After LS1: MAGNET POWERING WITH ZERO DOWNTIME – A DREAM?

Markus Zerlauth, CERN, Geneva, Switzerland

Abstract

Despite a number of improvements already applied in the course of the year, the magnet powering system of the LHC still accounts for around 50% of the premature beam dumps. This number might even further increase when moving to higher beam energies in the next years. With mitigations of radiation effects and the prospects for beam induced magnet quenches being discussed elsewhere, we aim at identifying possible mid- and long-term improvements within the various equipment systems to further reduce the number of equipment failures leading to a loss of the particle beams. Amongst others, this includes the sensitivity of equipment to external causes such as electromagnetic perturbations or perturbations on the electrical network. To conclude, the gain of the identified mitigations will have to be balanced against the potential impact on schedule and cost.

LHC MAGNET POWERING SYSTEM

The LHC magnet powering system represents a considerable part of the complexity of the LHC. A total of more than 1600 electrical circuits are needed to power the 10000 superconducting and normal conducting magnets installed in the LHC tunnel. In addition to the magnet and quench protection systems, the required infrastructure of powering equipment, cryogenics and electrical distribution will account for many 10000 interlock conditions capable of dumping the beam in case of equipment malfunctioning or failures in the magnet powering. Since the first Hardware Commissioning Campaigns in 2005, more than six years of operational data and experience with the magnet powering system have been acquired. While machine protection in case of powering failures is not only assured by the corresponding powering interlock systems but also through redundant protection by the beam loss monitors, the availability of the magnet powering system has been due to its complexity of primary concern since the very beginning of LHC operation.

Despite numerous improvements in the different equipment systems, the magnet powering system still accounts for the largest fraction of premature beam dumps, with 35% of the fills at 3.5TeV being aborted in 2010 and 46% in 2011. The increase in 2011 will be explained later in this paper in more detail, and is mostly correlated with the increase of luminosity and beam energy with respect to the operational year 2010. In addition the downtime following failures in the magnet powering system is often considerably longer than for other systems (e.g. after a failure of the cryogenic system). Analysing the 2011 dumps, four of the 5 most prominent causes of the beam dumps from physics energy are originating in the magnet powering system, namely from the Quench Protection System, Cryogenics, Power Converters and the Electrical Network [1]. Summing up the total physics time lost due to events in the magnet powering system one arrives at a total downtime of 35 days. With an average production rate of ~ 0.1 fb^{-1}/day in 2011 these additional days would have resulted in a considerable increase of luminosity production, highlighting the importance of additional efforts to increase the availability of the various equipment systems.

FAULT DEPENDENCIES

Failures in the powering system are showing a strong dependence on the energy level. While almost half of the premature dumps at 3.5TeV originate from the magnet powering, this number reduces to some 10% at injection level as shown in Figure 1 for the operational years 2010 and 2011. This is mainly due to the absence of radiation/luminosity related problems, lower loads and higher detection thresholds in e.g. the Quench Protection system.

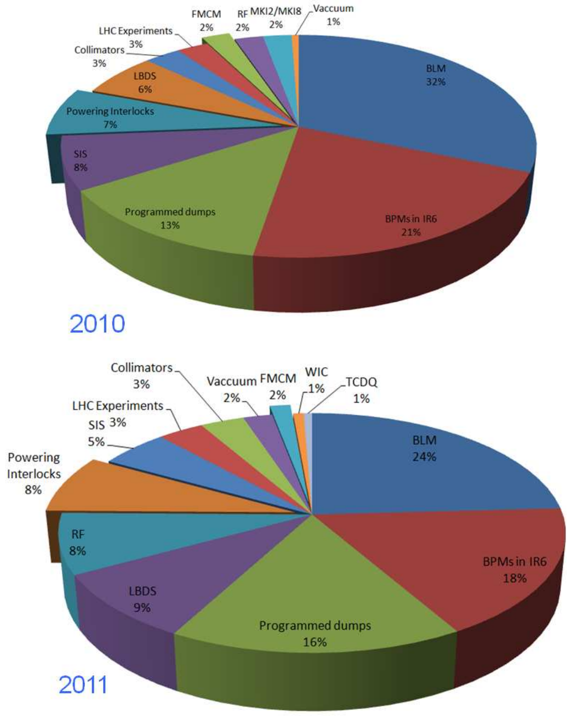

Figure 1: Causes of beam dumps for the past two operational years 2010 (left) and 2011 (right) at injection level of 450 GeV

