

Beam-related machine protection for the CERN Large Hadron Collider experiments

R. B. Appleby,* B. Goddard, A. Gomez-Alonso, V. Kain, T. Kramer, D. Macina, R. Schmidt, and J. Wenninger

CERN, Geneva, Switzerland

(Received 3 March 2010; published 10 June 2010)

The Large Hadron Collider at CERN, Geneva stores 360 MJ per beam of protons at the top machine energy. This amount of energy storage presents a considerable challenge to the machine protection systems designed to protect both the machine and the six LHC experiments. This paper provides an overview of the machine protection systems relevant to the protection of the experiments, and demonstrates their operation and level of protection through a series of injection and stored beam failure scenarios. We conclude that the systems provide sufficient coverage for the protection of the experiments as far as reasonably possible.

DOI: [10.1103/PhysRevSTAB.13.061002](https://doi.org/10.1103/PhysRevSTAB.13.061002)

PACS numbers: 29.27.-a, 29.20.-c

I. INTRODUCTION

The Large Hadron Collider (LHC) is a proton-proton collider at CERN, Geneva, and has a nominal beam momentum of 7 TeV/ c per beam. The design luminosity of the LHC (10^{34} cm $^{-2}$ s $^{-1}$ in 7 TeV p - p mode) requires 2808 bunches, of 1.15×10^{11} proton per bunch, in each of the two beams. This results in a beam energy of 362 MJ, which is enough to cause considerable damage to the normal and superconducting elements of the machine. Furthermore, the beam energy density is 3 orders of magnitude above currently operating machines, presenting a unique set of operational challenges to the machine and the associated experiments.

The LHC beam is prepared by a chain of injectors. A 26 GeV/ c beam is extracted from the proton synchrotron (PS), injected into the SPS and accelerated to the LHC injection momentum of 450 GeV/ c . When the LHC is filled, 12 batches per beam are transferred to the LHC ring through two transfer lines. The energy stored per batch is around 2 MJ, already exceeds that of an existing accelerator like the Tevatron, and is above the damage threshold of machine components. When the LHC is filled, the beams are accelerated to 7 TeV/ c per beam, with a total stored energy of 360 MJ. The beams now contain an unprecedented amount of energy, and 1 in 10^8 of the top energy beam is enough to quench a superconducting magnet. To deal with these levels of stored energy, a complex system of machine protection [1] has been designed to protect the machine, including beam loss monitors, current monitors, and quench monitors, which are connected to a beam dump and interlock system. The strategy of these systems is to manage the continuous losses from the beam, and handle failure scenarios over a large range of time scales. This protection system also needs to protect the delicate components of the six experiments of the LHC from accidental beam loss and failure scenarios, a task

complicated by the near-beam and movable silicon detectors of TOTEM [2] and Large Hadron Collider beauty experiment (LHCb)'s VETex LOcator (VELO) [3].

This paper gives an overview of the LHC machine protection system (MPS) relevant to providing protection to the experiments, with the intention of showing how the elements of the system work together. This will be done with a series of case studies and failure scenarios, which exercises the machine protection systems and demonstrates the degree of protection provided. The intent is to show the level of coverage and protection is sufficient to protect against a range of possible beam scenarios. With this goal in mind, the paper is organized as follows. The systems and strategy of the machine protection systems are discussed in Sec. II, including the injection and dump systems, and the failure scenarios and case studies are presented in Sec. III. A conclusion and summary are presented in Sec. IV.

II. OVERVIEW OF MACHINE PROTECTION SYSTEMS

The aim of the LHC machine protection systems is to protect the LHC and the experiments, which is achieved through many layers of machine protection systems and a strategy designed for the most likely, and as much as possible high risk, machine failures, and failure time scales. There is always the possibility of a multiple failure mode resulting in significant machine and detector damage, but the potential impact is reduced through MPS design and operational practice to reduce the impact and likelihood of these modes occurring. In this section the major subsystems and protection strategy are described. The philosophy of protection is derived from the analysis of the mechanisms for beam loss, during for the three broad phases of operation: injection, stored beam and dump.

1. Injection into the LHC from the SPS: the beam will be transferred from the SPS at 450 GeV/ c in several batches. One batch has up to 288 bunches, each with 10^{11} protons,

*robert.appleby@cern.ch

that are transferred from the SPS to the LHC within less than 10 μs . The energy stored in the batch is far above the damage limit, and beam losses during the transfer must be minimized. The beam must have the correct parameters (bunch structure, orbit, emittance) during extraction from SPS. Sources of beam loss are kicker magnet failures, wrong settings of magnets in the SPS extraction zone, the transfer lines or in the LHC, and aperture limitations (for example, closed vacuum valves).

2. Stored beam in LHC: At the end of injection, the energy stored in the beam is already about 26 MJ, 1 order of magnitude more than the beams at HERA or Tevatron at top energy. During the ramp the energy will increase to about 360 MJ. At a top energy of 7 TeV/ c the beams should collide for many hours. Any failure generating unacceptable beam loss must be detected and the beams must be extracted by the beam dump system into the dump blocks. Most failures during stored beam lead to fast movements of the orbit or beam size growth and to particle losses within some milliseconds to many seconds. Very rare failures leading to beam loss within one turn, such as the injection kicker accidentally deflecting the beam, must be avoided by the design of the hardware and the interlock system.

3. In the LHC, the full beam must never be lost around the accelerator, but always extracted into the beam dump blocks. We distinguish between programmed beam dumps, e.g., at the end of a physics fill, and emergency beam dumps after the detection of a failure. The extraction kickers must deflect the beam with the correct angle at the correct time synchronized with the revolution clock. During the rise of the kicker magnetic field there should be no beam in the 3 μs long particle-free beam abort gap, to prevent particles being deflected by a wrong angle. The beam must be transferred with minimum losses via a 800 m long transfer line to the dump blocks. Sources of beam loss during extraction are kicker and septum magnet failures as well as the presence of particles in the beam abort gap.

In the following sections, the mechanism of beam loss in the operating scenarios is reviewed and the components of the machine protection system described. The level of detail is sufficient for an overview of the philosophy and for the beam failure scenarios presented later, which are divided into the injection and stored beam phases presented in this section.

A. Regular and accidental particle losses

A particle moving in a circular accelerator has long-term stable motion only if its betatron amplitude is within a certain distance of the core of the beam. The amplitude limit for this stability, the dynamic aperture, is an important parameter of the optics of the machine. Particles with amplitudes larger than the dynamic aperture are eventually lost. Different phenomena make the particles drift away from the beam core; these include space charge effects,

scattering of particles with other circulating particles or with residual gas molecules, and beam-beam effects at the collision points. As a consequence, a small fraction of the beam is constantly being lost. These losses are unavoidable and define the need for collimators in the LHC.

Equipment failures, operation mistakes, or phenomena such as magnet quenches may generate accidental beam losses. This kind of loss is generally not a continuous process and occurs due to a deviation from nominal beam operation. Accidental beam losses in LHC can be due to beam deflection or defocusing when there is a change in the magnetic field. For example, after a failure of a magnet power supply or a magnet quench the beam can be deflected or defocused, leading to an increase of the losses. A further cause of accidental beam loss is aperture reduction. A total of 476 different objects can reduce or completely close the beam aperture at LHC, including collimators, vacuum valves, roman pots, injection and matching screens, rf and safety stoppers and alignment mirrors. Some movable objects such as wire scanners can also produce scattering of the circulating beam and generate localized losses. Finally, in case of a vacuum leak, the pressure in the vacuum chamber will locally rise. This will lead to abnormal losses due to scattering of the protons with the gas molecules.

Accidental beam losses arising from processes like these can occur in a single turn, over multiple turns in a short time (less than, say, 1 s), or during a longer time scale. The time constant of the losses is an important parameter as it determines the strategy for protection. The time scale for losses can thus be classified: (i) Ultrafast losses—significant beam losses develop in one turn or less. Ultrafast losses can be produced by failures at beam injection, at beam extraction, or during operation of dedicated kicker magnets. (ii) Very fast losses—losses develop in less than 5 ms (about 60 turns). Failures in some particular magnets in LHC can produce losses in this time scale during nominal operation. (iii) Fast losses—losses appear in less than 1 s (11236 turns). Most operational failures are likely to produce losses within this time scale. (iv) Slow losses—losses develop in times larger than 1 s. Failure of low strength multipole compensator magnets, vacuum leaks, and unexpected movements of slow-moving objects into the beam generate slow losses.

The protection strategy has to cover all the expected failure cases, which may produce losses at any speed.

B. Strategy for protection and beam interlocks

Protection of the LHC implies active protection (detection of the failure and extraction of the beam before damage thresholds are reached) as well as passive protection (reliability of equipment and collimators). The strategy is defined by the following principles: (i) The aperture of the machine shall be defined by the collimator jaws, with beam loss monitors (BLMs) close to the collimators to

monitor loss. (ii) Failures within the equipment that act on the beams should be detected early to generate a beam dump request before the beam is affected. (iii) The beam should be actively monitored with fast and reliable beam instrumentation, to detect abnormal beam conditions and generate a beam dump request within a very short time, down to a single machine turn. The BLMs play a predominant role in this protection. (iv) An active interlock loop signal is required for operation, and the absence of the signal is considered as a beam dump request or injection inhibit. (v) The extraction and transfer of beam from the SPS to LHC is only permitted when all parameters of the LHC and the transfer lines are correct. (vi) The beam dumping system shall, on receipt of a dump request or internal fault detection, operate reliably to safely extract the beams into the external dump blocks. (vii) Protection from specific failure cases should be provided by the passive protection of beam absorbers and collimators, in particular, during injection and extraction. (viii) The protection system should have redundancy, such that failures may be detected by more than just one system.

Not all of these strategies are efficient to protect against all possible failures. For failures producing ultrafast losses, there is no time to detect the failure and extract the beam. Effective protection is based on dedicated collimators and absorbers, and on the high reliability of the equipment that could lead to these losses. Very fast losses can be detected by the BLMs in the aperture limitations, but the time constant of the losses may be too short to allow the protection systems to react before quench or even damage levels are reached. Hardware monitoring equipment has been installed for the magnets that can lead to very fast losses in order to detect critical failures before the beam is affected. For failures leading to fast and slow losses, several protection mechanisms react in parallel in order to provide optimum protection.

