what should be observed when particular test cases are performed. However, this is not comprehensive. For example, Section 10.3.3 of the test procedure for commissioning the BMLS only refers to “beam permit transmission from all threshold comparators to the last combiner card” as a topic to be addressed by this test without any detail on what needs to be done. As reviewers, we are left with no indication of what, if anything, was done beyond simply checking that a beam permit exists at the last BLECS card when no dangerous loss has been detected. For instance, we cannot determine from the information available if the “negative case” of beam permit transmission is part of this particular step in the test procedure, i.e., does this step also include a demonstration that the beam permit is not transmitted (i.e., a beam dump is requested) for any one of several dozen reasons – and if so, how many of these different reasons are explicitly tested?

We have also considered the extent to which the BLMS has been independently verified. We are aware that the BLMS was reviewed by independent auditors in 2008. We also aware that highly qualified individuals from within the LHC Machine Protection community (but outside the BLMS development team) have participated in the review of various artifacts such as test procedures for commissioning. However, we do not see evidence of independence directly in the process of developing test procedures and the execution of these procedures. On the contrary, we believe that developers of different parts of the BLMS have been directly involved in the verification of those parts that they have developed and quite possibility they are the only people who have had a substantial role in the verification of these parts. This falls short of the best practices of industry which generally entail an expectation that verification of a critical system is performed independently, i.e., by individuals other than the developers. However, the unavailability of a comprehensive specification of the intended behaviour is a limiting factor here since it would be difficult for an independent party to verify a design if they do not have a description of its intended functionality.

As noted above, our confidence in the adequacy of the existing verification results for the BLMS, as it exists today, rests very much on our perception of the dedication and skill of the individuals who performed these verification activities. Indeed, we have an extremely high regard for the dedication and skill of these individuals. However, it is inevitable that the availability of these individuals to perform or support verification of the BLMS will diminish over time. While other individuals are likely to acquire sufficient understanding of the BLMS over time to perform testing activities, they will be at a disadvantage without a comprehensive specification of the required behaviour.

In general, the BLMS development team appears to have given priority to the automation of design testing over documentation. There are certainly some important benefits of test automation, especially in terms of regression testing for the purpose of checking that a change to the system has not had unintended effects on the behaviour of the system. Nevertheless, we are concerned about the ability to maintain test scripts over the lifetime of the system when changes are made to the system or the test environment. It may be too much to assume that the individuals responsible for implementing a change to the BLMS will be able to reverse-engineer an automated test procedure to take account of the how the functionality of the system has changed, especially when they do not have the benefit of a written specification of the system.

On the assumption that automation will continue to be given priority, a Master Verification Plan should be created to provide system maintainers with an overall understanding of what kinds of verification activities should be performed when a change is made to the design of a particular aspect of the system.

Recommendation #7: A Master Verification plan should be created to document the verification approach from a top-down perspective, e.g., what tests have been allocated to each component of the system. This document will guide system maintainers when changes are made to the system.

The Master Verification Plan should include guidance on how a suite of verification activities for a particular change should be defined based on the nature of the change and its potential impact on other parts of the system.

4.11 Proof Testing

The BLMS developers have established a very thorough approach to proof testing for the BLMS. In this context, the term "proof testing" refers to testing activities intended to reveal defects that arise during manufacturing or installation, damage, deterioration, or other unintended changes to the BLMS. This form of testing is distinct from design verification which is concerned with errors or other shortcomings in the design of the system. (It is possible however that proof testing might reveal certain kinds of those errors that design verification should reveal, and vice versa.)

One instance of proof testing for the BLMS is known as the "connectivity test". The operational procedure for initiating a new fill (i.e., an injection of particles into the LHC) currently requires this connectivity test to be performed sometime within the preceding 24 hours. Beam injection is blocked if non-conformities are found. This is a fully automated test that should reveal any loss of connectivity between the detectors and the BLMTC FPGA. In addition to demonstrating that a connection exists, the connectivity test provides a degree of qualitative data about the performance of the detectors. For example, the connectivity test has revealed defective solder connections in the detectors that affected the response of the detectors.

It might be possible for an operator in the control room to override or by-pass the portion of the operational procedure that ensures that the connectivity test has been performed. We believe that this check is implemented in an executable script which avoids the possibility that an operator can simply neglect to manually check this condition. However, we are not aware of any controls that would prevent a single individual from modifying the script so that this check is not performed. This modification may be unintentional or performed without a full understanding of the potential consequences.