Despite this strong energy dependence, the repartition of faults between the different equipment systems involved in magnet powering (QPS, Cryogenics, Power converters, electrical network, feedbacks, cooling and ventilation and powering interlock system) is surprisingly constant between the past operational years and between injection energy and flat top.

As for the energy, there is a clear dependency of powering failures on the total beam current present in the machine. As it can be seen from Figure 2, the density of faults throughout the year is strongly correlated to the intensity increase in the beginning of the year (until the second technical stop) and the luminosity production, especially for some systems such as the quench protection system and the powering interlock system. The fault rates appear to peak after every technical stop to decrease in the following by around a factor of two during the subsequent weeks of physics time. The much improved availability visible during the ion run at the end of the year confirms the potential gain of the R2E mitigations in the order of a factor 2-3.

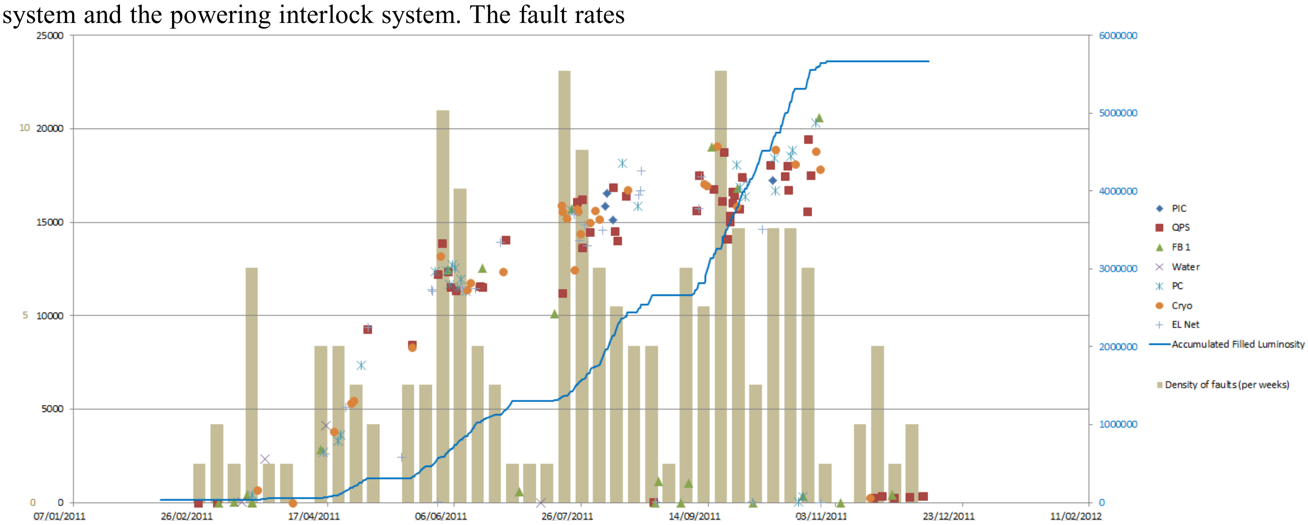

Figure 2: Beam dumps caused by the magnet powering subsystems as a function of time and intensity in the machine

SUBSYSTEM ANALYSIS

Power Converters

A number of weaknesses in the various types of power converters used for LHC magnet powering have already been identified and mitigated during the operational year 2011. This includes the re-definition of several internal fault states to become warning states (done during the 2010/11 Christmas stop), the resolution of an issue with air in the water cooling circuits of the main dipoles (resulting in cooling problems on the thyristor bridges and bypass diodes) as well as the deployment of a new FGC software version to increase the tolerance of these sensitive components to radiation effects. The remaining 26 faults recorded during physics operation at 3.5TeV were mainly caused by failures of the controls software, a known weakness of the auxiliary power supplies of the low-current converters and failures of the power module or the associated water cooling. The development of the FGC lite, radiation tolerant Diagnostics modules which are assumed to be ready to equip all LHC power converters between LS1 and LS2 and the study of redundant power supplies for 600A circuits are expected to further increase the availability in the coming years. On the contrary operation at higher energies after LS1 is expected to slightly increase the failure rates by some 10-20%.