A key part of the protection strategy is the beam interlock system (BIS). The BIS centrally manages the interlocks of the LHC, taking input from many passive and active systems and regulates the presence of beam in the LHC machine. The LHC BIS [4] is based on beam interlock controllers (BICs) distributed around the machine (17 in total) and linked by optical fibers to gather user permits (UPs) from different monitoring systems. Based on the status of the surveyed systems and the presence of the user beam permits, the beam interlock system gives the “beam permit” for each beam. Without the beam permit from the LHC ring, injection is inhibited. The beam dump is triggered to remove any circulating beam in a safe way. Therefore any problems in the machine or unsafe conditions means the injection or beam presence is not permitted. For example, to give the beam permit all vacuum valves must be open, all magnets powered, etc. In the injection regions IR 2 and IR 8, injection BICs are installed which are not part of the LHC ring interlocking loop. The

interlock system requires several subsystems for proper operation and for protection from failure scenarios on many different time scales. The architecture of the BIS is two BICs in each interaction region, which are connected to two beam permit loops around the machine. These loops, one for each beam, contain a 10 MHz signal and the presence of this signal is equivalent to a beam permit. A removal of a user permit in any BIC causes the BIS to break the beam permit loop and trigger a beam dump request. The BICs have a further input derived from the LHC beam intensity and energy, called the setup-beam flag. If this flag is TRUE, certain BIC inputs are maskable and can be ignored. This is not possible if the setup-beam flag is FALSE.

The LHC injection interlock system and the SPS extraction interlock system permit extraction of beam from the SPS, transfer, and injection into LHC. The systems monitor all relevant parameters in the SPS (e.g. beam position at the extraction point, magnet currents, beam energy), of the magnets and collimators in the transfer line, and of the injection setting in LHC (septum magnet current, injection collimators). Injection of high intensity beam is only permitted if there is already some beam circulating in LHC.

The powering interlock system (PIS) for superconducting magnets provides protection against uncontrolled release of the energy stored in the superconducting magnets and circuits. The system receives input from different hardware diagnostics systems (quench detection, cryogenics, powering) and allows powering only if the conditions are safe. When a failure in the powering system is detected, a trigger is transmitted to the BIS. The PIS for normal conducting magnets provides protection against overheating of the normal conducting magnets. The system receives inputs from temperature sensors and from the power converters for normal conducting circuits.

When a failure in the electrical circuit in any magnet is detected (e.g. a quench in a SC magnet, or a power converter failure), the power converters are stopped, the magnet energy is extracted into resistors, and a trigger is transmitted to the BIS. The PIS cannot guarantee that the failure detection is fast enough to ensure a beam dump before significant losses develop, and so fast-response active protection systems are used (see the discussion on fast magnet current monitors in the next section).

Figure 1 illustrates the complexity of the LHC beam interlock system, together with the inputs and the links to other systems. The inputs from the powering system (powering interlocks for superconducting and normal conducting magnets, magnet current monitor) and the beam loss monitor have already been mentioned. All elements that can move into the beam are interlocked (vacuum valves, screens, and mirrors), as are access systems to the machine. Since the collimation system must define the aperture for the beam, the position of the collimators is critical. In the case of a failure, e.g., the opening between

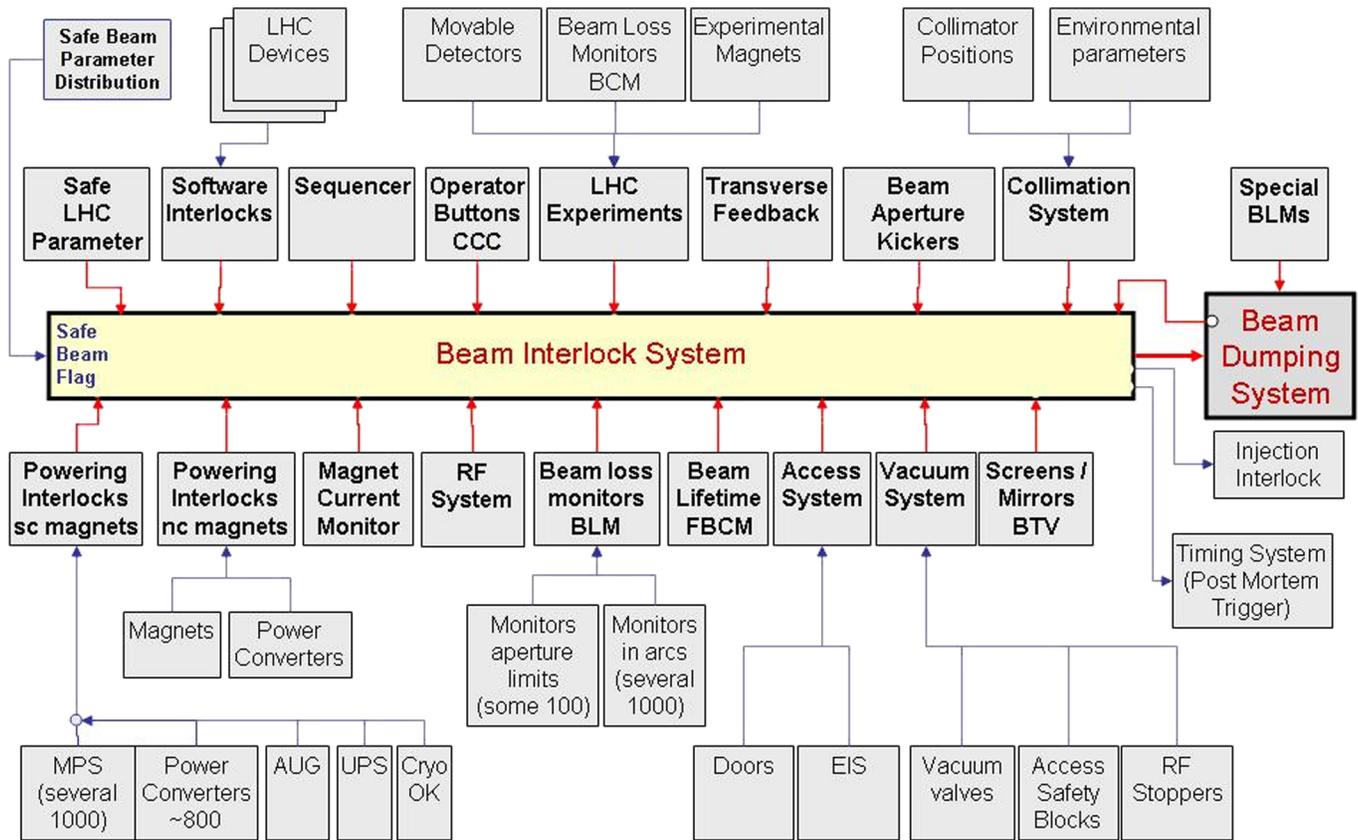


FIG. 1. (Color) Beam interlock system inputs in the LHC.

two collimator jaws are not as requested, or a jaw is not at the predefined position, the beams will be dumped.

Additionally, all LHC experiments provide input. All experiments use fast beam loss monitors which have an input to the BIS and if the experimental magnets trip a beam dump request will be generated. Furthermore, some experiments use movable detectors that should be far away from the beam during injection and during the energy ramp. When the beams are colliding and the beam is declared stable, the detectors are permitted to move closer to the beam. If this condition is violated, a beam dump request will be generated. Examples of movable detectors in the LHC are the roman pots of TOTEM [2] and the LHCb's VELO [3]. The former will be considered in Sec. III B for the circulating beam scenarios.

Another input to the BIS comes from the software interlock system (SIS), which is designed to monitor slow changing parameters and less critical systems. For example, the currents in all magnets are interlocked in the SIS for injection, and changes to orbit correctors in stable beam operation are SIS interlocked to prevent the buildup of local bumps. The reaction time of the SIS is longer than the BIS systems, due to the use of network communications rather than hard-wired connections.

There are a few other inputs, for example, allowing the operator to request a beam dump.

C. Active protection

Protection against failures developing on the time scale of ten or more turns (multiturn failures) is performed by active surveillance of equipment and by beam parameter monitoring. The active protection systems provide input to the BIS.

A large system of over 3600 beam loss monitors (BLM) distributed along the LHC circumference are used to detect any abnormal beam loss [5]. The BLMs consist mostly of ionization chambers (IC), and secondary emission monitors are used in regions where very large signals are expected, for example, near collimators and absorbers. Six ICs are placed around each superconducting quadrupole, three on either side, with their positions optimized by impact simulations. Other BLMs are placed near each collimator, absorber, special devices, etc. The system is designed with a sampling period and reaction time of $40 \mu\text{s}$, and so accumulates charge over this interval. Internally the BLM system uses twelve running averages with time windows ranging from $40 \mu\text{s}$ to over 100 s to be able to monitor losses on all time scales. For BLMs installed around superconducting magnets each of those 12 running averages has a dedicated dump threshold. The thresholds have a built-in energy dependence since, for example, the quench threshold for the main magnets is reduced by a factor of approximately 30 between injection

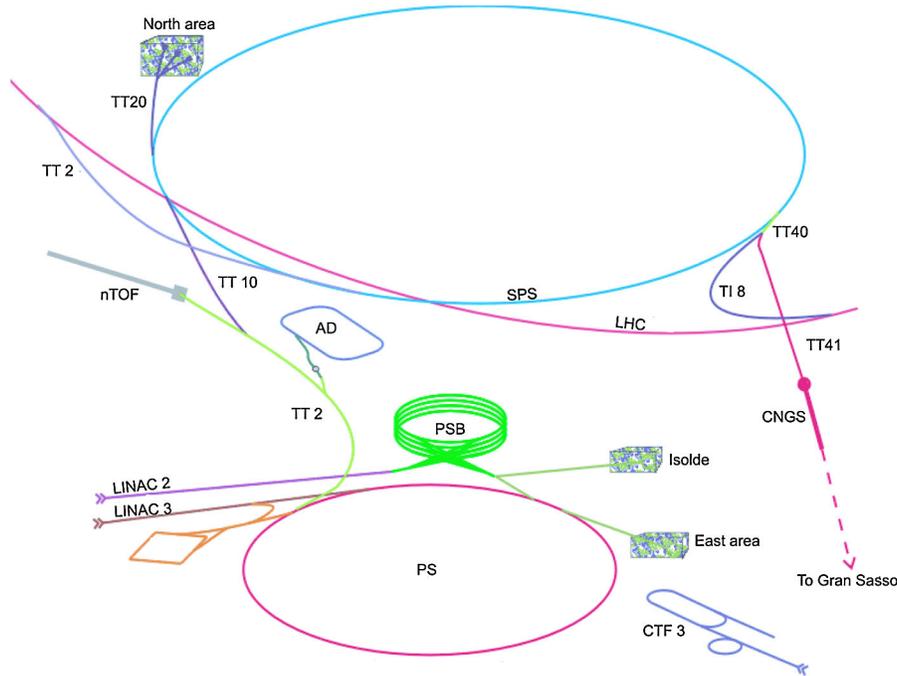


FIG. 2. (Color) Schematic view of the LHC preinjector chain.

energy and 7 TeV. The default thresholds of BLMs around superconducting magnets are set to 30% of the estimated quench threshold, but it is possible to scale the thresholds up or down. For example, the applied thresholds on the IC BLMs on an LHC long straight section (LSS) quadrupole are 75 and 24 μGy for the two BLM locations, which is equivalent to 4.1 and 1.3 nC of charge in the 40 μs integration time.