Recommendation #8: A formal policy should be established to impose limits on the ability of a single individual to modify the procedure for starting the accelerator in any way that would override, inhibit or otherwise interfere with the performance of the connectivity test.

Another form of proof testing used during development involves exposing each detector to a radioactive source and ensuring that this exposure results in an expected response. This form of proof testing is called the "cesium test" because the radioactive source used is cesium. Whereas the connectivity test uses high voltage to induce a current in the detectors, the cesium test is more "authentic" because it is testing the response of the detectors to particles. Unlike the connectivity testing which is automatically performed, the cesium test is a manual procedure that requires the radioactive source to be physically moved to the location of each detector. This is a very time consuming procedure that can only be performed when the LHC is not operational and underground access is possible. Due to the amount of manual effort involved, it does not seem likely that the cesium test will always be performed during shutdown periods when access underground is allowed. The fact that the cesium test is performed much less

regularly than the connectivity test (and might never be performed on a regular basis) does not greatly concern us for the following reasons:

1. The detectors are based on simple technology with a demonstrated record of high reliability; in particular, the same design for the detectors has been used for many years in the SPS at CERN.

2. Most failure modes of the detector involve failures that would also be revealed by the results of the connectivity test, e.g., a defective solder connection.

3. Most dangerous losses will be detected by more than one detector, and therefore, the failure of a single detector is an unlikely cause of a missed beam dump.

We also aware of other forms of proof testing such as the internal beam permit check on the transmission of beam dump requests with crates and between crates. However, we are unsure whether a schedule has been established for these other forms of proof testing and we suspect that some of this additional proof testing may only be performed discretionally.

Recommendation #9: A schedule for each form of proof testing should be defined along with a means by which operators would know if some form of proof testing has not been performed according to schedule, or if the results indicate an unresolved problem.

4.12 Known Limitations

The BLMS developers are aware of some limitations of the BLMS. We understand that one such limitation is the dynamic range of the detectors close to the injection sites of the LHC. We are confident that these limitations are being actively addressed by the BLMS developers.

4.13 Critical Data Settings

The BLMTC FPGA compares the particle losses reported by each detector over various durations to a set of specified threshold values. A beam dump request is generated when one of these thresholds values is exceeded. The operators in the control room can change thresholds within a range with an upper bound determined by values in a master database maintained by the BLMS developers. The threshold values are maintained in a different database and loaded into the BLMTC FPGA. The threshold values are defined in groups rather than individually. Hence, a change to a single value in the database would change the threshold used for multiple detectors. It is intended that the upper bound on the threshold values will limit the range of values that may be selected by the operators to safe values, i.e., below the level of a dangerous loss.

It is conceivable that a threshold value erroneously set too high could result in a failure of the BLMS to request a beam dump when a dangerous loss occurs. This could be caused by a relatively simple mistake such as an incorrect keystroke, e.g., typing “0” instead of “9” when modifying the database that sets the upper bound on the threshold values. Visual inspection is not a reliable method of detecting such errors, especially if the inspection is performed by the same individual performing the modification.

We are not aware of any controls that would limit the possibility of such errors, or the severity of such errors. In particular, the BLMS does not use relatively simple safeguards such as limiting how much the value of a threshold setting can be increased by a single human interaction. Such safeguards are commonly used in other critical systems and we are not aware of any reason why they could not also be used for the BLMS, other than the extra effort and time that would be required to make some changes.

In addition to threshold settings, we have similar concerns about the lack of safeguards to protect other critical data such as settings which determine whether individual detectors are masked or disabled.

Parenthetically, we suggest that the lack of safeguards limiting the extent to which an individual can change the master database in a single interaction is a potential security vulnerability.

Recommendation #10: Safeguards should be added to the tools used to modify critical data to reduce the likelihood of simple human errors such as incorrect keystrokes.

We note that a previous review of the BLMS in 2008 generated a similar suggestion that "an application should be deployed that provides means to minimize the introduction [of] erroneous values to this table, e.g., through human errors", which was motivated by a similar concern about the lack of safeguards.

4.14 Maintainability

Previous sections of this report have expressed concerns about the maintainability of the BLMS over its planned lifetime. These concerns include:

- The lack of explicit separation between parts of the system that implement critical functionality and parts that implement non-critical functionality increases exposure to the possibility that a change to the design of a non-critical part might unintentionally have an adverse effect on a critical part.

- The high utilization of logic elements in the FPGAs of the BLMS will limit what changes can be made to the configuration of the FPGA (e.g., the addition of new functionality).

- The dearth of meaningful comments in the VHDL code and other problems related to the understandability of the VHDL code will increase the difficulty of making changes to the VHDL code.

