

Protection of the CERN Large Hadron Collider

R Schmidt¹, R Assmann, E Carlier, B Dehning, R Denz,
B Goddard, E B Holzer, V Kain, B Puccio, B Todd,
J Uythoven, J Wenninger and M Zerlauth

CERN, CH-1211 Geneva 23, Switzerland

E-mail: rudiger.schmidt@cern.ch

New Journal of Physics **8** (2006) 290

Received 7 August 2006

Published 28 November 2006

Online at <http://www.njp.org/>

doi:10.1088/1367-2630/8/11/290

Abstract. The Large Hadron Collider (LHC) at CERN will collide two counter-rotating proton beams, each with an energy of 7 TeV. The energy stored in the superconducting magnet system will exceed 10 GJ, and each beam has a stored energy of 362 MJ which could cause major damage to accelerator equipment in the case of uncontrolled beam loss. Safe operation of the LHC will therefore rely on a complex system for equipment protection. The systems for protection of the superconducting magnets in case of quench must be fully operational before powering the magnets. For safe injection of the 450 GeV beam into the LHC, beam absorbers must be in their correct positions and specific procedures must be applied. Requirements for safe operation throughout the cycle necessitate early detection of failures within the equipment, and active monitoring of the beam with fast and reliable beam instrumentation, mainly beam loss monitors (BLM). When operating with circulating beams, the time constant for beam loss after a failure extends from \approx ms to a few minutes—failures must be detected sufficiently early and transmitted to the beam interlock system that triggers a beam dump. It is essential that the beams are properly extracted on to the dump blocks at the end of a fill and in case of emergency, since the beam dump blocks are the only elements of the LHC that can withstand the impact of the full beam.

¹ Author to whom any correspondence should be addressed.

Contents

1. Introduction	2
2. The LHC and its layout	3
3. Stored energy and risks of equipment damage	5
4. LHC magnets and powering	6
4.1. LHC machine cycle	6
4.2. Powering operation and quench protection	6
4.3. Quench levels	8
5. Accidental release of beam energy	8
5.1. Beam losses into material	8
5.2. Material damage test with 450 GeV LHC-type beam	9
5.3. Loss of the full LHC beam	9
6. Steady beam losses and beam cleaning system	11
7. Strategy for machine protection	11
7.1. Failure classification.	12
8. LHC machine protection elements	13
8.1. Equipment and beam monitoring systems	14
8.2. Beam Interlock System	16
8.3. Beam dumping system	18
8.4. Passive protection devices	20
8.5. Ring collimation system.	21
9. Examples of failure scenarios	23
9.1. Single-turn failure— injection kicker flashover	24
9.2. Single-turn failure— asynchronous beam dump.	24
9.3. Multiturn failure— normal conducting D1 magnets.	25
10. Operation, reliability and availability of the protection system	26
11. Conclusions	28
Acknowledgments	29
References	29

1. Introduction

The protection of accelerators with high beam power and high stored energy has become a topic of intense research in recent years [1]. Examples include the proton accelerator facility at PSI (Paul Scherrer Institute) in Switzerland [2], the spallation neutron source (SNS) at Oak Ridge, USA [3], the pulsed neutron and muon source ISIS, located at the Rutherford Appleton Laboratory, UK [4], future linear colliders (for example ILC), and the Large Hadron Collider (LHC), where both the momentum and the beam intensity reach unprecedented values.

For many accelerators, the high stored beam and magnet energy must be safely discharged when requested or after a failure. In case of uncontrolled release of energy, repair of damaged equipment would require long stops for repair. For example, in LHC the exchange of a superconducting magnet is expected to require about 30 days. In the worst case, many accelerator components (magnets, collimators, etc) could be severely damaged.

There are several general requirements for the protection systems:

1. Protect the accelerator equipment: the first priority is to protect equipment from damage, in the LHC ring, and during the transfer from the pre-accelerator SPS to the LHC. The second priority is to protect superconducting magnets from quenching.
2. Protect the beam: protection systems should only dump the beam when necessary. Unnecessary ('false') beam dumps should be avoided in order not to compromise availability.
3. Provide the evidence: in case of failure, complete and coherent diagnostics data should be provided to accurately understand what caused the failure and if the protection systems functioned correctly.

This paper introduces the LHC and addresses the risks during operation due to the large amount of energy stored in the magnet system and in the beams. The strategy for machine protection is discussed and the systems for protection are described.

2. The LHC and its layout

The motivation to construct the LHC at CERN comes from fundamental questions in particle physics. A deeper knowledge in particle physics today is linked to the understanding of particle mass scales: is there an elementary Higgs boson? The primary task of the LHC is to make an exploration in the TeV energy range. To reach the 1 TeV scale in the centre-of-mass of proton constituent collisions, a proton collider with two beams of $7 \text{ TeV } c^{-1}$ momenta is being constructed in the 23 km long tunnel that had previously been used for LEP (Large Electron Positron accelerator). A magnetic field of 8.33 T is necessary to achieve the required deflection of $7 \text{ TeV } c^{-1}$ protons, which can only be generated with superconducting magnets. The machine is also designed for collisions of heavy ions (for example lead) at very high centre-of-mass energies.

The LHC accelerator has been under preparation since the beginning of the 1980s, with a research and development program for superconducting dipole magnets and the first design of the machine parameters and lattice. The LHC accelerator is being constructed in collaboration with laboratories from both CERN member and non-member states and regular beam operation is scheduled to start in 2007. A design report was published in [5] and the main parameters are given in table 1.


The search for the Higgs boson and for other particles requires the LHC to provide proton–proton collisions at the centre-of-mass energy of 14 TeV and a luminosity in the order of $10^{34} \text{ cm}^{-2} \text{ s}^{-1}$, a significant challenge for the accelerator. The LHC will operate with high-field dipole magnets using NbTi superconductors cooled with superfluid helium at 1.9 K.

For the LHC, the operation with large stored energy in the beams in the presence of superconducting magnets with a very low quench margin is a particular challenge for the collimation system, which needs to have an unprecedented efficiency [6]. To achieve this efficiency, the position of close to 200 collimator jaws needs to be optimized in order to capture particles with large amplitudes that would otherwise be lost around the accelerator.

The LHC has an 8-fold symmetry with eight arc sections, and eight straight sections (also referred to as insertions) which contain experiments and systems for machine operation, see figure 1.

Table 1. Main LHC parameters.

Momentum at collision	7	TeV c^{-1}
Dipole field for 7 TeV	8.33	T
Luminosity	10^{34}	$\text{cm}^{-2} \text{s}^{-1}$
Protons per bunch	1.15×10^{11}	
Number of bunches/beam	2808	
Normalized emittance	3.75	μm
Typical rms beam size in arcs at 7 TeV	200–300	μm
Beam pipe diameter	56	mm
Nominal bunch spacing	25	ns

**Figure 1.** Schematic layout of the LHC machine.

Two counter-rotating proton beams will circulate in separate vacuum chambers installed in the twin-aperture magnets, with the beams crossing at four points. The total path length is the same for both beams. The beams will collide at a small angle in the centre of the main experimental detectors (ATLAS, ALICE, CMS and LHCb) installed in four of the straight sections. The injection elements are installed in the insertions where ALICE and LHCb are located. Of the remaining four insertions, two contain the beam collimation (also referred to as beam cleaning) systems, one contains the beam dumping systems, and one contains the radio-frequency (RF) acceleration system and specialized beam instrumentation.

The LHC beams are prepared in the CERN accelerator complex [7], pre-accelerated in the SPS from 26 to 450 GeV and transferred through two lines of about 3 km length to the LHC injection points.

Table 2. Energy stored in magnets and beams.