Since the BLMs cover all sensible locations around the LHC ring, they are a central component of the LHC MPS. Their fast reaction time and good machine coverage makes them sensitive to almost all failures, but the setting of the dump thresholds may require balancing the rate of unnecessary dumps with respect to the number of beam induced quenches.

Powering interlock systems are installed in the LHC both for the protection of superconducting and normal conducting magnet circuits and their associated power converters are also connected to the beam interlock system. The reaction time of the powering interlock systems ranges from approximately 10 ms to a few microseconds for the large and most critical circuits. For a large class of circuits and power converters, in particular for superconducting circuits with large inductance and small resistance, the powering interlock system is able to generate an interlock to dump the beams well before the decaying magnetic fields lead to beam loss on collimators or other LHC components.

Powering failures of magnets are among the fastest multiturn failures when normal conducting magnet circuits are involved, as the circuit time constants τ are generally

fast due to large resistances R ($\tau = L/R$) as compared to superconducting magnet circuits. The very first simulations of failures in the LHC [6], and the more recent studies [7,8], have indicated the powering failures of the separation/recombination dipoles (D1) in the insertions 1 (ATLAS) and 5 (CMS) would lead to beam impact on the collimators within a few tens of turns. Such time scales are too small to be effectively protected by the powering interlock systems. Furthermore, the rate increase depends strongly on the shape of the beam halo that is difficult to predict and that can change quickly with beam conditions. To provide redundancy for the BLM system for the fastest powering failures, a device able to detect powering failures on time scales of milliseconds and for relative field changes of around 10^{-4} was required. Such a device was developed for the HERA accelerator at the same time when the need of a better protection for the LHC was recognized. The HERA system was successfully adapted for LHC MPS as a fast magnet current change monitor (FMCM) with 12 devices installed in the LHC ring and 21 in the injection lines, where those devices are operational since 2006. FMCMs have been installed on the most critical normal conducting dipole and quadrupole magnet circuits installed in the LHC ring [7,8].

D. Passive protection

The passive protection systems of the LHC define the aperture of the machine, with the collimation system providing the smallest aperture restriction. The beam loss monitors are placed close to all aperture restrictions. The

LHC will be nominally operated with a collimation system to avoid quenching of superconducting elements and provide protection from beam damage. The inefficiencies when scattering a high-energy proton in a 1 m carbon collimator mean the LHC collimation system is designed to be multiphase and multiturn, with staged collimators at increasingly further distances from the beam, collimating protons over successive turns. The system is divided into momentum cleaning collimation, with the beam cleaned in momentum space in a dispersive region, and betatron cleaning designed to clean the horizontal and vertical phase space of the beam. The primary beam halo is intercepted and rescattered by betatron primary collimators made of carbon, which have jaws 6σ from the beam at top energy. The scattered secondary halo is intercepted by a betatron secondary collimators, again made of carbon with a jaw 7σ from the beam at top energy. The cleaning inefficiency of the secondary collimators produces a tertiary halo, which is cleaned by tertiary halo collimators located on the incoming beam on either side of each experiment. The tertiary collimators provide protection to the superconducting final triplet magnets from the tertiary halo, and also provide protection from abnormal dump events. The collimators are supplemented by several masks and spoilers.

The collimation system is designed to always provide the smallest aperture restriction, and so beam losses should be concentrated on the collimators. This provides a high level of passive protection to the machine and the experiments provided this requirement is maintained. There is no guarantee of complete phase space coverage over one turn, so a loss over one turn may not strike a primary or secondary collimator and may strike a sensitive region of the machine or a near-beam detector.

E. Injection protection systems

In this section the injection systems of the LHC are described. Almost all accelerators of the CERN accelerator complex participate in the production of the LHC beams and the last stage of preacceleration takes place in the SPS synchrotron. 12 injections—or 12 transfers from SPS to LHC—per ring are required to fill the LHC, with beam 1 transported along transfer line TI 2 and beam 2 transported along transfer line TI 8. The LHC preinjector chain is shown in Fig. 2.

Careful control of the trajectory and the preservation of the very small emittance during transfer and injection are of key importance, due to the limited mechanical aperture of the transfer line magnets, the high intensity and energy in the beam, and the tight tolerances on the beam parameters at injection into the LHC [9]. The machine protection elements in the lines should prevent the transfer of damaging beams into the LHC, either by interlocking or by the passive protection systems.

The 450 GeV beams coming from the SPS are injected into the LHC in the combined experimental and injection

insertions IR 2 and IR 8 [5,10]. In both insertions the beams are injected from the outside of the accelerator ring and from below, deflected in the horizontal plane by a series of septum magnets (denoted MSI).

The LHC will be filled with 12 SPS batches per ring and one gap free of bunches is $3\ \mu\text{s}$ long to allow for the rise time of the dump kicker and loss-free dumping. The available aperture of the LHC transfer lines and the LHC at injection energy is small, making injection a particular concern for machine protection. Active injection protection includes monitoring of critical settings and measurements through means of a hard-wired, very fast interlocking system based on the concept for the LHC ring interlocking system [4] in Sec. II B. Input signals come from many sources including the power converter surveillance (all transfer line power converters are checked 3–4 ms before extraction from the SPS), the kickers, the settings surveillance of the passive protection devices (jaw position), and the position of the beam stoppers in the transfer lines. In the case of a fast failure close to extraction there might not be enough time to detect the failure and react in time before the beam passage. Additional passive protection against mis-steered beam is foreseen with collimators and absorbers, as discussed in Sec. II D, with input into the BIS.

The TDI, TCDD, and TCLI are dedicated collimators against kicker failures on injection. Both the extraction kicker and the injection kicker with their very short rise times and a number of possible failure scenarios are followed by dedicated protection against kicker failures 90° betatron phase advance downstream. The vertical injection stopper TDI and the auxiliary collimators TCLI protect the LHC aperture against failures of the injection kicker MKI. The TDI is a movable two-sided 4.25 m long vertical injection stopper and is placed $\mu_y = 90^\circ$ (~ 70 m downstream of the kicker) from the MKI. It is 15 m upstream of the superconducting recombination-separation dipole D1. A 1 m long mask made of Cu, the TCDD, protects the superconducting coils of the D1 separation dipole from the showers generated in the TDI during impact, see Fig. 3. In case the phase advance between the MKI and the TDI is not exactly 90° , the protection against kicker failures is complemented by two double-jaw vertical auxiliary collimators, TCLIA and TCLIB [11]. The required protection setting of the system TDI-TCLI is 6.8σ for an LHC injection aperture of 7.5σ [6].

Injection protection starts at the SPS extraction and cannot be disentangled from the main LHC ring protection. Adequate injection protection implies signal exchange between the LHC beam interlocking system (see Sec. II B) and the SPS beam interlocking system. The LHC “beam permit” is input to the “injection permit,” and the injection kickers only pulse if the injection permit is TRUE.

For the LHC injection some additional conditions are evaluated by means of the status of a handful of flags in the

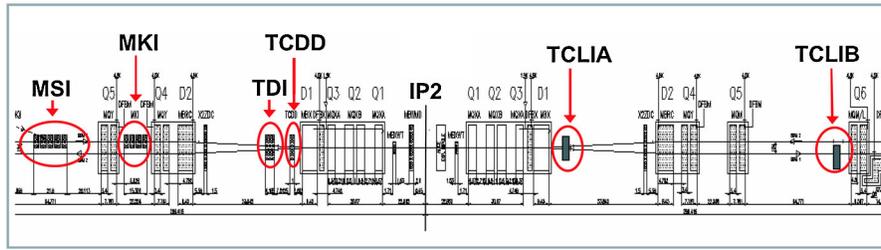


FIG. 3. (Color) Overview of the injection region in IR 2 (distance MSI-TCLIB: ~420 m).

master BICs. Only “probe beam intensity” can be injected into an empty LHC to avoid damage to sensitive equipment like inner detectors of experiments. The probe beam intensity is settable through the system of management of critical settings, and the default is 10^{10} protons. This value is set from damage simulations and is an intensity visible to the beam instrumentation. The probe beam flag is generated from a beam current transformer in the SPS and the beam presence flag by a fast beam current transformer in the LHC signalling whether the LHC is empty or whether there is beam circulating. For more details see [12]. Intensity above the setup beam intensity can only be injected into the LHC if the LHC setup beam flag is FALSE, ensuring there is already some beam in the machine. The setup beam intensity at 450 GeV is 1×10^{12} protons.

More information on the interlocking logic between the SPS extractions and the LHC can be found under [13].

F. Beam dump systems

The function of the beam dumping system (LBDS) is to fast extract the beam in a loss-free way from each ring of the LHC and to transport it to an external absorber, positioned sufficiently far away to allow for appropriate beam dilution in order not to overheat the absorber material. To minimize the losses in the extraction system will require the presence of the abort gap, during which the field of the extraction kicker magnets can rise to the nominal value. Given the destructive power of the LHC beam, the dumping system must meet extremely high reliability criteria, which condition the overall and detailed design.

IR 6 of the LHC is dedicated to beam dumping, and the layout is shown in Fig. 4. The main components of the LBDS comprise (per beam) 15 extraction kicker magnets (MKDs), 15 Lambertson dump septum magnets (MSDs), four horizontal and six vertical dilution kickers, and the external beam dump block. There is a vacuum system for the two beam lines, with an entrance window before the

dump block. The fixed TCDS diluter, mobile TCDQ, and standard LHC collimators protect the septum magnets and LHC machine against unsynchronized dumps. The lattice quadrupole Q4 also contributes to the deflection required to extract the beam.

Each LHC beam is initially deflected by the MKD kicker magnets which provide a total horizontal deflection of 0.285 mrad or 5.25 mm in the center of quadrupole Q4 located right behind the MKD. The dipole component at this distance from the neutral axis of Q4 further increases the horizontal deflection angle by 0.09 mrad, to a total angle of 0.37 mrad.

The protection elements TCDS in front of the MSD and the TCDQ installed in front of Q4 downstream of IP6 serve to protect machine elements from a beam abort that is not synchronized with the particle-free beam gap, or from particles in the abort gap itself.

