MACHINE PROTECTION SYSTEMS

D. Wollmann*, CERN, Geneva, Switzerland

Abstract

During Run 2 the LHC reached a stored beam energy of about 300 MJ. Its safe operation requires highly reliable and complex machine protection systems. In this paper it is critically reviewed if the LHC has been running safely, the results of the regular intensity ramp-ups are analyzed, the main machine protection relevant issues during Run 2 operation are discussed and conclusions as input for Run 3 derived.

RE-CAP RUN 2 - MP VIEW

Operation with beam in Run 2 started in July 2015 with the intensity ramp-up following an intense hardware and beam commissioning campaign at the end of long shutdown 1 (LS1). 2015 was a commissioning year, which important to re-discover the LHC after LS1, to iron out left over issues in the protection and other hard- and software systems. Therefore, a conservative β^* of 80 cm was chosen. The energy stored in the two beams was slowly increased during two intensity ramp-ups. One was performed with 50 ns bunch spacing and the second one with 25 ns bunch spacing. The latter continued until the end of the proton run and was used to condition the LHC machine.

In the following year β^* was reduced to 40 cms, and, finally, reaching a minimum of 25 cm in 2018. With the deployment of the lumi server in 2017 a high level software tool was increasingly used to reduce the crossing angle and in 2018 also β^* during stable beams. This meant a dynamic change of important beam parameters, while critical hardware settings like position limits of tertairy collimators (TCT) remained quasi static. To counter balance the reduced aperture margins new software interlocks were introduced. The main once are the interlocking of the beam positions in the TCTs and TCSP via the DOROS BPMs and SIS and the interlocking of the phase advances between MKDs and TCTs via the PC interlock.

A non-exhaustive list of failure cases observed during Run 2 is given below.
Known or expected failure cases:

- UFOs
- Beam induced quenches
- Shorts in circuits
- MKI erratics
- MKD and MKB erratics
- Injection of high intensity beam in empty machine
- Injection into wrong beam
- ...

New or unexpected failure cases:

- 16L2 UFOs with fast instabilities
- Spurious firing of multiple quench heaters (following inj. losses)
- Symmetric triplet quench with orbit offset
- N_2 leaks in the beam dump blocks
- Ultra fast kicks in beam due to quench heater firing in main dipoles
- ...

However, no damage to accelerator equipment or the experiments due to beam and no damage to circuits due to powering failures or quenches has been incurred during Run 2. This required reliable machine protection systems, diverse redundancy, vigilant hardware experts, machine protection experts and OP teams.

In the following a few machine protection key events are summarized for each operational year of Run 2. Although, the choice of events was derived based on the input from different machine protection experts, it still remains subjective and has to remain non-exhaustive.

Key machine protection events in 2015

Very early after the beginning of beam operation in 2015, the LHC experienced its first asynchronous beam dump with circulating beam, a MKD erratic in B2 triggered by generator C. As there were only a few bunches stored in the machine, no beam was kicked onto the dump protection absorber TCDQ. Following this event the concerned generator was replaced. In addition the generator cleaning procedure was reviewed and improved.

The jaw material (hBN) of the injection protection absorber TDI experienced accessiv heating and outgassing. This required to limit the length of the injected batches with 25 ns bunch spacing to 144 bunches. The TDIs were replaced by new once with different jaw materials during the year-end-technical-stop (YETS) 2015/16.

At the end of the first ramp to flat top with pilot beams, the so-called unidentified-lying-object (ULO) in cell 15R8 induced important beam losses causing a quench in a neighbouring main dipole magnet. Two more quenches were caused by the ULO in the coming fills until this issue was mitigated by introducing a closed orbit bump around the ULO aperture restriction and the lowering of the BLM thresholds for close by BLMs, to avoid an excessive number of quenches in the same dipole magnet. Following the opeing of the beam vacuum during LS1, an increased rate of beam losses due to so-called un-identified flying objects (UFOs) was observed. These are macro-dust-particulates entering the beam and causing very fast beam losses up to a few milliseconds. Besides many protection dumps, UFOs also caused three quenches in main dipoles in the first year of the run. In the following the BLM thresholds were increased above the quench limit of the main dipoles to reduce the number of unnecessary protection dumps and allowing a

* daniel.wollmann@cern.ch

few UFO induced quenches for the run in 2016. Furthermore, the QPS thresholds were tightened for the individually power quadrupole magnets (IPQ) in the straight sections, to improve their protection in case of symmetric beam induced quenches. These strategies have proven to be successful and were applied for the full Run 2.