Energy stored in magnet system	10	GJ
Energy stored in one main dipole circuit	1.1	GJ
Energy stored in one beam	362	MJ
Beam power averaged over the length of a fill of 10 h, both beams	20	kW
Beam power averaged over one turn, one beam	3.9	TW
World Net Electricity Generation (2002)	1.7	TW
Energy to heat and melt one kg copper	700	kJ

3. Stored energy and risks of equipment damage

The energy stored in each dipole magnet powered at a current I of 12 kA is 7.6 MJ, given by $E_{\text{magnet}} = 0.5 \cdot L \cdot I^2$, where L is the magnet inductance. The magnetic energy stored in all LHC magnets is about 10 GJ as shown in table 2. An energy of 10 GJ is equivalent to the energy content of about 230 kg gasoline [8]—with this amount of energy, it is possible to heat up and melt nearly 15 tons of copper.

In order to safely handle the energy stored in the magnet system, the LHC magnets are powered in several independent powering sectors. This reduces the energy stored in one electrical circuit by about one order of magnitude, and the systems for the protection of the superconducting magnets are thus similar to other accelerators with superconducting magnets (HERA, TEVATRON, RHIC).

In the LHC, the proton momentum at 7 TeV c^{-1} is a factor of seven to sixteen above accelerators such as SPS, TEVATRON and HERA. The energy stored in the beams, for each beam 362 MJ (corresponding to the energy content of about 100 kg TNT), is larger by a factor of about 200 due to the very high beam intensity, see table 2 and figure 2. Due to the small beam dimensions, the maximum energy density as the relevant factor for equipment damage is a factor 1000 higher than for other accelerators. Compared to accelerators producing high beam power (e.g. the 1 MW cyclotron at PSI [2]), the LHC beam power averaged over the length of a fill of 10 h is with 20 kW relatively small.

At other accelerators with large energy stored in the beams (HERA [9], TEVATRON [10] and SPS [11]) accidents have already happened involving uncontrolled beam loss that have caused damage of accelerator components and led to significant downtime.

The exact beam intensity required to damage equipment depends on the impact parameters, beam size and on the equipment hit by the beam. An uncontrolled loss of even a very small fraction (below 10^{-4}) of the 7 TeV LHC beam could already cause damage to equipment. The beam that is injected from the pre-accelerator SPS at 450 GeV can already damage equipment. Protection must be efficient from the moment of extraction from the SPS, throughout the LHC cycle.

The beams must be handled in an environment with superconducting magnets that could quench in case of fast 7 TeV beam losses of 10^{-8} – 10^{-7} of the nominal beam intensity. This value is orders of magnitude lower than for any other accelerator with superconducting magnets and requires very efficient beam cleaning.




Figure 2. Energy stored in the beams for different accelerators (based on a figure by R Assmann, CERN-AB-2003-008-ADM).

4. LHC magnets and powering

4.1. LHC machine cycle

The beams are accelerated in the SPS from 26 to 450 GeV c⁻¹, and then transferred to the LHC. This energy corresponds to a magnetic field of the LHC dipoles of about 0.54 T [12]. During the injection phase, 12 batches per beam (one batch has either 216 or 288 bunches) are injected from the SPS into the LHC (see the LHC operational cycle in figure 3). Injection of the two beams will take about 15 min. Then the field of the LHC dipole magnets is ramped within 28 min to 8.33 T corresponding to a beam momentum of 7 TeV c⁻¹. Normally, the beams will collide for several hours during a physics fill. At the end of a fill or after the detection of a failure, the beams are extracted from the accelerator into specially designed absorbers (beam dump blocks) thus discharging the energy. Finally the magnets are ramped down to a field slightly below injection level to prepare for the next injection, and just before the next injection sequence the magnets are ramped up to their injection level.

4.2. Powering operation and quench protection

The LHC is complex, with more than 10 000 magnets powered in 1612 electrical circuits, large and distributed cryogenics and vacuum systems etc.

The electrical circuits with the largest stored energy power the 1232 main dipole magnets. There are eight power converters for the main dipoles, each powering 154 dipoles in one arc. During the magnetic ramp, the energy is taken from the powering grid. The time for normal discharge (ramping down the magnets) is about the same as for ramping up and the energy is delivered back to the grid.




Figure 3. LHC cycle, showing the various phases of operation of the accelerator at different energies (courtesy O Brüning *et al* 2005 LHC Project Note 313).

For a given magnetic field and current density in the cable, the temperature must not exceed the critical temperature, otherwise the magnet quenches. There are several mechanisms that can lead to a quench. Quenches are initiated by an energy deposition in the order of mJ, for example due to beam losses. Other mechanisms for inducing a quench are movements of the superconductor by several micrometre (friction and heat dissipation), or a failure in the cooling system.

Without protection, after a quench, the temperature in the resistive zone would increase within less than 1 s to 1000 K, and the magnet would be destroyed. Since quenches cannot be avoided, sophisticated protection systems are required [13].

In the first place the quench has to be detected and the power converter has to be switched off. The energy of the quenching magnet is then distributed in the magnet coils by firing quench heaters that force-quench the superconductor. The heaters are stainless steel strips mounted along the magnet coils. After a quench at 7 TeV, the current in a dipole magnet decays with a time constant of about 200 ms, approximately following a Gaussian. The temperature of the magnet coil increases by some 10 K.

If the magnet in an electrical circuit is powered in series with other magnets, the energy stored in the non-quenching magnets must also be extracted. For the LHC arc dipole circuit where 154 dipole magnets are powered in series, the natural time constant for the current decay is many hours. To reduce the time constant for the discharge to about 100 s, resistors are switched in series with the magnet string. During the discharge, the current passes through the cold bypass diode of the magnet that quenched. The energy is safely absorbed in the resistors, heating eight tons of steel to about 300°C.

While quench heaters and diodes will be used for the protection of the LHC main dipole and quadrupole magnets [14], many other stand-alone magnets require protection by quench heaters only.

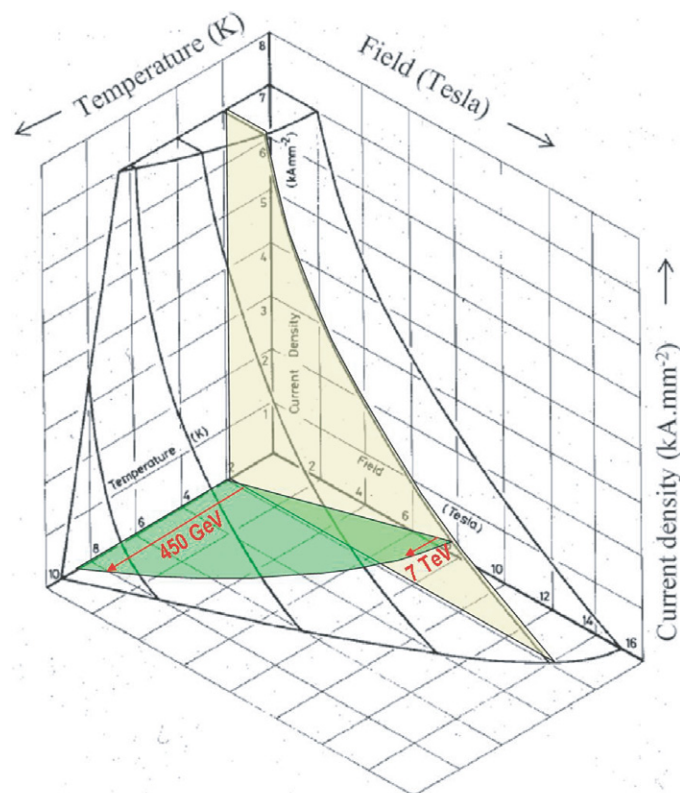


Figure 4. Operational margin (green surface) of a dipole magnet, showing difference between 450 GeV and 7 TeV.

4.3. Quench levels

To operate the LHC, beam losses must be kept below the quench level of the magnets, since recovery after a quench to re-establish the conditions for beam operation will take several hours. At injection energy of 450 GeV an instantaneous loss of about 5×10^9 protons is expected to quench a dipole magnet, while at top energy of 7 TeV only about 5×10^6 protons can cause a quench [15] due to the reduced temperature margin of the superconductor (see figure 4) and the increased momentum of the protons.