- The absence of a comprehensive specification of the BLMS requirements will increase the difficulty of maintaining the BLMS especially when changes to the design are needed.

- The lack of traceability data from higher levels of documentation down to the VHDL code will make it more difficult to assess the impact of a change to the design of the BLMS.

- The BLMS does not incorporate safeguards that would reduce likelihood of human errors when modifying critical data.

It is widely recognized in the discipline of system safety that maintenance activities have frequently resulted in unintended changes to a system (or have failed to prevent unintended changes to a system) that have resulted in accidents. The likelihood of such problems could be reduced by resolving the above-mentioned concerns. A number of recommendations and suggestions presented in earlier sections of this report are intended to resolve some of these concerns.

4.15 Operational Monitoring

During operation of the LHC, we have wondered how operators and system maintainers will become aware of a problem in the BLMS. We have seen evidence (in a presentation) of a software tool that appears to report the status of various forms of proof testing. However, the manner in which this tool will be used is not clear to us. Will it only be used when an operator or system maintainer decides to use the tool? Or does some operational procedure necessitate the use of this tool?

Similarly, we are uncertain about how operators and system maintainers will become aware of a problem in the BLMS that is indicated by anomalous events recorded in log files generated by the BLMS. Some anomalous events, such as corruption of a single message on one of the two redundant optical links from the tunnel electronics to the surface electronics will not trigger a beam dump request. Is there some automatic way in which an operator or system maintainer will be alerted when a certain threshold of such events is exceeded? Otherwise, is there a maintenance procedure that would require a system maintainer to periodically analyze the log files for indications of such problems?

It is very evident that the BLMS developers have created a powerful set of tools for proof testing and diagnosis. However, it unclear how these tools will be effectively used as part of a systematic approach to maintenance of the BLMS.

Recommendation #11: A comprehensive procedure for monitoring the status of the BLMS should be defined to ensure that operators and system support maintainers will become aware of faults and other problems in a systematic and timely manner. This procedure should include review activities to be performed after each LHC run.

Beyond consideration of whatever automation is involved in the operational monitoring of the BLMS, the above recommendation refers primarily to the level at which human awareness and action is necessary.

4.16 Dependence on other Systems

Different threshold values are used by the BLMS for different levels of beam energy. Measurements of the beam energy are generated by the Beam Energy Tracking System (BETS) and distributed throughout the LHC by the General Machine Timing (GMT) link. These measurements are received by the Combiner and Survey cards of the BLMS which converts the measurement to a value indicating the level of beam energy and distributes the result to the Threshold Comparator cards.

If the beam energy value received by the BLMTC FPGA is incorrect, then the threshold values used for checking the current beam loss might be too high which, in turn, might inhibit the BLMS from requesting a beam dump in response to a dangerous loss. Hence, the safety of the BLMS is partially dependent on the reliability of other systems, namely, the BETS and GMT link. This dependency leaves the BLMS developers with an obligation to ensure that other stakeholders, in particular, the LHC Machine Protection committee, are fully informed about the extent to which the safety of the LHC depends on other systems that provide a measurement of beam energy to the BLMS.

Recommendation #12: The LHC Machine Protection Committee should determine whether the risk associated with the dependency of the BLMS on other systems has been adequately controlled.

5. Summary and Conclusion

The salient findings of this review may be briefly summarized as follows:

1. The design of the BLMS is conservative, except where novel solutions are needed to overcome limitations of known solutions.

2. The design of the BLMS includes very substantial provision for error detection. Fault tolerance is used appropriately.

3. The critical path for beam dump requests includes some single points of failure.

4. The BLMS serves other purposes besides machine protection. This might bring other interests into conflict with the safety objective of the BLMS.

5. Separation of critical parts of the design from non-critical parts is not an explicit feature of the design.

6. Very high utilization of the logic elements in the BLECF and BLMTC FPGAs could be problematic as it evolves over the lifetime of the LHC.

7. The algorithm and circuit design for the BLMTC Successive Running Sums (SRS) calculation is very complex.

8. The lack of a comprehensive specification of the BLMS requirements will likely have an adverse effect on the maintainability of the system.

9. System maintainers could have difficulties understanding the VHDL code for several reasons, especially the lack of meaningful comments.

10. While impressively diverse, it is difficult to objectively assess the adequacy of the verification results.

11. Proof testing is very substantial, but operational procedures concerned with proof testing are unclear.

12. There are known limitations of the BLMS that need to be addressed, e.g., insufficient dynamic range of the detectors near the injection sites.