Key machine protection events in 2016

Following a few suspicious quench events, an intermittent inter-turn short was suspected in MB.A31L2 in summer 2016. Machine operation was immediately paused for powering test on the main dipole circuit in sector 12. The risk of beam induced quenches and fast power aborts in this dipoles circuit was reduced by the deactivation of the so-called Global Protection Mechanism in this sector, the reduction of BLM thresholds for the main dipoles in sector 12 and the increase of the QPS thresholds on MB.A31L2. The magnet was replaced during the extended-year-end-technical-stop (EYETS) 2016/17 and the inter-turn short was confirmed in the test station.

A N_2 leak was detected in the beam dump block in UD68. As mitigation a SIS interlock and BigSister warning was implemented and the operational pressure was iteratively adapted in the following weeks. Furthermore, intensive simulation studies were performed to define the maximum allowed intensity in case of a severe N_2 pressure drop.

Finally, an erratic trigger of the injection kickers (MKI) lead to the impact of four batches from the SPS on the TDI. The showers from the TDI into the triplet of IP2 triggered the onset of a quench in Q2.

Key machine protection events in 2017

Following a vacuum incident during the cool-down of sector 12 after the replacement of MB.A31L2, UFO like events were observed in cell 16L2 causing several quenches of the neighbouring dipole. Besides local losses, these so-called UFO type 2 events caused the fastest beam instabilities (< 10 ms) observed in the LHC so far. The reaction time of the machine protection systems proofed to be sufficient to issue protection beam dumps before critical beam losses occurred in the collimation region. To understand the origin of these events additional diagnostics - multiple beam loss monitors, diamond detectors - and a solenoid were installed around the interconnect in 16L2. Furthermore, the BLM thresholds were lowered around the affected cell to avoid unnecessary quenches of the main magnets. Finally, beam with 8 bunches followed by 4 empty bunch-slots, the so-called 8b4e scheme, were successfully used to reduce the rate of 16L2 UFO events.

Following a period with many consecutive high intensity beam dumps due to UFO type 2 events in 16L2, the beam dump block in UD62 developed a severe N_2 leak and was finally found at ambient pressure. Therefore, beam operation with high intensities was paused until a continuous flow of N_2 could be ensured by the installation of a stack of N_2 bottles. This stack was replaced by a fixed supply line from the surface during YETS 2017/18. Studies of the behaviour of the core material at high temperature in air were initiated.

In September 2017 the abort gap keeper length was made adjustable to allow the optimization of the filling patterns based on the use of bunch trains shorter than 288 bunches. Due to a wrong manipulation of the abort gap keeper parameters beam was injection into the abort gap, which has to stay free of bunches to allow for the rise of the dump kickers (MKD). To avoid future issues with this, the procedure for the change of the abort gap keeper length was reviewed and improved and the flat top length of the injection kicker (MKI) was shortened. Furthermore, an SIS check was introduced, which inhibits the injection of too long trains into the LHC.

To counterbalance the reduced aperture margins due to the reduced β^* an interlock of the phase advance between TCDQ and tertiary collimators in IP1 and IP5 was introduced. This interlock is based on the PC interlock and triggers a beam dump via the SIS, in case the defined current windows in the different quadrupole circuits are violated.

Key machine protection events in 2018

A new type of flash-over was observed in the vertical dilution kickers (MKBV) of B2, leading to a loss of kick strength equivalent to more than two dilution kicker magnets. As immediate follow-up the voltage in the MKBV.C and D of B2 were lowered by 20% to reduce the probability of the repetition of such an event. Furthermore, studies were performed to identify the new worst case failure scenarios of dilution kicker flash-overs also in the horizontal plane. In addition, it was investigated if and how the electric insulation of the high voltage bus-bars can be improved between the magnets. Finally the installation of additional two horizontal dilution kickers for the High Luminosity LHC era has been proposed.

Due to a regulation issue in the cryogenic system a symmetric triplet quench was triggered, causing a fast developing orbit offset in B1 before the beams were dumped by the BLMS in IP7. No orbit movement was observed in B2. As the quench developed fairly symmetrically it was only detected by the QPS 40 ms after the start of the event and 23 ms after the beams were dumped. The concerned triplet circuit was re-powered after the correct behaviour of the circuit protection elements had been verified. Further studies revealed that the non-symmetric behaviour of the two beams was most likely caused due to current re-distribution in the magnet during the quench.

Following the use of crystal collimators during the high-β run, multiple injections of high intensity beams were performed with crystal collimators in the beam, causing non-usual loss patters. After successful verification that no damage had occurred the normal operation was re-started. To avoid these type of events the operational procedures for MDs and special runs using prototype devices will be strengthened. Sequencer tasks will be prepared, which ensure that all prototype devices are in their out of beam position, before the injection of high intensiy beams is allowed.