Cryogenics

Since the beginning of operating the cryogenic system in 2005 its availability has been a driving factor of the commissioning and operation of the magnet powering system. The system has thus been closely monitored by the cryogenic operations team and considerable efforts have gone in the consolidation and upgrade of the system with the target to minimize in particular events that would lead to long recovery times such as stops of cryogenic compressors. While stops of the these cold compressors have been the cause of only one quarter of the observed events, the majority of the events seen in the past operational year were due to quickly recoverable problems such as issues with temperature sensors, valve controllers, single event upsets and the electrical distribution. Additional actions foreseen during the present Christmas stop include the relocation of PLCs in IR4/6 and 8 into radiation free areas (ULs), an SEU mitigation for the temperature sensor cards installed in the tunnel areas as well as efforts to move towards redundant PLC architectures after LS1. Further long-term improvements will depend on future hardware upgrades and the strategy chosen for the purchase of additional spares [2].

Quench Protection System

The quench protection system has shown to be the one suffering most from radiation induced effects during the 2011 run as shown in Figure 3 and Figure 4, triggering numerous mitigation measures that will be put in place during the present Christmas stop as well as during LS1 [3]. The considerable number of spurious triggers is another area of ongoing work and understanding of these events is hoped to be much improved through the development of additional diagnostics and monitoring tools.

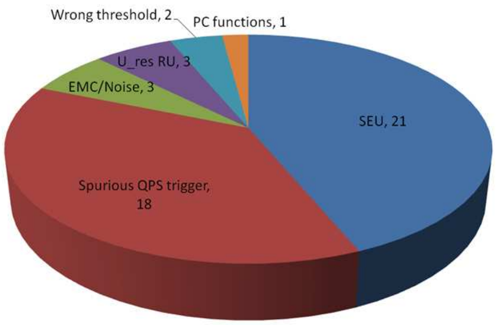

Figure 3: Origins of QPS triggers from 3.5TeV in 2011

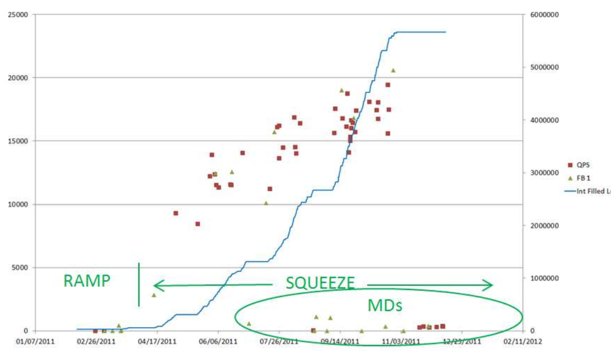

Figure 4: Distribution of QPS triggers from 3.5TeV in 2011 over time

A clear priority will be the development of a dedicated threshold management tool for the various components of the quench detection system as well as additional pre- and post-operational checks in order to assure the integrity of protection level throughout the whole operational year in perspective of the many changes and interventions necessary on this complex system. For longer-term upgrades, alternative architectures (with respect to the current 1 out of 2 architecture) should be investigated, as for example voting architectures like 2 out of 3 provide more immunity against single equipment failures, noise and SEUs whilst maintaining the required high level of safety.