5. Accidental release of beam energy

5.1. Beam losses into material

Beam impact in material produces particle cascades due to nuclear and electromagnetic interactions. The local energy deposition and temperature increase depend on the material and on the number, energy and spatial distribution of the impacting particles. The energy deposition of a particle distribution as a function of material and geometry is calculated with programs such as FLUKA [16]. The temperature increase in the material is then estimated using the temperature-dependent specific heat capacity.

A simple approximation for the temperature increase in material for a 7 TeV beam impact is given in the following example: for copper, the maximum longitudinal energy deposition for a




Figure 5. Damage of a copper plate by 450 GeV c^{-1} proton beams with different intensities.

single 7 TeV proton at about 25 cm inside the material is $E_{\text{dep}} = 1.5 \times 10^{-5} \text{ J kg}^{-1}$ (calculation with FLUKA). The energy required to heat and melt copper is $E = 6.3 \times 10^5 \text{ J kg}^{-1}$. Assuming a pencil beam, the number of particles required to damage (melt) copper is in the order of 10^{10} . For graphite, the number of particles needed to cause damage is about one order of magnitude larger. More refined and complete calculations are made to determine real-world scenarios on a case-by-case basis, where the distribution of the impacting particles and the details of the material and geometry are very important.

5.2. Material damage test with 450 GeV LHC-type beam

The design of LHC protection elements is based on detailed energy deposition simulations and an assumption for the damage levels. A dedicated experiment was carried out to cross-check the validity of this approach by trying to damage material in a controlled way with beam [17]. The impact of a 450 GeV beam extracted from the SPS on a specially designed high-Z target was simulated for a simple geometry comprising several typical materials used for LHC equipment. The beam intensities for the test were chosen to exceed the damage limits of parts of the target, between 2×10^{12} and 8×10^{12} protons. The transverse r.m.s. beam dimensions were about 1 mm.

The geometry of the target was modelled in FLUKA and the target heating was estimated. The temperature rise was obtained from the energy deposition using the temperature-dependent heat capacity for each material. The results of the controlled damage test show reasonable agreement with the simulations. Zn-, Cu- and INCONEL-plates are damaged at the predicted locations within the error of the simulation, see figure 5. The transverse extent of the damaged area on the Zn- and Cu-plates was predicted to within 30%. The outcome of the experiment gives confidence that beam-induced damage limits for simple geometries can be adequately predicted with simulations.

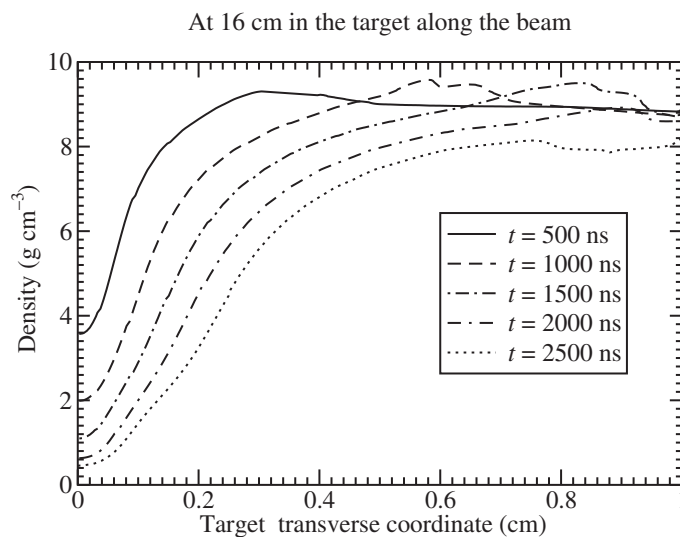
Table 3 summarizes present estimates for the beam intensity below quench and damage levels at the LHC.

5.3. Loss of the full LHC beam

One of the worst case failure scenarios, an accidental release of the entire LHC beam energy into equipment, has been considered in the study described in [18]. The damage has been

Table 3. Bunch intensities, quench and damage levels for the LHC.

Intensity one 'pilot' bunch	5×10^9 protons
Nominal bunch intensity	1.1×10^{11} protons
Nominal beam intensity,	2808 bunches, 3×10^{14} protons
Nominal batch from SPS (at 450 GeV)	216/288 bunches, 3×10^{13}
Intensity below damage level for fast losses at 450 GeV	$\approx 1-2 \times 10^{12}$ protons
Intensity below damage level for fast losses at 7 TeV	$\approx 1-2 \times 10^{10}$ protons
Intensity below quench level for fast losses at 450 GeV	\approx some 10^9 protons
Intensity below quench level for fast losses at 7 TeV	\approx some 10^6 protons

**Figure 6.** Density distribution along the transverse direction at different times at a longitudinal position of 16 cm into the target [18].

estimated for a solid copper target hit at normal incidence by the full LHC beam, by carrying out three-dimensional (3D) energy deposition calculations and 2D numerical simulations of the hydrodynamic and thermodynamic response at different longitudinal positions in the target.

If instantaneous energy deposition were assumed, the energy density deposited in the material would exceed the energy that is required for vaporization by several orders of magnitude. However, the beam energy is deposited over $86 \mu\text{s}$, long enough to change the density of the target material. This density change strongly affects the energy deposition of the impacting beam. The calculations indicate that the target density around the beam axis can be reduced by more than a factor of 10 within $2.5 \mu\text{s}$, due to the transverse shock wave moving outwards from the region heated by the beam. The material in a 0.5 cm radius around the region heated by the beam is subject to substantial expansion (figure 6). The material in this hot inner zone is in a plasma state, with the surrounding target in a liquid state. The density is reduced by a factor of 10 after only 100 LHC bunches out of 2808 have been delivered. The protons in the following bunches will therefore penetrate into the target more deeply as they will encounter material with reduced density—the penetration depth is estimated at up to 40 m.

Table 4. Lifetime of the LHC beams.

Single beam lifetime	Beam power lost into accelerator (per beam)	Comment
100 h	1 kW	Healthy operation
10 h	10 kW	Acceptable operation - collimators must absorb large fraction of losses
12 min	500 kW	Operation only possibly for short time - collimators must be very efficient - not necessary to dump the beam
1 s	330 MW	‘Slow’ failure e.g. trip of superconducting magnets - detect equipment failure, beam must be dumped
15 turns	Several 100 GW	‘Fast’ failure e.g. trip of normal conducting magnets - detect beam losses, beam dump as fast as possible
1 turn	Several TW	‘Ultra-fast’ failure, e.g. at injection/beam dump - detection and reaction impossible - passive protection relies on collimators

6. Steady beam losses and beam cleaning system

As is known from other proton storage rings, the lifetime of a non-colliding LHC beam could exceed 100 h (table 4). Such a lifetime corresponds to a power deposition due to lost particles of about 1 kW distributed over the 27 km long LHC circumference. Such a power deposition does not pose any problem to the cryogenic cooling system. Nevertheless, a beam cleaning system is required to capture a large fraction of the beam losses thus preventing magnets from quenching due to local peaks in the particle loss distribution.

During colliding beam operation with a luminosity of $10^{34} \text{ cm}^{-2} \text{ s}^{-1}$, the lifetime of the beam is dominated by the collision losses due to the high proton–proton cross-section. 10^9 protons are lost per beam per second in each of the two high luminosity insertions (ATLAS and CMS experiments). This corresponds to a local power deposition of several kilowatts. Part of the power is deposited into the close-by superconducting quadrupole magnets that have been designed to accept high heat load. Heavy shielding is required for these insertions.

For nominal beam intensity at 7 TeV, assuming a single beam lifetime of 10 h, the collimation system must capture more than 99.9% of the beam losses [6]. If the lifetime decreases down to 0.2 h, the power deposition increases to 500 kW. Since periods with such short lifetimes must be anticipated during operation of such a complex accelerator like the LHC, the collimators are designed to accept such lifetime drops for 10 s. If the lifetime becomes even shorter, in particular after equipment failure, the beams must be dumped.