MPP AND rMPP

The SPS & LHC Machine Protection Panel (MPP) cotains experts from all machine protection, protection related systems and operation. It has met 73 times after the end of LS1. During the meeting topics concerning machine protection mainly for the LHC were covered. This included parameter changes in protection systems but also near misses or issues discovered in systems relevant for machine protection. In several occasions machine protection topics concerning the SPS and Linac4 were addressed. In the latter case the topics concerned the interlocking strategy and the BIS topology in Linac4 and damage to a bellow. In addition, the MPP endorsed changes to machine protection systems and their commissioning. New system designs for machine protection and protection relevant system were validated. Furthermore, new failure cases observed in the LHC and expected in HL-LHC, either by extrapolation from LHC or due to new accelerator equipment or operational parameters were studied, their criticality analysed and the appropriate interlocking strategy proposed.

The restricted Machine Protection Panel (rMPP) is the reactive smaller brother of the MPP and contains a sub-set of machine protection system and operations experts. Its main role is to support LHC operation and equipment teams concerning protection related questions, requiring a timely response. In addition, it follows-up the intensity ramp-up and periodic check-lists, which summarize the readiness of crucial machine protection systems for the next intensity steps. Finally, the rMPP reviews machine protection critical MD requests and provides recommendations on how to minimize the machine protection risks during the MD periods. This body met 33 times since 2015. In the early part of the run the reactive nature of this body was more relevant than later, where the meetings were dominated by MD related topics.

INTENSITY RAMP-UPS

The intensity ramp-ups after long shutdowns or YETS has the purpose to identify and mitigate issues in machine protection relevant systems, which remain after the individual system tests, hardware and beam commissioning with as low as possible stored beam energy. Furthermore, they allow to identify issues and limitations related to stored beam intensity and other beam related parameters as well as to establish mitigation measures. Finally, they are important to establish the operational cycle and train the OP teams.

In the established intensity ramp-up scenario, which has been successfully used during Run 2, each intensity steps contains of free fills, a minimum of 20 hours in stable beams and is finalized by a check-list summarizing all issues in the machine protection systems during this step and verifying the readiness of all machine protection systems for the next intensity step. Due the about a factor 20 lower stored beam energy in the Ion runs, each intensity step requires two fills and 6 hours in stable beams, followed by a check-list.

The check-lists cover the behaviour in the past intensity step and the readiness for the next step of the magnet powering system, the interlocks, the RF, beam instrumentation, operation, orbit and feedbacks, the injection system, the beam dumping system, heating of accelerator equipment and the status of the machine vacuum. The check-lists are documented in EDMS.

For the re-start after other types of machine stops during a run, like technical stops, stops due to hardware issues etc. three standard ramp-up scenarios, depending on the length of the stop and the abundance and criticality of hard- and software changes, were established in 2017 and successfully applied since then.

Intensity ramp-up 2018

Figure 1 shows the intensity ramp-up after the YETS 2017/18. The beam intensity in B1 and B2 is shown in blue and red, respectively. The stored number of bunches is indicated in orange and the issuing of a check-list by a dashed vertical orange line. The date of the check-list is given in black next to the dashed vertical line. Original intensity steps of 3, 12, 72, 300, 600, 900, 1200, 1800, 2400 up to the maximum of 2550 bunches were foreseen. The first three steps are dominated by the establishing of the cycle, the intensity range from 300 to 1200 bunches is machine protection dominated and the intensity steps above 1200 bunches are dominated by beam intensity related issues.

Figure 1: Intensity Ramp up 2018: the beam intensity in B1 and B2 is shown in blue and red, respectively. The stored number of bunches is indicated in orange and the issuing of a checklist by a dashed vertical orange line. The date of the check-list is given in black next to the dashed vertical line.

Comparison of intensity ramp-ups 2015 - 2018

The time required for the intensity ramp-ups following the regular YETS has been reduced from year to year during Run 2. Figure 2 compares the number of operation days between the end of the beam commissioning and reaching a stored beam intensity of 2000 bunches. As 2015 was a so-called commissioning year, it covered two intensity-ramp ups, one with beams using 50 ns and the second one with beams using 25 ns bunch spacing. Figure 2 contains the latter, which lasted until the end of the operational year and reached a maximum number of 1825 bunches. It can be clearly seen that the required time to reach 2000 bunches could be significantly reduced over the years. The 14.5 days achieved in 2018 is close to the theoretical minimum, when taking into account the required 3 fills, 20h of stable beams per intensity step and adding the time for the turn-around between fills. Between 2016 and 2018 the total time

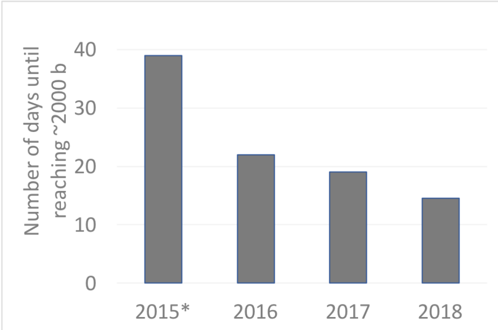

Figure 2: Comparison of the number of operation days between the end of the beam commissioning and reaching a stored beam intensity of 2000 bunches. Note, that for 2015 the bar shows the number of days required for the 25 ns intensity ramp-up until reaching 1825 bunches at the end of the 2015 run.