Magnet Interlock Systems

The interlock systems were at the origin of five lost physics fills during the past operational year, all of them which have been traced down to spurious stops of programmable logic controllers (PLCs) of the powering interlock system, leading to a fast abort in all magnet powering circuits controlled by this PLC. All of the observed failures happened during beam operation and in areas where the PLCs are exposed to a certain amount of radiation (UJ14, UJ16 and UJ56). Although the diagnostic buffers have been retrieved and sent to the supplier for closer analysis the true origin could not be identified. In view of the circumstances and operational conditions at the time of failure they are believed to be radiation induced effects. Therefore all 10 PLCs still installed in these areas have been relocated during the TS4 and the Christmas stop to radiation free areas. An additional measure that could be envisaged to attempt a further increase of the availability of the magnet powering system is to decrease the number of corrector circuits that preventively dump the beams in case of equipment failure through the powering interlock system. Currently this configuration is done by circuit family, irrespective from the actual current used in each of the circuits during nominal operation (and thus the effect on the particle beams in case of powering failures). This has to be done with great care however as this would mean to rely to a larger extend on the beam loss monitors to ultimately capture powering failures in case the loss of a circuit has a bigger effect on the beam than expected.

Electrical Distribution Systems

The LHC magnet powering system critically depends on the quality of the mains supply. Especially the high current thyristor power converters are susceptible to perturbations on the electrical network due to the lack of intermediate energy storage in the conversion stage. The majority of the events happen during the summer period, when lightning strikes on high voltage lines outside the CERN network result in voltage dips of typically some 10-20% on a single phase of the AC input of the power converter. Figure 5 is showing a typical distribution of the perturbations as seen inside the CERN network. While major events (of larger voltage variation or duration) mostly result in trips of accelerator components, minor perturbations are typically caught by the LHC machine protection system (typically the Fast Magnet Current Change Monitors installed on the normal conducting magnets of the LHC and the SPS-LHC transfer lines). Mitigations to decrease the sensitivity of the powering equipment against minor network perturbations are currently being investigated, including an increased rejection by the power converter as well as a potential increase of detection threshold on magnets with lower β function and thus less impact on the circulating beams.

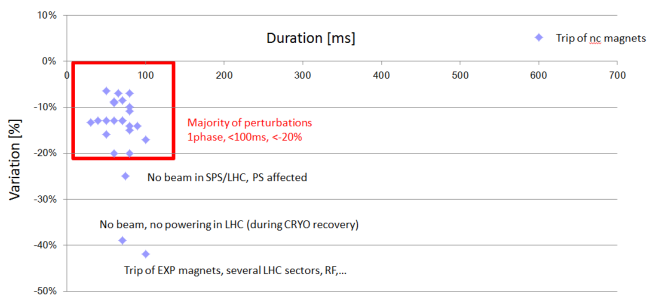

Figure 5: Distribution of typical network perturbations

CONCLUSIONS AND OUTLOOK

In view of the major contribution of the magnet powering system to the number of premature beam dumps during the past two operational years, all equipment groups are already undertaking serious efforts to further enhance the availability of their systems. Apart from a few systematic failures, most systems are already within or well below the predicated MTBF numbers, where further improvements will become increasingly costly. Failures in the magnet powering system in 2011 were dominated by radiation induced failures, and the low failure rates in early 2011 (when operating with lower beam intensity) and during the ion run indicate a considerable potential to further decrease the total failure rate. It is expected that the various mitigation measures deployed during 2011 and the present Christmas shutdown will further reduce the premature beam dumps caused by the magnet powering system in 2012 by around 30%.

Efforts for further mid- and long-term consolidations of systems to improve availability should be globally coordinated to guarantee the maximum overall gain. For this it is suggested to create a dedicated working group along the lines of the Reliability Sub Working Group which has proven of vital benefit to the development of the machine protection systems.

REFERENCES

[1] A.Macpherson, '2011 availability analysis', LHC Performance Workshop, Chamonix 2011

[2] L.Tavian, 'Cryogenic system: strategy to achieve nominal performance and reliable operation', LHC Performance Workshop, Chamonix 2011

[3] R.Denz, 'QPS upgrade and machine protection during LS1', LHC Performance Workshop, Chamonix 2011