7. Strategy for machine protection

Protection at the LHC relies on a variety of systems with strong interdependency (figure 7). The main principles of the strategy are:

1. Definition of the aperture by the collimator jaws, with beam loss monitors (BLMs) close to the collimators.

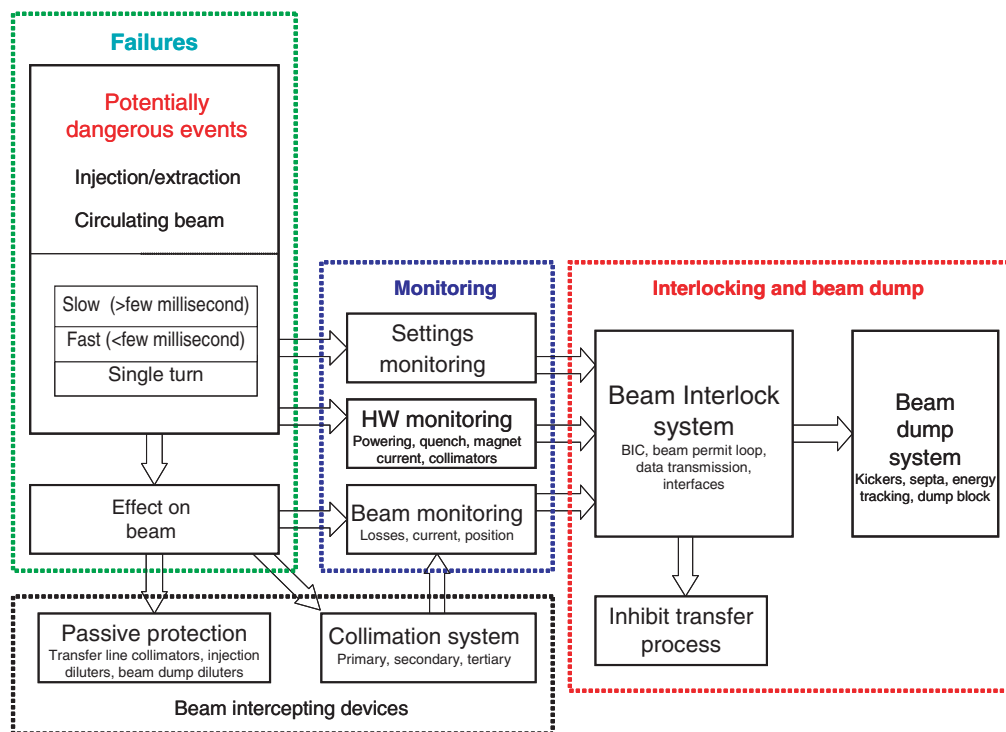


Figure 7. The functional blocks of the LHC machine protection system.

2. Early detection of failures within the equipment that acts on the beams, to generate a beam dump request before the beam is affected.
3. Active monitoring of the beam 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 ($89 \mu\text{s}$).
4. Reliable transmission of a beam dump request to the beam dumping system by a distributed interlock system. For all interlocks, an active signal is required for operation, and the absence of the signal is considered as beam dump request and injection inhibit.
5. Reliable operation of the beam dumping system on receipt of a dump request or internal fault detection, to safely extract the beams on to the external dump blocks.
6. Passive protection by beam absorbers and collimators for specific failure cases.
7. Redundancy in the protection system such that failures may be detected by more than just a single system.
8. Very high safety and reliability standards that are applied in the design of the core protection systems.

7.1. Failure classification

Beam losses due to a failure can occur in a single turn, over multiple turns or during a much longer timescale and are therefore classified as ultra-fast losses (from one turn failures), very fast losses in less than 5 ms or fast losses in more than 5 ms (from multi-turn failures), and steady losses of one second or more (from slow losses) [19].

7.1.1. Single turn failures (ultra-fast losses). Such events are due to failures at beam injection, at beam extraction or during operation of dedicated diagnostic kicker magnets. The probability for kicker magnet failures is minimized by designing the systems for high reliability and by interlocking kickers (for example, the injection kickers will be switched off when the LHC is not at injection energy). The strength of the kickers for diagnostics is limited to a small amplitude, and the operating conditions are constrained by interlocking. The magnet systems where the field could change rapidly must be monitored with tight tolerance windows, and the designs of the kicker systems must be as reliable as possible. Despite these precautions, it is known from other accelerators that failures such as pre-firing of kicker magnets cannot be excluded. In these cases, active protection based on fault detection and reaction is not possible because the failure occurs on a timescale that is smaller than the minimum detection and dump time. Protection from such specific failure cases relies therefore on beam absorbers and collimators that need to be correctly positioned close to the beam to capture the particles that are deflected accidentally.

Another source of single-turn failures at injection are the magnets in the LHC, since a wrong current value or aperture problem (for example due to a closed vacuum valve) would cause the injected beam to be lost. To prevent such an eventuality, only a ‘pilot’ beam with low, non-damaging intensity can be injected when there is no beam circulating in the LHC. Injection of beam exceeding this intensity requires some circulating beam to be present [20], verified by beam current measurements just before injection.

7.1.2. Multiturn failures (very fast and fast losses). The majority of equipment failures lead to beam ‘instabilities’ (fast movements of the orbit or beam size growth) that must be detected on the timescale of 1 ms or more. Protection against such events relies on fast beam loss monitoring and beam position monitoring. Equipment monitoring (e.g. quench detectors, and monitors for failures of power converters) provides redundancy for most cases.

7.1.3. Steady losses. During normal operation, there will be unavoidable losses around the LHC machine. The beam losses and heat load at collimators will be monitored, and if the losses or heat load become unacceptably high, the beam will be dumped. The temperature of the superconducting magnets will be monitored, and if a temperature increase risks leading to a quench, the beam will be dumped.

8. LHC machine protection elements

Protection of the LHC relies on several systems which must work reliably together to detect faults, to transmit beam dump requests, to dump the beams and to intercept any mis-steered particles. This section presents the various elements that together provide the machine protection for the LHC:

1. Systems that monitor either directly equipment or beam parameters and that are able to generate a beam dump request.
2. A beam interlock system to provide a highly reliable transmission of the dump request from the monitoring systems to the beam dumping system.

3. A beam dumping system that is designed to extract the beams from the LHC ring and dispose the beams on to the beam dump blocks.
4. The collimation system defines the aperture and protects against quenches of superconducting magnets. Beam absorbers and collimators for specific failures must be positioned correctly with respect to the beam.

8.1. Equipment and beam monitoring systems

8.1.1. Quench protection system. As described in subsection 4.2, when a main magnet quenches and the resistive voltage across the magnet exceeds a threshold of 0.1 V for 10 ms, the energy stored in the quenched magnet is discharged into the coils by firing quench heaters and the energy extraction process is initiated. At the same time a beam dump request is issued. Since it takes between 15 and 300 ms until the quench heaters become effective and the extraction switch opens, the beam is normally dumped before the magnetic field decays.

The quench protection system will be tested prior to and fully commissioned during the first powering of LHC without beam. It will be fully operational for the start-up with beam.

If the event of a failure in the quench protection system, for example when the heaters are fired without prior quench, this will lead to a beam dump before the beam is affected.

8.1.2. Fast magnet current change monitors. A failure in normal conducting magnets can lead to a change of the magnet current and to beam losses in a very short time. It is of particular importance to ensure that the conditions for beam injection or extraction are correct just before the firing of kicker magnets. To detect fast powering failures in a sufficiently short time, a device developed for HERA is being adapted for the LHC [21]. This instrument is able to detect current changes at the level of 0.03% within about 0.7 ms for the LHC, and at the level of 0.1% within 50 μ s for the extraction septum magnets that are used for beam transfer from SPS to LHC. Consequences of thunderstorms on the electrical distribution that have an impact on the magnet current should also be detected by these monitors, which are used for 15 of the electrical circuits powering the most critical magnets.