The MKD and dilution kickers are triggered from the BIS via the triggering and synchronization unit (TSU).

The LBDS interfaces directly with several different LHC systems, as indicated schematically in Fig. 5. The LBDS is connected to the beam interlock system, to set the beam permit TRUE/FALSE depending on the state of the LBDS and to trigger the beam dump action if the beam permit is read as FALSE. This occurs via a dedicated module which monitors the 10 MHz beam permit signal on the optical fiber loop. There is also a connection to the safe LHC parameters distribution system and to the BICs of the injection interlocking system, to be able to inhibit the LHC injection for arming the beam dump, via dedicated optical fiber links from IR 6 to IR 2 and IR 8.

The function of the LBDS is to cleanly dump the LHC beam. There is a delay between the request to dump the beam by the BIS and the end of the beam dump. The BIS needs to inform the beam dump of the request by interrupting the beam permit loop, which takes up to $50 \mu s$ if the request needs to travel halfway around the ring. It then

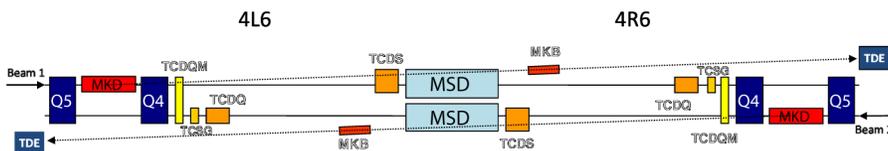


FIG. 4. (Color) An overview of the LHC beam dump region.

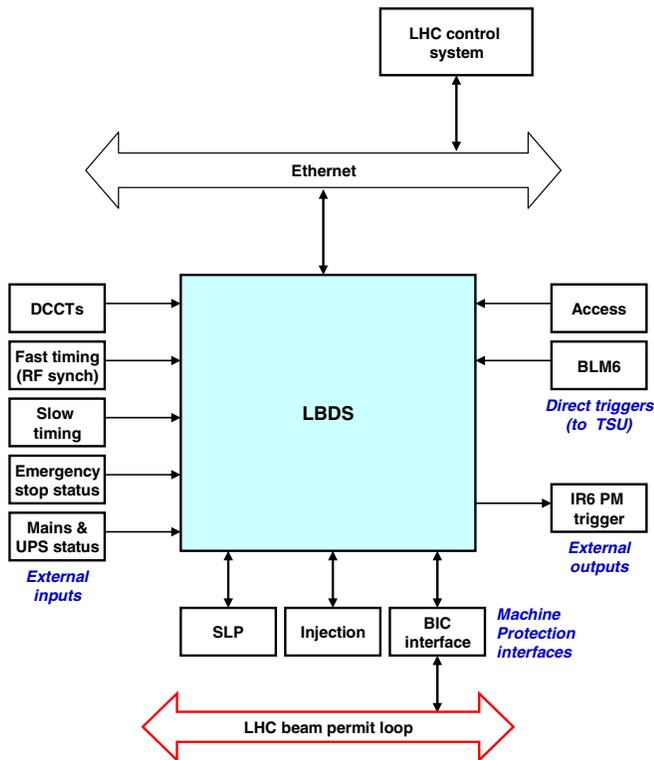


FIG. 5. (Color) LHC beam dump system interfaces with other LHC systems.

takes a maximum of one turn to synchronize the extraction kicker with the abort gap, and a further turn to fully dump the beam. Therefore the maximum beam dump response time is approximately $270 \mu\text{s}$, or three turns.

III. FAILURE SCENARIOS

The LHC machine protection systems and strategy described in this paper protect the LHC machine and experiments as far as reasonably possible. In this section, the machine protection systems are exercised to demonstrate that the systems described provide functional protection and show how they respond to real beam scenarios. The scenarios presented here focus on the machine protection relevant to the experiments of the LHC, and cover both injection and stored beam operational modes of the machine. The scenarios are particularly sensitive to the machine protection temporal flow, and how the time scales of the parts of the systems work together.

A. Injection

In this section the machine protection systems for LHC injection are discussed and the level of protection analyzed. A series of scenarios relevant for the experiments are then discussed showing how the experiments are protected against injection turn magnet mis-settings. The injection scenarios are illustrated using LHCb, but the

analysis of protection and conclusions apply to all the LHC experiments.

The risks to the experiments when injecting beam arise from the incorrect fields of magnet in the arcs or long straight sections. These ultrafast failure scenarios, at a time scale of less than 1 turn, can arise from incorrectly set magnets on injection, due to a communication error or from faulty hardware. The subsequent first-turn distorted trajectory may strike part of the machine or detector regions, causing partial loss or a direct strike of an injected bunch. The potential impact of this class of accident is reduced by the machine protection logic described in the previous section, through the use of a low population of injected protons—the probe intensity—where the injected number of protons is limited by the management of critical settings system. This limit is set to minimize the damage caused by beam strike, while still allowing operation of the machine elements. Further protection from injected beam loss is provided by the beam interlock system, which only permits injection when the LHC ring is ready for beam, and the magnet current interlocks, which are done in software and controlled by the software interlock system (SIS).

The injection turn accident scenarios depend on the layout of the experimental region, in terms of vacuum chambers, locations of subdetectors, and interaction region accelerator elements. The details of the interaction region vary from experiment to experiment, and the key features can be seen from a detailed discussion of LHCb. The relevant differences at other IRs will be discussed later in the section. This interaction region contains an 18 m conical vacuum chamber [14], consisting of three cone-shaped segments running through the detector. This vacuum chamber provides the principle aperture restriction in this region and the VELO sets the aperture restriction at the interaction point [3]. The distance of approach to the beam at injection is 30 mm (the VELO is inserted to 5 mm for collisions).

The next significant aperture restriction begins 2.25 m from the IP, where the vacuum chamber becomes 50 mm in diameter for 0.25 m, and then approximately 20 m from the IP, where the apertures of the machine elements MBXWS (dipole) and Q1 (final triplet quadrupole) begin. The resulting aperture model is shown in Fig. 6, where the solid line shows the vacuum chamber and the stars show the aperture restrictions from magnetic elements. Also shown are the magnets in the interaction region relevant to this study, for example, the final triplet quadrupoles Q1 around IP8, MQXA.1L8 and MQXA.1R8.

The potential sources of first-turn orbit distortion are also different in each IR, and again LHCb is taken as an example. The LHCb spectrometer dipole (MBLW) is located about 5 m to the right of IP 8, and is designed to give a deflection of $181 \mu\text{rad}$ at the top energy of 7 TeV. The field is orientated in the vertical plane, and hence the magnet gives a horizontal deflection. The integrated field

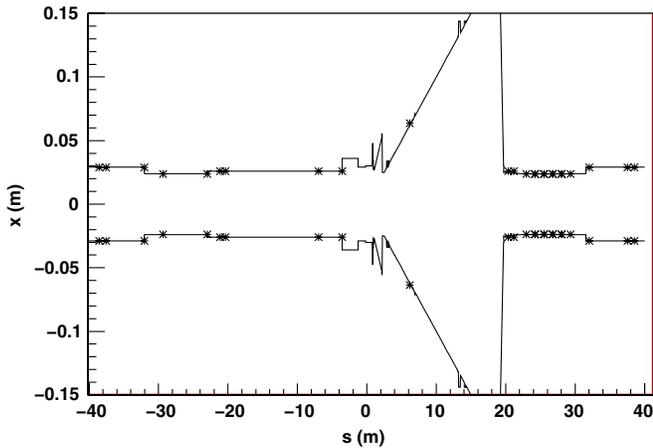


FIG. 6. (Color) The aperture model of the LHCb interaction region. The Q1 magnet on the right-hand side is located at the dense row of stars.

is 4.2 T m, which would give a deflection of 2.82 mrad at injection. The spectrometer is also required to work with the opposite polarity for reduction of LHCb analysis of systematic errors. The strong effect on (both) beams is compensated by three additional horizontal corrector magnets which, when acting with the experimental dipole magnet, give a closed asymmetrical bump across the IP. This bump is independent of the optics as there are no magnetic elements between the magnets. The bump magnets are MBXWH, located at -5 m, (with a bend angle of $181 \mu\text{rad}$, positive and opposite MBLW), and two much weaker magnets at ± 20 m, called MBXWS with a bend strength of $\pm 46 \mu\text{rad}$. The bump gives a residual IP crossing angle of $135 \mu\text{rad}$. Note there is an additional crossing angle bump imposed on LHCb in the horizontal plane (the outer bump), to reduce parasitic bunch collisions [15], and a vertical plane parallel separation bump for injection.

The interaction region of ALICE, IR 2, is very similar to LHCb, although ALICE has no near-beam VELO to provide an aperture restriction. The IRs of ATLAS and CMS differ from LHCb and ALICE, apart from vacuum chamber differences, in two main ways. The first is the presence of the forward region fixed collimator called the TAS. The purpose of the TAS is to collimate the collision debris from the high luminosity interaction point which would otherwise quench the coils of the superconducting final triplet quadrupoles. The TAS has a half aperture of 17 mm and starts 21 m from the IP. The second significant difference is that ATLAS and CMS have no strong compensator dipole in the interaction region, and hence no inner experimental bump. This reduces the number of strong IR magnets to consider in the failure studies [16].

The injection errors arise from the incorrect fields seen by the beam on injection. For example, consider a mis-setting, undetected failure or communication error to the LHCb strongest compensator magnet, MBXWH. This magnet, nominally set to $181 \mu\text{rad}$ in the horizontal plane

at injection, and could be set to any strength up the maximum of 4.2 T m, and/or with reversed polarity. Hence, the orbit error about the closed orbit will be in the horizontal plane. Further errors can occur when the corrector coils attached to the low- β quadrupole Q1, MCBXH and MCBXV, are incorrectly set on injection. These orbit correctors play a role in setting the beam crossing angle and parallel separation on injection, with MCBXH involved in creating the beam crossing angle at the IP, and MCBXV part of creating the vertical parallel separation of 2 mm between the beam at the IP for injection. Furthermore, the incorrect settings of the D1 and D2 horizontal separation dipole magnets can also cause beam accident scenarios to hit elements of the interaction region. These magnets are used to separate and recombine the beams, and cause the transition from separate beam pipes to a shared beam pipe. They are both 9.45 m long and are superconducting, with a single set of coils in the cryostat (in contrast to the magnets in points 1 and 5, where D1 is normal conducting). All the separation dipoles are strong magnets with the potential to cause significant orbit distortion on the first turn.