13. There is a lack of safeguards to protect against human error when critical data such as threshold settings is modified.

14. There are a number of ways in which the maintainability of the system could be improved.

15. It is unclear how operators and system maintainers will become aware of problems that arise while the system is operating.

16. The safety of the BLMS depends on other systems including the BETS and GMT link.

While it is appropriate for our review to focus the reader's attention mostly on concerns and possibilities for improvement, this should not over-shadow the fact that we have formed a very favorable impression of the BLMS and the expertise of the BLMS developers. As remarked earlier, it is a remarkable accomplishment to have created a solution that addresses all of the many inherent challenges of the BLMS.

With one minor reservation, we have found no reason to be concerned that the current configuration of the BLMS might fail to request a beam dump in response to a dangerous beam loss. This conclusion is based on the following assumptions: (1) appropriate threshold settings are used; (2) the current operating procedures, including regular execution of the "connectivity test", are maintained; (3) no element of the BLMS such as individual detectors is disabled without a decision by the appropriate authority that the safety risk is acceptable; (4) known limitations of the BLMS are adequately addressed before the machine is used at higher energies; and (5) the risk associated with the dependency of the BLMS on other systems for an accurate indication of bean energy is deemed by the appropriate authority to be acceptable. Changes to any aspect of the design may invalidate this conclusion.

The one minor reservation associated with the above conclusion is the fact that the critical real-time datapath responsible for generation of a beam dump request is not fully redundant. Hence,

there are single points of failure in the design of this critical portion of the BLMS. In this regard, we are concerned mostly about non-redundant elements of this datapath within the Threshold Comparator Field Programmable Gate Array (FPGA).

The above conclusion is also supported by actual experience to date with the operation of the BLMS, especially the fact that there has not been any time when the BLMS failed to detect and respond appropriately to a beam loss that exceeded the preset threshold.

We encourage the BLMS developers and others stakeholders who share in the responsibility of machine protection for the LHC to consider the recommendations and suggestions that we have presented in this report. Finally, we suggest that the appropriate CERN authority should expect a response to all of the recommendations within six months of the receipt of the final version of this report.

Appendix A

The following recommendations are presented in Section 4 of this report:

Recommendation	Section
Recommendation #1: There should be a clear policy statement by the Director for Accelerators and Technology that changes to the BLMS design for the purpose of enhancing or modifying the capability of the BLMS to provide measurements of beam loss must not compromise the safety of the BLMS, i.e., its ability to reliably detect a dangerous loss.	4.4
Recommendation #2: The engineering documentation for the BLMS design should be enhanced to distinguish critical from non-critical elements of the design, and to record an evidence-based argument that the critical functionality is sufficiently isolated from the non-critical functionality (if such an argument can be made).	4.5
Recommendation #3: Any proposal to modify the BLMS in a manner that would increase the amount of utilization of the FPGAs should be strongly resisted unless some previous change to the BLMS has reduced the utilization to a level that easily accommodates additional functionality without increasing the utilization back to a very high level.	4.6
Recommendation #4: Options to reduce the utilization of the FPGAs should be identified and evaluated.	4.6
Recommendation #5: A comprehensive specification of the required functional behaviour of the BLMS should be created and maintained using a style and format that is amenable to modification as changes are made over its lifetime.	4.8
Recommendation #6: The understandability of the VHDL code should be significantly improved with the addition of meaningful in-line comments.	4.9
Recommendation #7: A Master Verification plan should be created to document the verification approach from a top-down perspective, e.g., what tests have been allocated to each component of the system. This document will guide system maintainers when changes are made to the system.	4.10
Recommendation #8: A formal policy should be established to impose limits on the ability of a single individual to modify the procedure for starting the accelerator in any way that would override, inhibit or otherwise interfere with the performance of the connectivity test.	4.11
Recommendation #9: A schedule for each form of proof testing should be defined along with a means by which operators would know if some form of proof testing has not been performed according to schedule, or if the results indicate an unresolved problem.	4.11
Recommendation #10: Safeguards should be added to the tools used to modify critical data to reduce the likelihood of simple human errors such as incorrect keystrokes.	4.13
Recommendation #11: A comprehensive procedure for monitoring the status of	4.15

the BLMS should be defined to ensure that operators and system support maintainers will become aware of faults and other problems in a systematic and timely manner. This procedure should include review activities to be performed after each LHC run.	
Recommendation #12: The LHC Machine Protection Committee should determine whether the risk associated with the dependency of the BLMS on other systems has been adequately controlled.	4.16