8.1.3. Hardware diagnostics. For many systems a beam dump request is issued in case of hardware failures via the beam interlock system:

1. Vacuum valves or other movable devices leaving the 'OUT' position and moving towards the beam.
2. Trip of a power converter (typical detection time > 10 ms).
3. Too high temperature for normal conducting magnets.
4. Failure in the RF system.
5. Wrong position or excessive heat load for critical beam absorbers and collimators.

Anticipated failure in the beam dumping system will also generate a beam dump (for example: if one out of 15 extraction kicker magnets failed, the beam would be dumped since 14 magnets would still ensure a clean beam dump).

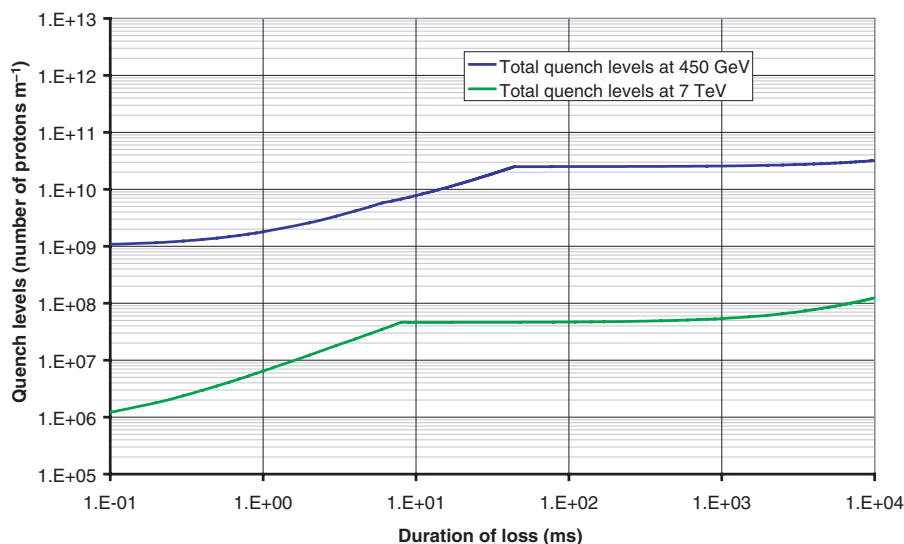


Figure 8. Quench levels versus loss duration [22].

8.1.4. BLM system. BLMs are installed at each quadrupole magnet of the LHC (more than 3000 gas ionization chambers) and all aperture limitations around the machine, in particular at the collimators (several hundred gas ionization chambers and secondary emission monitors) [22].

Figure 8 shows the expected quench levels for superconducting magnets, as a function of the duration of the beam loss, for 450 GeV and for 7 TeV [22]. For the BLMs that will be installed in the superconducting part of the LHC, the threshold will be adjusted in order to request a beam dump before the beam loss quenches a magnet.

The electronics integrates the signal from the detectors [23, 24]. There are 12 integration windows, ranging from about 40 μ s to many seconds. The monitors can detect accidental beam loss within one turn. The integrated value in each window is compared with predefined thresholds that depend on the LHC energy. If the threshold is exceeded, a beam dump request is issued.

Since collimators must always delimit the aperture, beam loss should always be observed with the monitors close to collimators, for example if the emittance grows or if the beam becomes unstable. Particles scattered from the cleaning insertions into the arcs will be detected by BLMs in the arcs.

Together with the BLMs at aperture limitations, the arc monitors cover the entire accelerator, detecting beam loss independently of collimator settings. Arc monitors detect losses that appear only in the regular lattice cells, for example due to closed orbit bumps. The complete system has a large complexity and risks producing a number of false beam dumps [25]. The monitors are therefore grouped into two categories:

1. Critical monitors that must be fully operational to inject beam into the LHC. The BLMs installed at the collimators and at selected aperture limitations belong to this category. The number of monitors in this class is limited to a few hundred.
2. Standard arc monitors that are considered less critical. Due to their large number (over 3000) and complexity, it is acceptable to operate the LHC with beam when a few monitors are not operational.

Four additional BLMs will be installed close to beam absorbers in the beam dumping insertion. The loss monitors will be hardware configured with relatively high thresholds, with direct inputs to the triggers of the LHC beam dumping system. The trigger signals will not pass via the Beam Interlock System. The interlock threshold levels are independent of the LHC energy and must be chosen with care, to provide adequate protection while avoiding spurious beam dumps. This system is intended to alleviate the effects of a ‘catastrophic’ failure, probably restricting damage to the LHC collimation system.

8.1.5. Beam position change monitors. For failures that lead to a fast movement of the closed orbit, a system using Beam Position Monitors (BPMs) for fast surveillance of the beam position is under development. The BPMs are grouped in two pairs for redundancy. Two pairs are separated in betatron phase by 90° to allow detection of orbit oscillations of any phase. The betatron function at both locations is large, giving an enhanced sensitivity.

To avoid perturbations from slow orbit movements, the interlock logic will be based on the detection of a position change with respect to the last measured closed orbit. At 450 GeV, the fastest orbit movement during normal operation by an orbit corrector magnet is in the order of some millimetres per second. At 7 TeV, the fastest movement is less than 1 mm s^{-1} . If the change of the orbit exceeds substantially these values, the beam should be extracted. To provide protection against the fast orbit movements due to the most critical powering failures, the position interlock thresholds must be set to 1 mm ms^{-1} . The reaction time of the system will be around 1–2 turns [26].

The advantage of this system is that it provides a direct measurement of fast orbit drifts, possibly detecting failures before the beam touches the aperture. The system is independent of collimator settings and has a limited complexity.

8.1.6. Beam current decay measurements. The damage limit for fast beam impact on materials such as copper on stainless steel is in the range of some 10^{12} protons at 450 GeV, and around 10^{10} protons at 7 TeV (see table 3). As an alternative to the detection of local losses by BLMs, it is envisaged to monitor directly the total beam intensity in the machine. At 450 GeV, a beam current transformer able to reliably detect a loss of 10^{11} protons on the timescale of one to several milliseconds would detect any dangerous beam loss and dump the beam before damage. At 7 TeV, provided the collimators are correctly positioned and are hit first, detection of beam loss would prevent serious damage of the LHC, although beam loss could damage the surface of some collimators before extraction is complete. If the collimators are not in the correct position, then detecting such a loss would still protect the LHC at 7 TeV for many failure scenarios, and could reduce any damage by orders of magnitude.

A similar system is operating at HERA [9]. It remains to study what level of detection can be achieved with such a system with the required reliability.

8.2. Beam Interlock System

The role of the Beam Interlock System is to concentrate beam dump requests from the various monitoring systems described above and to transmit such requests to the beam dumping system [27].