The first-turn beam strikes for the experiments have been studied for LHCb in [17,18], for ALICE in [17,18], and for ATLAS in [16] (where the results are applicable to CMS), and will now be illustrated using the LHCb interaction region already introduced in this section. As an illustration, the scenario for excess current in the compensator magnet MBXWH with beam 1 is shown in Fig. 7, where the magnet is set from the nominal injection strength (6.4% of maximum) to maximum strength. This corresponds to a change of angle from $+181 \mu\text{rad}$ to $+2.82$ mrad. The range of magnet settings shown by the cone show those which could be dangerous to the interaction region beam pipe or elements. The calculations show

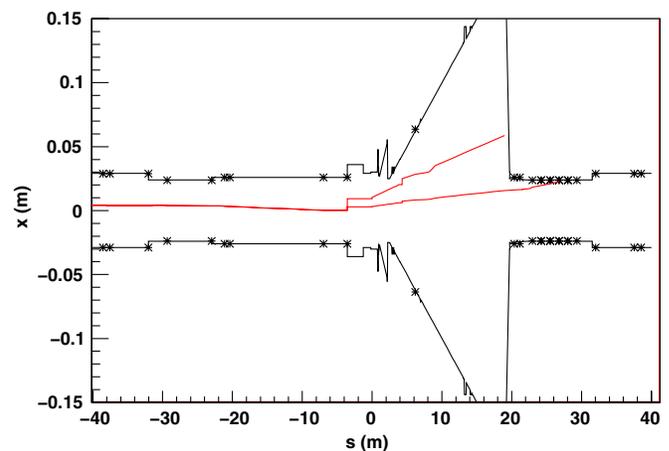


FIG. 7. (Color) The range of compensation dipole corrector (MBXWH) settings which are dangerous for the LHCb beam pipe and interaction region magnets for beam 1 and excess current in the magnet. On this plot $s = 0$ corresponds to the interaction point and the center of the experimental cavern.

this dangerous region is defined by MBXWH being set to an angle of $981 \mu\text{rad}$ to the maximum angle, which is equivalent to 34.8% to 100% of maximum strength (6.4% is nominal at injection). The figure shows the beam can hit the LHCb conical beam pipe, the spectrometer corrector MBXWS, or, for a few settings of the magnet, the beam pipe of Q1 (the beam trajectory can hit the element MQXA). The situation for LHCb can be contrasted to a similar study performed for ATLAS [16], where a similar range of beam accidents was considered for the main orbit correction magnet MCBX (attached to Q1). It was found the mis-setting of the magnet resulted in pilot beam loss in the ATLAS beam pipe or the TAS collimator. There is no TAS collimator in LHCb, and hence there is possible beam loss in MBXWS or MQXA.

An analogous calculation for ALICE is shown in Fig. 8 for the case of a first-turn field error in the separation dipole D1, which can cause injection trajectory changes due to the high field strength. The impact on the beam strike scenarios with injection jitter (1.5σ in both planes) can be seen in Fig. 9, where a Monte Carlo calculation of the beam size and envelope shows potential beam strike on LHCb for errors in the final triplet corrector MCBXH.

The protection from the first-turn beam strikes on the experiments shown in these examples is provided by several different parts of the machine protection systems: (i) The BIS will only permit injection if there is a full set of user beam permits. (ii) Injection into an empty machine is only permitted with a reduced proton bunch (probe beam intensity). Intensity above the setup beam intensity can only be injected into the LHC if the LHC setup beam flag is FALSE. Therefore the potential damage impact of injection failure scenarios is significantly reduced. The probe beam is replaced with a higher intensity beam through overinjection. (iii) The SIS has a set of software

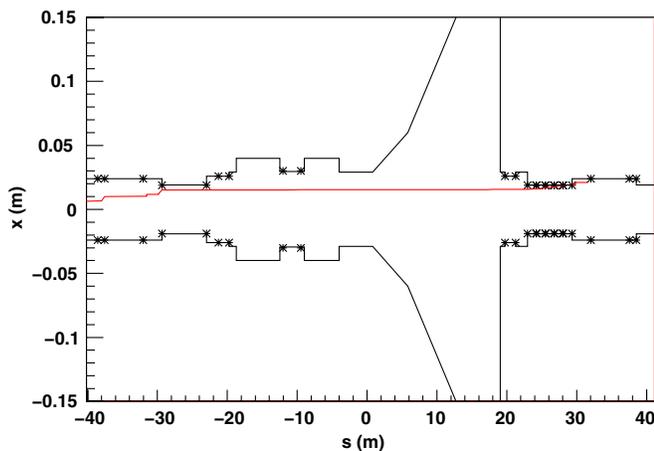


FIG. 8. (Color) A possible MBX.4L2 (D1) dipole setting which is dangerous for the ALICE beam pipe and interaction region magnets, where $s = 0$ corresponds to the interaction point and the center of the experimental cavern.

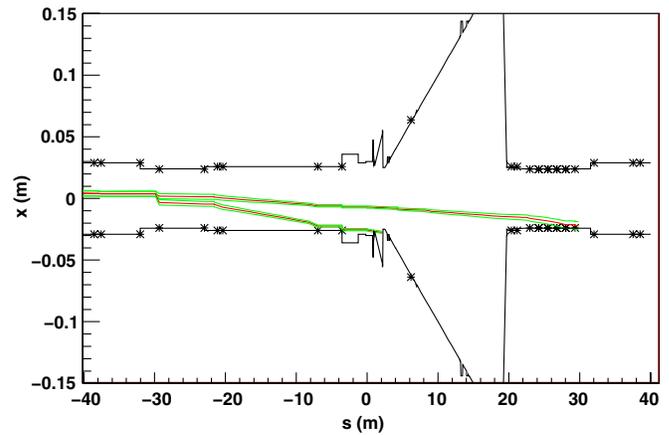


FIG. 9. (Color) The range of the corrector on Q1, MCBXH settings, which are dangerous for the LHCb beam pipe and interaction region magnets for beam 1 and excess current in the magnet. The effects of injection jitter and 3σ beam envelope scraping have been included. In this plot, the beam centroid is shown in red and the beam envelope is shown in green.

interlocks on the magnet currents, providing a significant level of protection.

The magnet current interlocks will provide protection from magnet mis-settings provided they are set sufficiently tightly to prevent the beam strike scenarios illustrated by Figs. 7 and 8. The resulting magnet current thresholds for beam 1 in LHCb to avoid beam orbits striking the vacuum chamber are shown in Table I, as a fraction of the maximum field and expressed as integer percentiles. In this Table, the trends in the current thresholds are consistent with the optics and apertures and these current thresholds should be considered as maximum permissible currents to avoid injection turn beam accidents. Therefore these thresholds should be considered as part of the current software interlocks to avoid beam strikes on the aperture restrictions of MQXA.1R8, MQXA.1L8 or the conical vacuum chamber of LHCb.

The current thresholds to avoid beam loss calculated for the various accident scenarios can be compared to the magnet current interlocks set in the SIS for injection. At the nominal settings, the orbit correctors are interlocked to a tolerance of $100 \mu\text{rad}$ around the nominal current, until

TABLE I. The required thresholds of the magnets to avoid beam accident scenarios on injection, rounded to a integer percentile, for beam 1.

Magnet	Threshold (nominal polarity)	Threshold (reversed polarity)
MBXWH	35% ($987 \mu\text{rad}$)	-14% ($-395 \mu\text{rad}$)
MCBXH	35% ($-354 \mu\text{rad}$)	-55% ($556 \mu\text{rad}$)
MCBXV	30% ($-313 \mu\text{rad}$)	-28% ($292 \mu\text{rad}$)
MBX.4L8	8.5% ($-2026 \mu\text{rad}$)	4.9% ($-1168 \mu\text{rad}$)
MBRC.4L8	8.8% ($2098 \mu\text{rad}$)	4.7% ($1120 \mu\text{rad}$)

the injected beams have been steered. This is equivalent to about 10% of nominal current. The separation dipoles (D1 and D2) have an injection current tolerance of 1.5% of the nominal injection current. Furthermore the SIS constrains the sum of the corrector angles of the spectrometer bump to ensure a closed bump.

For the beam separation dipoles for beam 1 and beam 2, a current interlock of 3% of nominal injection current would correspond to a bend angle change of $23 \mu\text{rad}$, or 0.11% of maximum current. Consideration of Table I shows there is no danger to the experimental region if this software interlock is maintained. For the corrector magnets, a tolerance of $100 \mu\text{rad}$ corresponds to approximately 10% of maximum current. Again, consideration of Table I shows there is no danger to the experimental region if this software interlock is maintained. In [17,19] multi-point failures were also considered and the injection current thresholds were found to be sufficient.

The injection failure scenarios have been discussed and illustrated using the LHCb interaction region as an example, and complete calculations for all the IRs can be found in [16,17,19]. In summary for these injection scenarios for the IR region corrector dipoles, the software current interlocks are initially set to an equivalent kick angle of $100 \mu\text{rad}$, which is shown in [17,19] to be sufficiently tight to avoid beam accidents. Similarly, the separation dipole interlocks are initially set to 1.5% of nominal current, which again is sufficient to protect against injected beam accidents [17,19]. Furthermore, the injection interlocking provided by the BIS ensures only probe-level beams can be injected into an empty machine. Note that none of these failure scenarios are compatible with circulating beam.

The results and level of protection for ATLAS and CMS are similar, with some differences due to the TAS, and the interlocks set for the nominal machine are sufficient to protect against the injection failure scenarios described in this paper. Finally, a further source of injection turn orbit distortion can arise from misalignment of the final triplet, where a kick of $100 \mu\text{rad}$ can be caused by a misalignment of the order of 1 mm. The arguments for magnet errors apply to this distortion and the experiments should be protected from a kick of this magnitude, and further protection is provided by a SIS interlock on the positions.

B. Stored beam

A particular concern for machine protection is changes in magnetic fields around the ring arising from equipment failures and quenches, which can cause abnormal beam dynamics of the circulating beam and potential losses in elements of the machine or experiments. The impact of these circuit and magnet failures on the circulating beam can be characterized in terms of a time constant for the change, which sets the speed of the failure and which of the machine protection systems detects it first. The resulting

beam evolution calculations are in the time domain. Study of the warm separation dipole failure, evolution of the stored beam, and protection of the machine have been studied in [20], studies for the experiments have been made in [7,18], and some studies form part of [8].