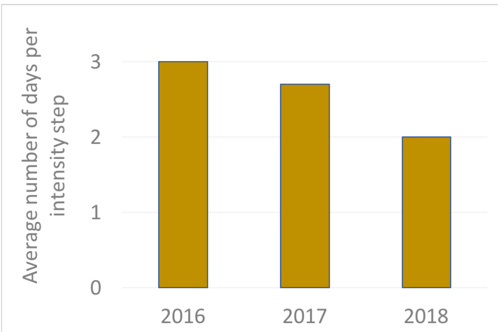

Figure 3: Comparison of the average number of days per intensity step. In 2016, 2017 and 2018 seven intensity steps were required to reach 2000 or more bunches following the end of the beam commissioning.

required for the seven intensity steps to reach 2000 or more bunches was reduced by 35%. This development reflects the increasing maturity of the accelerator, its systems and the hardware and OP teams as well as a smaller number of hard- and software changes in critical system during the YETS in the later part of Run 2. This development can also be seen in Fig. 3, which compares the average number of days required per intensity step in 2016, 2017 and 2018. The theoretical minimum per intensity step is 1.5 days in case of an accelerator availability of 100% and without any issues, requiring interventions.

Issues discovered during intensity ramp-ups

To emphasize the importance of intensity ramp-ups after a long shut-down or YETS, the list below gives an overview of the issues discovered during intensity ramp-ups in Run 2.

- Post Mortem / XPOC: missing data, data mis-aligned, missing files & synchronization, PM event builder stuck
- BIS timing mis-alignment
- Unbalanced rupture of the QPS internal quench loop
- Setup Beam Flag: glitches
- Beam Loss Monitors: communication issues BLM-SIS, un-physical readings in post mortem data
- Direct BLMs (IP6): connected to LBDS of wrong beam
- Collimators & Roman Pots: LVDT position drifts, LVDT/resolver faults
- Orbit Feedback: orbit jumps at optics changes, zeroing of reference, offsets due to BPM calibrations
- Dump Line: screen remaining in dump line, BTVDD images missing
- MKD and MKB erratics
- MKI – flashovers, MKI kicking last bunch of circulating beam
- Abort Gap Cleaning: in-sufficient cleaning, not functioning due to software issues / wrong parameters
- Radio Frequency: wrong low level settings
- QPS_OK flickering
- Beam current change monitor: false dumps
- BIS: too much attenuation of signal in fibre causing three protection dumps
- Earth fault in circuits (RB, RCS)
- Training quenches in RB circuits
- QPS: single event upsets causing protection dumps & communication issues
- UFOs: 16L2 events causing beam dumps & quenches
- Collimators: un-physical temperature readings wrong jaw positions
- TDI: vacuum issues and heating (2016)
- Insufficient cooling of a collimator
- Decrease of bunch length in the cycle below 1 ns
- Beam instabilities
- Injection: high losses and satellites leading to beam dumps

CONCLUSION - HAVE WE BEEN RUNNING SAFELY?

The machine protection systems worked very well in Run 2, avoiding damage to accelerator equipment, experiments and circuits. However, the LHC experienced a whole list of protection critical events like wrong parameters in protection systems, interlocks not acting as expected, operational mistakes, running with un-validated machine configurations, high level controls software commissioning with hundreds of circulating bunches, use of un-validated coupling knobs with strong impact on β^*, undetected quench heater firing, masking of critical interlocks during hardware commissioning and not correctly following of procedures. Besides these issues, no damage happened, which is due to the diverse redundancy in the machine protection systems and vigilant hardware experts, machine protection experts and OP teams. As shown in greater detail above intensity ramp-ups have been essential to identify and mitigate issues before physics production with stored beam energies of 300 MJ. Their length was reduced from 2016 to 2018 by 35% indicating reduced number of issues, due to increase maturity of the accelerator and less changes to critical systems during the YETS. New failures leading to very fast beam losses were experienced in Run 2 and confirmed the importance of the BLM system as last safety net. Therefore, this last safety-net should be complemented with an independent system - the beam current change monitor, which will issues a beam interlock in case of too high global beam losses.

ACKNOWLEDGEMENTS

The author would like to thank Z. Charifoulline, S. Redaelli, B. Salvachua, R. Schmidt, A. Siemko, J. Uythoven, J. Wenninger, C. Wiesner and M. Zerlauth, as well as the members of the MPP and rMPP for numerous discussions, inputs and contributions.