Two fundamental requirements for the Beam Interlock System dictate architecture and design, namely signal transmission delay and dependability. To protect the LHC against fast

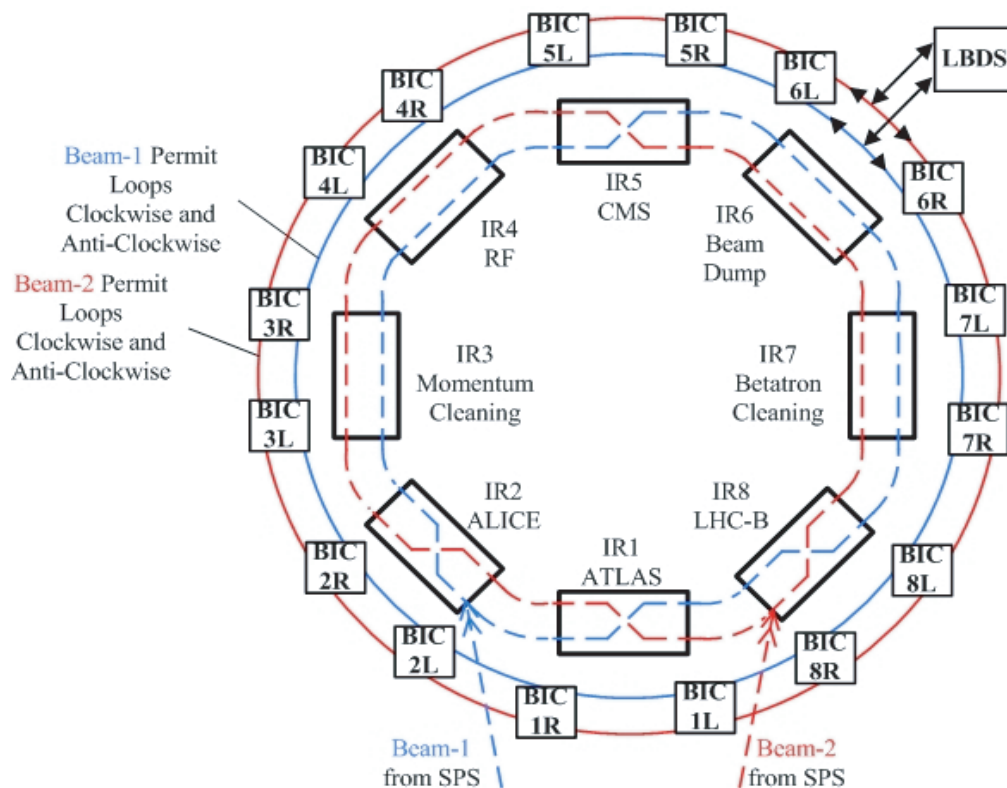


Figure 9. Architecture of LHC Beam Interlock System, showing Beam Interlock Controllers (BIC) and Beam Permit Loops for all LHC insertions [29].

beam losses, the system must transmit dump requests from connected systems around the LHC circumference to the beam dumping system as fast as possible, within less than $\approx 100 \mu\text{s}$.

The system must have very high reliability and is specified to be better than Safety Integrity Level-3 (SIL-3) [28]. The interlock system is based on fail safe logics providing a continuous beam permit signal to the beam dumping system. In order to give beam permit and inject beam, all connected systems must provide their permit signals. A missing beam permit leads to a beam dump request. Redundant signal paths have been used for the beam permit, making the system both safe and allowing a full system test to be carried out on demand, by closing only one beam permit loop for each signal path at the time. The mean time between a failure that prevents the beam dump request to be correctly transmitted has been evaluated between 1000 and 10 000 years [29].

There are around 140 distributed ‘user systems’ that can request a beam dump. To accommodate the user inputs, the LHC Beam Interlock System has 16 Beam Interlock Controllers (BIC), one situated to the right and to the left of each Insertion Region (IR). The controllers communicate with the beam dumping system situated in IR6. The permissions for beams to be present in the LHC, called Beam Permits, are carried around the machine between these controllers by beam permit loops (figure 9). The ‘Beam Permit’ signals may be in the state TRUE (signal present) or FALSE (signal absent).

The communications from one point to another are carried out over four dedicated fibre optic channels. To reach the high level of safety, a clockwise and anti-clockwise link exists for each beam. The request for a beam dump always takes the shortest path from a controller to the

beam dumping system. On each of the Beam Permit Loops, a 10 MHz square wave signal is generated at IR6 beside the beam dumping system, this signal travels around the Beam Permit Loop and through each BIC. Each BIC can monitor the beam permit loops and open the loops to request a beam dump. A correct frequency being detected at the end by the beam dumping system represents 'Beam Permit = TRUE' given for all user systems connected to that loop.

If the beam dumping system detects a change of Beam Permit from TRUE to FALSE, then the corresponding beam is extracted into the beam dump blocks. The beam permit loops representing the two beams operate completely independently. Connected systems can dump the two beams either independently, or simultaneously. The relevant beam permit signals are ANDed together to allow or deny the propagation of the 10 MHz square wave. The BIC also returns beam permit status information back to the connected systems.

As upgrade for the pre-accelerators and to verify the design and the expected safety and reliability of the Beam Interlock System, an interlock system with identical hardware has been installed for the CERN SPS and for the transfer lines from SPS to LHC. The system is also being used for the CERN Neutrinos to Gran-Sasso (CNGS) project that started operation recently.

8.3. Beam dumping system

The role of the LHC beam dumping system is to safely dispose of the beam when beam operation must be interrupted for any reason. Fifteen fast kicker magnets with a pulse rise-time of less than $3 \mu\text{s}$ deflect the beam by an angle of $280 \mu\text{rad}$ in the horizontal plane (figure 10). To ensure that all particles are extracted from the LHC, the beam has a particle free abort gap with a length of $3 \mu\text{s}$ corresponding to the kicker rise-time. The extraction kicker is triggered such that the field increases from zero to the nominal value during this gap when there should be no particles. Downstream of the kicker the beam is deflected vertically by 2.4 mrad towards the beam dump block by 15 septum magnets. A short distance further downstream, ten diluter kicker magnets are used to paint the bunches in both horizontal and vertical directions to reduce the beam density on the dump block (figure 11). The beam is transferred through a 700 m long extraction line to increase the transverse r.m.s. beam size from approximately 0.2 to 1.5 mm and to spread the bunches further on the dump block. The projection of the beam on the beam dump block is shown in figure 12 [30, 31]. The overall shape is produced by the deflection of the extraction and dilution kickers. For nominal beam parameters, the maximum temperature in the beam dump block is expected to be in the order of about 750°C .

The beam dumping system must function with utmost reliability, since any failure during the dump process can have catastrophic consequences, and is specified as having SIL-4 [28], which has been confirmed by dependability calculations performed on a functional model of the system [32, 33]. To ensure safe extraction of the beam, several conditions have to be met:

1. At least 14 of the 15 separate extraction kicker magnets must operate correctly.
2. The beam dump kicker must be synchronized with the $3 \mu\text{s}$ long beam abort gap.
3. The field of the extraction and dilution elements must track the beam energy.
4. The closed orbit errors in the dump insertion must be limited to about 4 mm since the aperture of the beam dump channel is tight.

Although nominally empty, the abort gap can contain particles due to de-bunching caused by RF noise, intra-beam scattering, etc [34]. The number of particles in the abort gap should be




Figure 10. Layout of the beam dumping system, for both LHC beams (courtesy M Gyr, private communication).




Figure 11. Layout of the beam dump block and shielding (courtesy L Bruno, private communication). The graphite dump block weighs about 10 tons, the concrete shielding about 1000 tons.

below the level at which quenches might be provoked. To avoid problems of this nature, after a failure in the RF system, the beam is dumped immediately. At 7 TeV, non-captured protons will lose energy by synchrotron radiation and therefore be captured in the momentum cleaning insertion. In addition, the particle density in the gap can be measured using an abort gap monitor and, if necessary, particles removed using the transverse feedback system [35].




Figure 12. Map of temperature increase in the central part of the beam dump block after full beam impact. The maximum temperature rise is about 750°C . The actual graphite block is cylindrical with a diameter of 70 cm.

8.4. Passive protection devices

A few well-defined failure modes exist which can cause single-turn failures, associated with injection or extraction kickers. For these failures, dedicated passive protection devices are designed to intercept the wrongly steered beam and to safely absorb or dilute the energy to a level which does not pose a threat to the downstream accelerator equipment.

8.4.1. Transfer line collimators. The beam is accelerated in the SPS to 450 GeV. To minimize risk of quenches in the LHC, it will be possible to shape the beams in the SPS by scraping the tails at $3.0\text{--}3.5\sigma$, where σ is the r.m.s. transverse beam size. At 450 GeV, the beam is then extracted and sent through two ≈ 2.8 km long transfer lines to the LHC. During this transfer process the beam, if mis-steered, could damage extraction elements, transfer line magnets and the LHC injection septum magnets [36]. An interlock system verifies a few milliseconds before extraction from the SPS the correct settings of all elements: orbit in SPS before extraction, strengths of kicker and septa magnets, magnet strengths in the transfer line, position of vacuum valves, collimators etc [29].