There are a large number of failures that can impact a stored beam, with the variety of time constants discussed in Sec. II A. The most critical failures are those which happen on the shortest time constants and strongest magnets, for example, the failure of a normal conducting separation dipole in a region of large β function, the quench of a string of arc dipole magnets, or the failure of power converter during injection and the subsequent setting of the output voltage to the maximum value [7,8]. The resulting magnet field change will cause an orbit distortion around the LHC ring for the case of a dipole failure (including in the aperture restrictions of near-beam experiments and in collimators), a β beat and tune shift for the case of a quadrupole failure and higher order beam dynamical effects for the case of higher order element failure. For the fastest failures, particles will begin to touch the aperture as soon as several turns after the start of the failure. The time constant and the distribution of the loss depends on the class of the failure and the location in the ring. The protection against these losses takes the form of passive absorbers installed at key locations, together with active monitoring through beam loss monitors located close to aperture restrictions and monitoring devices such as magnet current monitors and FMCs (see Sec. II C). These monitoring devices are capable of issuing a beam dump request through the BIS, which can dump the beam within several hundred microseconds. These were described in Sec. II B.

We will now consider the possible stored beam accident scenarios arising from magnet circuit failures and quenches which could impact the near-beam experiments, as it is these which provide the most vulnerable small aperture restrictions. In particular, the near-beam movable detectors of TOTEM, with the distance of closest approach being 10σ , and the VELO of LHCb will be considered. For each of these experiments, the worst possible magnet circuit failures and quenches will be studied, with the time evolution of the beam calculated over many turns to assess the risk to the detectors, speed of loss, and the loss patterns. The impact of local closed bumps on the experiments will also be considered, where local corrector magnets can create a closed bump across an experimental region. These bumps reduce the available aperture of the machine and could potentially cause direct beam strike and increase susceptibility to the failure scenarios.

1. Classes and origins of failure

The magnets in the LHC are divided into families (circuits) of related magnets, with magnets in a given family connected in series. For example, the warm separation

dipoles for ATLAS (D1) are connected together in series in a circuit called RD1.LR1 (R is the code for electrical circuits). The implication is that a circuit failure results in the field decay of all magnets in that circuit. The list of circuits, along with the magnets they contain, can be found in the LHC layout database [21]. The failure of a circuit results in the decay or rise of the currents in the magnets powered (a decay from a technical failure and a rise from a circuit being anomalously driven to maximum voltage), which is modeled as a resistive-inductive circuit with some circuit-dependent time constant. Some data on the real current change exists but is not expected to change conclusions based on the analytic models [7].

The power circuits are connected to power converters, which provide the voltage across the circuit, and the peak voltage is taken to be the maximum steady voltage of the power converter. It corresponds to the applied voltage for a 7 TeV beam. In this paper the worst-case circuit failure is considered, which corresponds to a current decay from the top power converter voltage to zero if the circuit is operating for a 7 TeV beam, and an increase from the 450 GeV nominal voltage (reduced from the top voltage by a factor of 7000/450) to the top voltage of the power converter for a 450 GeV beam. The exponential decay or rise of a circuit is modeled by a scaled exponential,

$$I(t) = \left[\left(1 - \exp\left(-\frac{t}{\tau}\right) \right) \left(\frac{V_f}{V_n} - 1 \right) \right], \quad (1)$$

where τ denotes the time constant of the change, V_n denotes the nominal (initial) voltage, and V_f denotes the final voltage. The time constant for a particular circuit is found in the layout database [21] for a circuit failure, and the maximum voltage in a circuit is taken from the maximum steady voltage of the associated power converter, plus a 10% safety margin. The nominal voltage for a 7 TeV beam is taken to be the top voltage, and the nominal voltage for a 450 GeV beam is taken to be the steady voltage with a factor of 7000/450. The changing current is applied to the integrated strength of all magnets in the circuit.

The magnet field change resulting from a quench of a magnet follows the decay of the current, which is modeled by Gaussian decay with a temporal width σ of 200 ms at 7 TeV [22] and 2000 ms at 450 GeV, and decay factor defined as

$$I(t) = \exp(-t^2/2\sigma^2). \quad (2)$$

The long width means losses from quenches are very slow at 450 GeV. The decay of the current produces a corresponding decay of the integrated strength of the quenched magnet. At the same time, the magnets in the same circuit as the quenched magnet start to decay exponentially, as the quench protection circuit [5] activates to protect the magnetic systems. This decay has a time constant at 7 TeV of

40 s for quadrupoles and 104 s for main superconducting bend magnets [5] (although the time constants have been changed to 10 and 50 s, respectively, for 3.5 TeV running in 2010), and results in a corresponding drop in integrated strength of all magnets in the circuit according to Eq. (1). The decay of the current in the protection circuit follows the decay of the power converter voltage, with the top voltage being set to the maximum steady voltage from the power converted, plus a safety margin of 10%.

The time-dependent response of the LHC stored beam under the change of a powering circuit and resulting field change requires a turn-to-turn based time-dependent particle tracking code. These time-dependent studies of beam evolution and proton loss were performed using the methods in [7] with the LHC optics v6.5. This allows a time-dependent study of the beam dynamics, proton loss evolution, and time-dependent loads on collimators for a power converter failure or magnet quench.

2. Case study: Failure scenarios for the TOTEM experiment

The TOTEM [2] roman pot installation around LHC point 5 consists of two stations of silicon detectors, housed in roman pots, and located symmetrically around the IP at 147 m (145 to 149 m) and 220 m (216 and 220 m). The 147 m station is located after the TAN absorber and before the D2 separation dipole, and the 220 m station is located immediately before Q6. At both the 147 and 220 m locations, there are two roman pot units (separated by a distance of 4 m at 220 m and less than 4 m at 147 m) allowing precise determination of proton angle. Each unit consists of three pots, two which move vertically and one horizontally, with the horizontal detector on the outside of the ring. The distance of approach of the window of the roman pots to the beam is constrained by the beam halo and the setting of absorbers, and will be a minimum of 10σ from the beam. For the nominal collision optics, this is equivalent to a closest physical distance of 1 mm for the horizontal pots in the 220 m station. The details of machine optics are crucial to the success of TOTEM, and special sets of high β^* optics have been developed [2]. However, the failure scenarios presented here are done for nominal machine conditions with $\beta^* = 0.55$ m at 7 TeV and $\beta^* = 10$ m at 450 GeV, and can be redone once all the information for high β^* running conditions becomes available.

To demonstrate the protection of the near-beam experiments from hardware failures, Fig. 10 shows the time-dependent proton distribution for a 450 GeV beam with collision optics at the TOTEM roman pots located 220 m from the interaction point of CMS. At $t = 0$ the power converter RD1.LR1 powering the LSS1 separation dipoles D1 fails, causing an exponential change in the separation dipole fields and a subsequent deviation of the closed orbit around the LHC ring. In general, the decay of a bend angle will result in an orbit distortion $x(s)$ around the ring due to

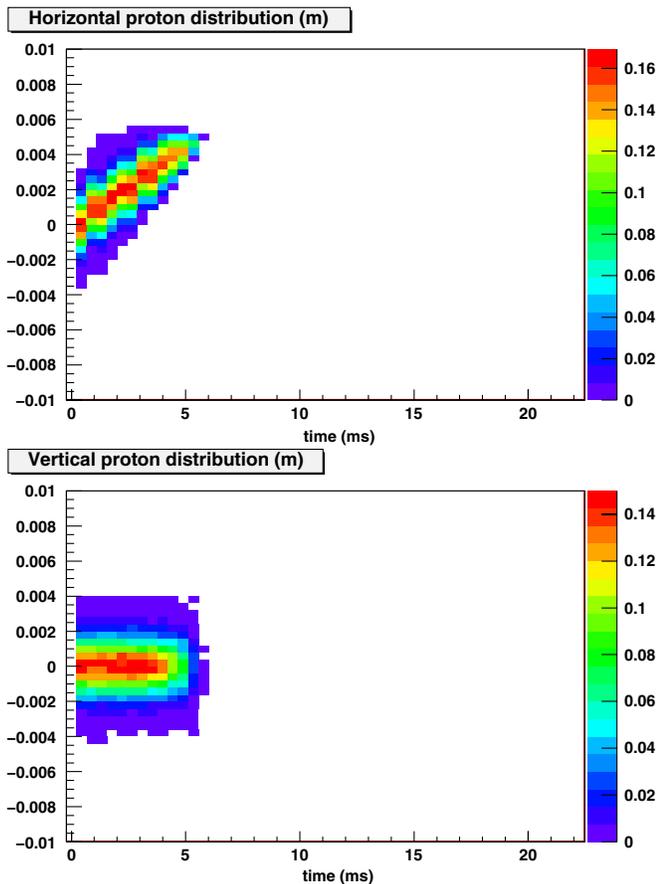


FIG. 10. (Color) The horizontal and vertical proton distribution at the TOTEM roman pots as a function of time, for a circuit error in RD1.LR1. The calculation was made for a 450 GeV stored beam with nominal collision optics. The color scale shows, in arbitrary units, the particle density in the beam.

angular kick θ_j at position j ,

$$\Delta x(s) = \frac{\sqrt{\beta(s)\beta_j}}{2 \sin(\pi Q)} \theta_j \cos[|\Psi(s) - \Psi_j| - \pi Q], \quad (3)$$

where $\beta(s)$ denotes the β function at position s and Q denotes the tune (the horizontal tune is 64.31 for collision optics). The calculations show the orbit deviation at TOTEM occurs very rapidly and beam loss begins within a few turns on the collimators. The beam is lost after about 5 ms and does not scrape the aperture of TOTEM in this time. The location of the beam loss around the ring is shown in Fig. 11, where the peaks correspond to losses on primary and secondary collimators (passive protection) in the betatron cleaning section LSS7. Therefore this passive protection relies on the correct alignment of the collimators, as described in Sec. IID. This loss would be detected on adjacent beam loss monitors, and when the applied BLM threshold was exceeded (see Sec. IIC for details) the user beam permit would be removed and the

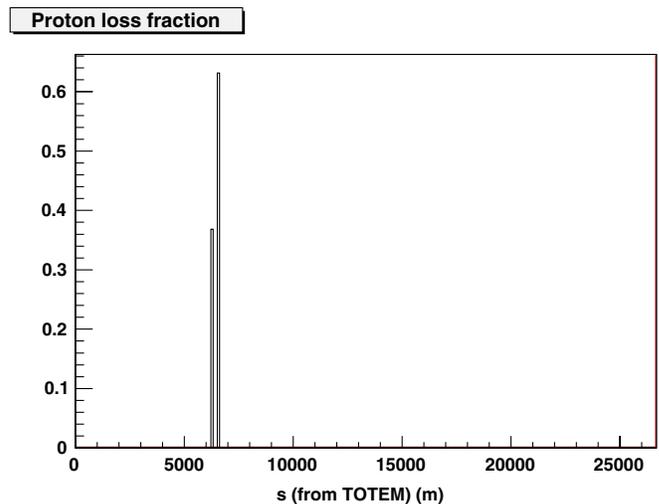


FIG. 11. (Color) The proton loss map around the ring for a circuit error in RD1.LR1. The calculation was made for a 450 GeV stored beam with collision optics. The major loss spike corresponds to loss on a primary collimator.