Despite this comprehensive equipment and beam surveillance in the SPS and in the transfer lines, failure modes exist, especially for the SPS extraction kickers, which could result in uncontrolled beam loss during the transfer process. A set of six collimators per transfer line [37], three per plane, will be used to limit the excursions of mis-steered particles into the LHC. The jaws of these collimators must be set to about 4.5σ , to ensure that the LHC aperture is protected [36]. A seventh collimator will be used to intercept particles with large momentum offsets.

8.4.2. Injection absorbers. The main LHC injection elements comprise injection septa and injection kickers. In case of a kicker failure, the whole injected beam could be deflected into the LHC with the wrong angle, and would impact some distance downstream on the aperture. To prevent this, three families of dedicated passive beam absorbers [38] are located downstream of the injection kicker magnet and will be set to about 7σ during injection, to avoid damage of elements in the injection region and the LHC where the aperture at injection is about 7.5σ [36]. These elements also intercept any circulating beam which is accidentally deflected by the injection kickers, for example, due to a synchronization error.

8.4.3. Beam dump absorbers. The beam dump extraction kickers are designed to deflect the whole of the LHC beam, and should be synchronized with the particle-free abort gap. To protect against failures in the synchronization, and against spontaneous kicker firing, the beam dump insertion has been designed [39] with dedicated devices to intercept particles mis-steered by the extraction kickers. A fixed absorber protects the extraction septum magnets, with a graded composition of graphite, high- and low-density carbon composite and titanium, for optimum absorbing power and robustness. A 6 m long single-jawed mobile graphite absorber is installed further downstream, to intercept any particles with amplitudes above about 8σ . This large object is supplemented by a two-jawed carbon composite collimator, which allows the beam position to be better constrained. These devices must be positioned with respect to the beam axis with an accuracy of better than $\pm 0.5\sigma$, or ± 0.3 mm at 7 TeV. The objects must safely absorb several tens of undiluted 7 TeV bunches, and the subsequent severe thermal loading results in difficult engineering challenges, in particular regarding the dynamic mechanical stresses (figure 13). The conceptual and mechanical designs have been the subject of extensive energy deposition and finite element dynamic stress simulations, requiring several iterations in the element designs [40].

8.5. Ring collimation system

Collimators will protect the machine against quenches and damage. The role of the collimation system is to capture a large fraction of the particles with large amplitudes that could otherwise be lost around the accelerator [6]. In particular beam losses into the superconducting magnets must be avoided.

Because of the very high stored energy in the beams, the LHC will be the first accelerator requiring collimators to define the mechanical aperture throughout the entire magnetic cycle. For efficient beam cleaning, a large number of collimators (about 43 per beam) are located at specific phase advance locations and are adjusted to define an opening corresponding to $5-9\sigma$, with σ the r.m.s. beam size at the collimator, as illustrated in figure 14 [41].

At 7 TeV this corresponds to an opening of only 2–4 mm. For operation at 7 TeV with nominal beam parameters and 10 h beam lifetime, more than 99.9% of the protons in the beam halo must be captured in the cleaning insertions to minimize the number of protons impacting on the superconducting magnets. For reduced lifetime, the efficiency must even be better. For the first year(s) of operation, the parameters are relaxed due to operation with beams below nominal intensity.

In one of the two insertions for beam cleaning, collimators in locations with nonzero dispersion capture protons with momentum deviations that are too large. In the other insertion, collimators capture protons with large betatron amplitudes.




Figure 13. Maximum stress ratio in 25 cm long carbon-composite beam dump absorber block, in the interval 0–200 μ s after beam impact. Regions of high localized stresses are visible in the block mid-plane, corresponding to the beam impact.




Figure 14. Typical setting of the collimators for nominal luminosity at 7 TeV. The aperture of the vacuum chamber is 56 mm (illustration drawing with the superposition of two collimators, one for the horizontal and one for the vertical plane).

The LHC collimator jaws are blocks of solid materials. Primary collimators intercept the beam halo with an impact parameter of the order of 1 μ m (figure 15). They scatter the protons into the so-called secondary halo that is intercepted by secondary collimators downstream, with an impact parameter of about 0.2 mm. Most of the protons are absorbed, but there is some




Figure 15. Impact of beam halo on collimator for two-stage cleaning.

leakage to a tertiary halo that extends to large amplitudes. Particles in the tertiary halo are lost into accelerator components, or, if the local losses risk to provoke quenches, into adequately positioned absorbers and tertiary collimators.

In case of equipment failures, collimators will be the first devices to intercept the beam and must absorb part of the energy until the beams are extracted. As discussed previously, if the beam lifetime drops exceptionally down to 0.2 h, the power deposition in the cleaning section could reach 0.5 MW at 7 TeV [42].

Studies of possible failure scenarios showed that up to about five bunches at 7 TeV can impact on a collimator jaw in case of beam dump pre-firing, or up to 288 bunches for a failure at injection. The simulated increase of the temperature shown in figure 16 indicates that materials such as aluminium or copper cannot be used. Only very light material with low number of protons in the nucleus (low-Z material) would not be damaged. This is the motivation to use graphite-based materials (fibre-reinforced carbon CFC) for the collimator jaws during the initial years of LHC operation [43], although a drawback of CFC is its higher electrical resistivity that can lead to high impedance and to beam instability at high bunch currents [44]. The robustness of a prototype collimator has been tested with beam impact at the SPS [45].

The collimators for cleaning and protection will be installed in several phases. For the start-up of the LHC, all primary and secondary collimators will have CFC jaws. In a second phase, it is planned to install additional advanced secondary collimators, to reduce the effect on the LHC impedance and to improve efficiency and operational handling.

9. Examples of failure scenarios

This section presents some example scenarios to illustrate possible failures, and the interplay of the various systems used for machine protection at the LHC.

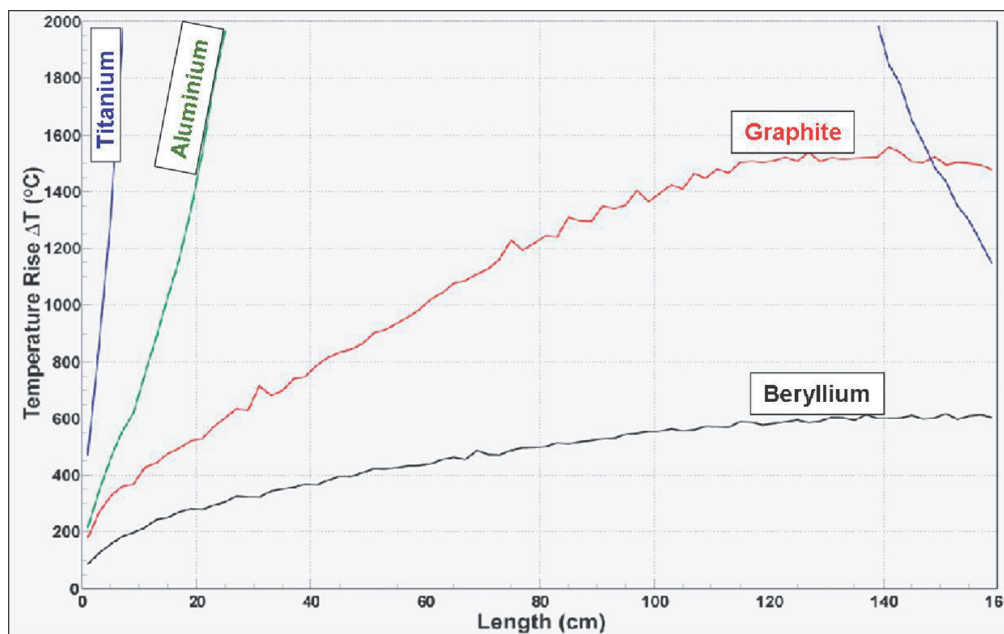


Figure 16. Temperature increase after about 20 bunches at 7 TeV hit the collimators versus length of the jaw (courtesy P Sievers *et al* 2003 LHC Performance Workshop, Chamonix XII and [6]).