BIS would trigger a beam dump. Therefore the active protection systems and the BIS would detect the loss and trigger a beam dump. A further level of protection is achieved with the FMCMs, which are described in Sec. IIC. These devices have a connection to the BIS and would trigger a beam dump when the current change was detected. The resulting information flow and beam dump trigger were described in Sec. IIF. Note that there are a series of collimators positioned after the principle loss collimator TCP.B6L7.B1, with $\sim 0^\circ$ and 180° phase advances, which will catch any protons rescattering from this collimator and, hence, the probability of such a proton surviving to an orbit excursion of 10σ at the TOTEM roman pots is small. Therefore Figs. 10 and 11 show that the TOTEM pots are in the shadow of the collimations system and are protected against the failure of the power converter RD1.LR1 provided the collimators are correctly aligned. The overall protection relies on the interplay of the BIS, passive and active protection, and the magnet current monitors.

To further illustrate the role of the protection system in protecting the experiments, we will now consider the possible danger to TOTEM from a quench of a main arc dipole (with time constant 200 ms). The scenario is considered at 7 TeV, with the arc dipole chosen with a phase to TOTEM which maximizes the orbit distortion at TOTEM, Eq. (3). One such magnet is MB.B12L5.B1, which is in the circuit RB.A45. The quench has been modeled as a Gaussian decay with time constant 200 ms, and the protection circuit decay has been modeled with an exponential time constant of 104 s [5] and a top voltage of 190 V [21].

Figure 12 shows the horizontal and vertical proton density at the TOTEM 220 m roman pots for this scenario. The beams are shown as a function of time, with the failure

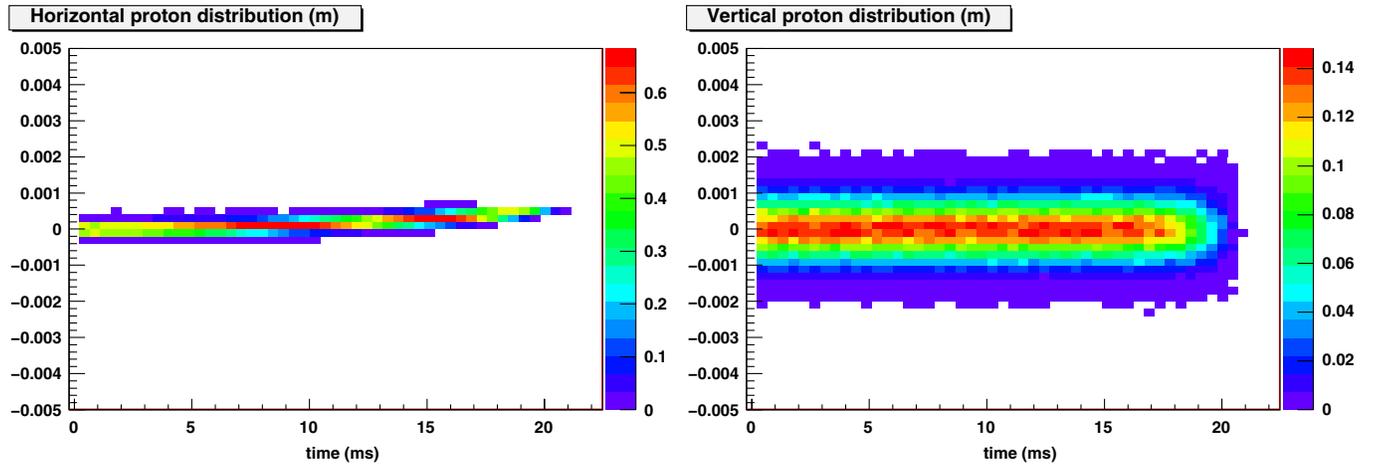


FIG. 12. (Color) The horizontal and vertical proton distribution at the TOTEM roman pots as a function of time, for a quench in MB.B12L5.B1 and resulting circuit decay in RB.A45. The calculation was made for a 7 TeV stored beam with nominal collision optics. The color scale shows, in arbitrary units, the particle density in the beam.

occurring at $t = 0$, for a stored 7 TeV beam with collision optics ($\beta^* = 0.55$ m at IP5). Hence, the crossing angles and spectrometer dipoles are turned on, the separation bumps are turned off, and the collimator apertures are set to the phase 1 values. The calculations show the beam orbit rapidly distorts in the horizontal plane, deviates to positive x (as the beam is not bent onto the design orbit as the magnet turns off and drifts to the outside of the ring), and is lost after around 20 ms. The beam does not reach the 147 or 220 m roman pot apertures in this time. Figure 13 shows the location of the proton losses around the ring, with $s = 0$ corresponding to the 220 m roman pots. The principle loss spike occurs at the secondary collimator (TCSG.A5L7.B1) and the rescattered protons from this

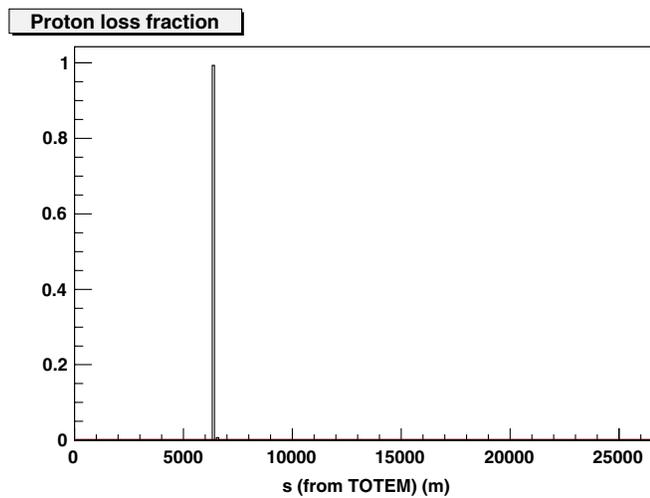


FIG. 13. (Color) The proton loss map around the ring for a quench in MB.B12L5.B1. The calculation was made for a 7 TeV stored beam with collision optics. The spike corresponds to loss on a secondary collimator.

collimator will be in the shadow of following collimators. This implies the TOTEM roman pots are in the shadow of the collimation system for a 7 TeV stored beam if the circuit worst-case magnet MB.B12L5.B1 quenches.

The loss of the beam on the collimators would be detected on adjacent beam loss monitors, and when the applied BLM threshold was exceeded (see Sec. II C for details) the user beam permit would be removed and the BIS would trigger a beam dump. Therefore the active protection systems and the BIS would detect the loss and trigger a beam dump. There are BLMs at every collimator and connected to the BIS, and the hadronic cascade of the lost proton spreads the signal over many of the downstream BLMs, with a dump time scale of several turns. A further level of protection is achieved with the quench protection system on the quenching dipole, which would also trigger a beam dump.

The implication of beam loss on the secondary collimator TCSG.A5L7, as opposed to a primary collimator, means the error scenario and resulting orbit distortion has reversed the role of the primary and secondary collimators (the jaw difference is typically $200 \mu\text{m}$ at 7 TeV, requiring this level of orbit distortion to achieve this swap). This occurs from the phase advance from the bend to the collimators, and the turn-by-turn phase advance at the primary and secondary collimators.

So the failure scenarios illustrated in this section exercise the systems described in Sec. II, and demonstrate the protection provided by the machine protection systems against failures of the considered types. A combination of signals into the BIS from active protection systems like BLMs and FMCs, together with the passive protection of the collimation system, gives good protection against the scenarios discussed and demonstrates the functionality of the systems. In [7], studies have been done for all the near-

TABLE II. Sweep parameters for asynchronous dumps under nominal conditions.

Time	Duration	Intensity	Comments
0–475 ns	475 ns	19 bunches; 2.2×10^{12}	Beam swept over machine aperture
475–1185 ns	710 ns	28 bunches; 3.2×10^{12}	Beam swept over TCDQ
1185–2250 ns	1065 ns	43 bunches; 4.9×10^{12}	Beam swept over TCDS
2.3–90.5 μ s	88.2 μ s	Rest of beam	Beam extracted

beam experiments for 450 GeV and 7 TeV collision optics stored beams, where it was shown that the experiments are always in the shadow of the collimation system for a large range of failure scenarios.

C. Beam dump failure scenarios

Several failure scenarios for the beam dump have been studied to identify and verify risks for machine elements in the ring and the extraction lines [23]. The dump system layout and components can be seen in Fig. 4 and were described in Sec. II F.

Beam dump requests can come from three different sources: the BIS systems for emergency cases, the machine timing system for scheduled dumps, and the LBDS itself. Issued dump requests from the machine timing system will be synchronized with the abort gap within the trigger synchronization units (TSU) and result in a synchronized dump. All MKD generators are connected to the retrigger system which creates a trigger signal for all generators in case of any spontaneous generator firing. This retrigger system is also triggered every 89 μ s after a dump request was issued, which bypasses the TSU and would lead to an asynchronous dump if the TSU and subsequent systems failed [24].

Two kinds of asynchronous dumps can occur during normal operation. A simple asynchronous dump is when the kickers are not fired synchronized with the 3 μ s abort gap. Therefore the rising edge of the MKD kicker waveform sweeps particles into the ring. The second kind is called a prefire dump, where one kicker fires erratically and as a consequence all others have to be fired immediately without any further synchronization to avoid that the pre-fired MKD kicks the beam directly into the TCDQ. Thus, for the prefire case the rising edge of the MKD waveform becomes shallower and more particles are transmitted into the ring than in the simple asynchronous dump case.

The rising edge of the extraction kicker waveform means around 475 ns of beam (19 bunches) is deflected with low kick angles, bypassing the TCDQ, and is swept into the machine. These bunches may, in principle, impact aperture restrictions like collimators and near-beam experiments around the ring. Therefore the simple asynchronous dump case and the prefire dump case present a possible accident scenario for the experiments. Table II is a summary of the extraction sequence and shows durations and intensities for an asynchronous dump process sorted by time. The 19 bunches swept over the machine aperture are shown in the first row.

For both the asynchronous and prefire cases, with all other systems and parameters within tolerance, beam dynamics calculations show no losses are seen anywhere in the LHC ring except for the collimators [24]. This includes the near-beam experiments of TOTEM and LHCb VELO. Therefore all bunches swept into the machine are intercepted by the primary and secondary collimators. This can be seen in Fig. 14, which shows a typical beam 1 loss distribution along the LHC for a prefire case at 7 TeV and the loss pattern shows losses only at the TCDQ and in the collimation insertion in IR7. In these cases, with the TCDQ at its nominal position, the beam load on the primary and secondary collimators remains within the specified level for this accident scenario [25]. Therefore the near-beam experiments are protected from beam loss during an asynchronous dump of beam 1 from the dump specific protection hardware and the primary and secondary collimation systems. This protection relies on the presence and alignment of the passive elements of the machine protection systems.

D. Local bumps across experiments

The available aperture for the beam in the region of a near-beam experiment can be reduced by the formation of

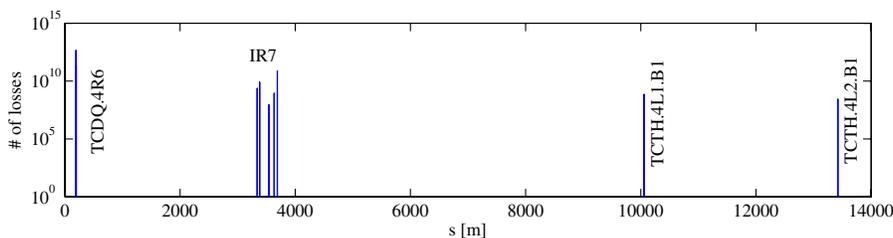


FIG. 14. (Color) Losses around the LHC for the 7 TeV prefire case, average over ten different orbit seeds.

a closed local bump, which can be created by the correctors available for global orbit correction, crossing angle, and separation bump creation. The bumps can be made by an operator or, albeit unlikely, as a result of an orbit correction algorithm. The 10σ beam distance for TOTEM corresponds to 1 mm for the horizontal pot in the 220 m station, which is the most vulnerable to local bumps. Here we consider bumps at 7 TeV, when it is most difficult to bump the beam due to the beam rigidity. The available global orbit correctors have a maximum bend angle of $90\ \mu\text{rad}$ (MCBC) or $96\ \mu\text{rad}$ (MCBY), and corrector strengths for the crossing angle have a maximum of $1010\ \mu\text{rad}$ at 450 GeV. Note that some of this strength is used for global orbit correction (typically around $10\ \mu\text{rad}$ for the global orbit correctors), leaving the rest available for local bumps.

The first plot of Fig. 15 shows a three-magnet horizontal closed bump across the 220 m TOTEM station, which is opened 199 m from the IP and closed after the TOTEM station. This bump, with corrector strengths consistent with

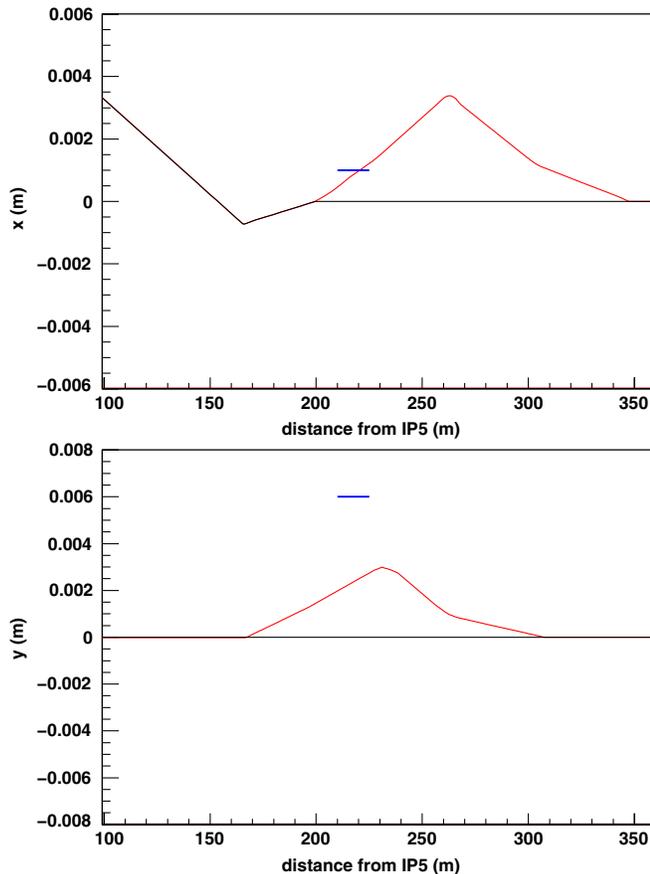


FIG. 15. (Color) The formation of horizontal and vertical local closed bumps across the TOTEM 220 m station. The available correctors can make a dangerous horizontal bump, but not a vertical bump. The black line shows the ideal orbit, the red line shows the bumped orbit, and the blue line denotes the physical boundary of the TOTEM pot.

the orbit correction, creates a horizontal orbit distortion of +1 mm at the 220 m TOTEM station, sending the beam into the horizontal pot on the outside edge of the beam pipe. The second plot in Fig. 15 shows the possible vertical bumps at 220 m, which cannot impact the vertical detectors if the vertical correctors are kept within their limits (this excludes the allowance of global orbit correction, which provides a further safety margin). The correctors used to apply the local bumps across TOTEM are slow, and change at approximately 1 microrad/s at 7 TeV and so a bump of 1 mm at TOTEM will be applied over a long time scale of many seconds. However, it should be noted that there are many ways to create closed bumps across the experiment and the subsequent reduction of aperture in this region is potentially very dangerous to the near-beam detectors.

To illustrate how the machine protection systems in this paper could be employed for protection against this beam-based scenario, the possibilities for detection and interlocking of such a closed bump are: (i) The corrector magnets around the near-beam detectors are interlocked through the SIS, to permit only a small relative change once the orbit is corrected and the moveable detectors flag is enabled in the BIS. (ii) The downstream BLMs will see a signal when the bump is being applied, and a threshold for the BLMs could be set to only permit the pot to scrape the beam halo. The BLM would remove the user permit once this threshold is exceeded, and the BIS would dump the beam. (iii) The SIS could, in principle, monitor the near-beam detector distance to the current beam orbit and give a warning if the available aperture is reduced. However, the required operating tolerance for the pots mean this is unlikely to work in practice.

Note, bumps at 450 GeV across TOTEM are easier and quicker to create than the 7 TeV case, increasing the need for detection mechanisms and interlocks. A further dangerous situation is the combination of a local bump across an experiment, reducing the available aperture for beam, combined with a fast failure of a dipole or quadrupole around the ring. This could lead to beam loss in the bumped region before loss on the collimation system, as concluded for the calculations earlier in this section. However, this combined scenario is unlikely and would be mitigated by the bump detection mechanisms discussed in this section.

IV. CONCLUSIONS

The unprecedented amounts of energy stored in the top momentum LHC beam provides a considerable risk for the elements of the machine and for the experiments. A complex machine protection system has been designed and partially tested for the LHC to provide an adequate amount of protection for the experiments. The subsystems, discussed in this paper, work together to provide a range of active and passive beam monitoring and protection, linked to a fast-response beam dump system.

This paper has presented a series of injection and circulating beam failure scenarios, shown how the machine protection systems react, and demonstrated that for the cases studied, the protection of the LHC experiments is sufficient. It is clearly very hard to protect against all possible scenarios but the level of redundancy built into the systems is high. In conclusion, the experiments of the LHC are as safe as could reasonably be expected from a wide range of failure scenarios.

There are plans to test the machine protection systems with beam once stable running of the LHC is achieved, including at a beam energy of 3.5 TeV. This is expected to give even further reassurance of the quality and reliability of the systems. Of course, the MPS systems of the LHC are very complex and will rely on the development of operational experience for successful operation. This includes the gradual increase in intensity as the machine develops, and does not rule out unexpected surprises along the way.

ACKNOWLEDGMENTS

We would like to thank all members of the LHC Machine Protection Panel, who have contributed over many years to the design and implementation of the LHC machine protection system. One of us (R. B. A.) would also like to thank members of all the LHC experiments for useful discussions throughout this work.

-
- [1] R. Schmidt *et al.*, *New J. Phys.* **8**, 290 (2006).
 - [2] CERN-LHCC-2004-002, 2004.
 - [3] LHCb TDR 5 CERN/LHCC 2001-0011, 2001.
 - [4] B. Puccio and R. Schmidt, Reports No. LHC-CIB-ES-0001 and No. EDMS 567256.
 - [5] The LHC Design Report, Vol. 1 (2004).

- [6] V. Kain, Ph.D. thesis, Vienna University, 2005.
- [7] R. B. Appleby, LHC Project Report No. 1176, 2009.
- [8] A. Gómez-Alonso, Ph.D. thesis, University of Barcelona, 2009.
- [9] O. Bruning and J. B. Jeanneret, LHC Project Note No. 141, 1998.
- [10] A. Hilaire, V. Mertens, and E. Weisse, in *Proceedings of the 6th European Particle Accelerator Conference, Stockholm, 1998* (IOP, London, 1998).
- [11] O. Bruning *et al.*, in *Proceedings of the 18th Particle Accelerator Conference, New York, 1999* (IEEE, New York, 1999).
- [12] V. Kain, J. Wenninger, and R. Schmidt, Reports No. LHC-CI-ES-0004 and No. EDMS 810607.
- [13] V. Kain and J. Wenninger, Reports No. LHC-OP-MPS-0003 and No. EDMS 889343.
- [14] A. Rossi, Vacuum Technical Note No. 04-08, 2007.
- [15] W. Herr, LHC Project Workshop Chamonix XV.
- [16] D. Bocian, LHC Project Note No. 335, 2004.
- [17] R. B. Appleby, LHC Project Report No. 1174, 2009.
- [18] R. B. Appleby, in *Proceedings of the 23rd Particle Accelerator Conference, Vancouver, Canada, 2009* (IEEE, Piscataway, NJ, 2009).
- [19] R. B. Appleby, LHC Project Report No. 1175 2009.
- [20] V. Kain, LHC Project Note No. 322, CERN, Geneva, 2003.
- [21] <http://layout.web.cern.ch/layout/>.
- [22] R. Schmidt and J. Wenninger, in *Proceedings of the 21st Particle Accelerator Conference, Knoxville, 2005* (IEEE, Piscataway, NJ, 2005).
- [23] T. Kramer, Ph.D. thesis, University TU-Graz, 2010.
- [24] E. Carlier *et al.*, in the Proceedings of the 10th International Conference on Accelerator and Large Experimental Physics Control Systems, Geneva (ICALPECS 2005), PO2.020-2.
- [25] R. Assmann, Damage Limits for LHC Collimators, Collimation WG Note, CERN, Geneva (2008).