9.1. Single-turn failure—*injection kicker flashover*

The injection kicker magnets are travelling-wave systems with an applied voltage of 25 kV, and during operation may experience a high-voltage breakdown or flashover inside the magnet structure, which can substantially alter the deflection imparted to the injected beam. Apart from the high-voltage aspects of the magnet design, the only protection against this eventuality is provided by the passive injection absorbers, which must intercept enough of the mis-steered beam to avoid damage to downstream components. The performance of the injection absorbers has been evaluated numerically, to derive the system settings and tolerances. Figure 17 shows the results of a particle tracking study where the required absorber settings of $\pm 7\sigma$ were determined.

9.2. Single-turn failure—*asynchronous beam dump*

A possible failure of the beam dump is the pre-firing of one beam dump kicker magnet, leading to a sweep of the beam across the LHC aperture (asynchronous beam dump). The other 14 kicker modules are immediately triggered after such a failure, but about 60 bunches will not be extracted correctly:

1. Twenty bunches receive small enough deflections to travel once around the machine, come back and are again deflected by the extraction kicker. For the nominal LHC settings these bunches will be extracted, but with different trajectories to nominal.
2. Four bunches reach the cleaning insertion and impinge on the collimators.




Figure 17. Fraction of injected particles transmitted into the LHC with amplitudes greater than $7.5\sigma_y$, as a function of the injection absorber system setting and kick error.

3. Seventeen bunches are intercepted by the 6 m long graphite absorber in the dump insertion some 300 m downstream from the kicker magnet.
4. Twenty bunches are intercepted by a fixed absorber in front of the extraction septum magnet.

Figure 18 shows the results of particle tracking and FLUKA energy deposition studies for the case of a pre-fired kicker magnet, where beam is swept across the absorbers and particle cascades result in energy deposited in the downstream superconducting magnet coil. The maximum energy density in this case is several orders of magnitude above the quench limit of 5 mJ cm^{-3} , but still below the assumed damage threshold of about 100 J cm^{-3} .

When such a failure occurs during extraction, bunches will oscillate with large amplitudes around the closed orbit. Hence, the closed orbit around the machine must be well controlled. At collision energy, the maximum 4 km beta functions in the low beta triplet at the high luminosity experiment insertions lead to a local reduction of the available aperture and hence increased probability for particle impact. Tertiary collimators will be used to shadow the superconducting triplet apertures against the tertiary halo and to provide local protection for irregular beam loss [6].

A failure in the synchronization with the beam abort gap has slightly less severe consequences than an erratic firing of the kicker, because the number of bunches swept at low amplitudes is smaller.

9.3. Multiturn failure—normal conducting D1 magnets

Quenches of superconducting magnets and other failures in the powering system are likely causes of beam instability. Combined failures may occur on certain occasions, for example after disturbances of the electrical network inside or outside CERN. Several failures were considered in [19]. A failure in the powering systems of the normal conducting separation dipole (D1) magnets that are used to deflect the beams close to the high luminosity insertions with physics experiments (ATLAS and CMS) is the most critical of all LHC powering failures. The deflection

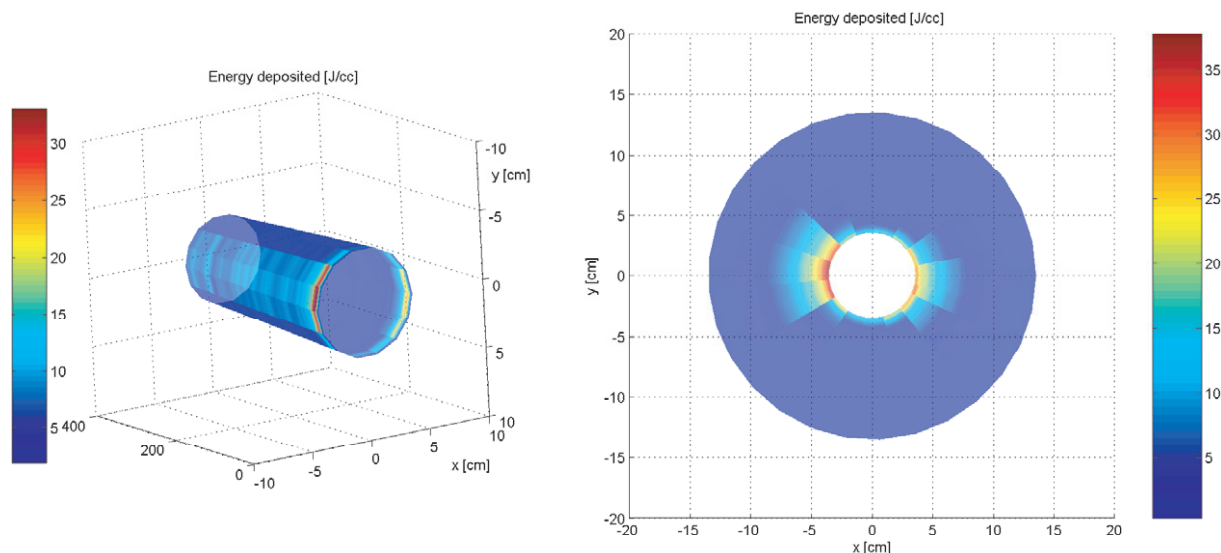


Figure 18. Energy deposited in a 400 cm long coil of the superconducting magnet downstream of the beam dump diluters, following an asynchronous beam dump. The left figure shows the deposited energy in J cm^{-3} as a function of length at $R = 3.7$ cm and the right figure shows the energy deposition across the cross-section at $Z = 12$ cm.

by the D1 magnet installed at a location with a large β -function of more than 4000 m leads to a fast change of the closed orbit. Since the collimators are very close to the beam, protons in the tails of the distribution would touch the collimator jaws already after several turns, exceeding more than 10^9 protons after about 15 turns [46]. The resulting losses will be detected by the BLMs installed around the collimators. The BLMs in the arcs will detect particles scattered out from the collimator jaws. Assuming that the collimators can withstand a beam loss of about 10^{12} protons, the jaws could be damaged already after 30 turns. The failure will also be detected by the fast magnet current change monitor that should request a beam dump before the beam touches the aperture. In addition, the fast beam position change monitors should detect the orbit movement.

10. Operation, reliability and availability of the protection system

The protection of the LHC relies on a large number of components, including about 4000 quench detectors, about 3800 BLMs, about 100 collimators and a variety of other systems, including the beam interlock and beam dumping systems. Common principles are used for the design of the protection systems:

1. A single equipment failure should not compromise the functioning of the protection systems and lead to beam-induced equipment damage.
2. Systems should be 'fail safe': a failure in the protection systems leads to a beam dump and downtime of the accelerator, but no equipment must be damaged.
3. Erroneous manipulation of the protection system that could compromise the accelerator safety must not be possible.

Table 5. Results for unsafety and false dumps for the baseline LHC scenario. Note that the unsafety numbers do not add linearly, as they are combined according to the functional model [48].

System	Unsafety (y^{-1})	False dumps (y^{-1})
Beam dump	$1.8 \times 10^{-7}(2\times)$	$3.4 \pm 1.8(2\times)$
Beam interlock	1.4×10^{-8}	0.5 ± 0.4
BLMs	1.4×10^{-3}	17 ± 4
Powering interlock	0.5×10^{-3}	1.5 ± 1.2
Quench protection	4.0×10^{-4}	16 ± 4
Overall protection system	$2.3 \times 10^{-4}(SIL - 3)$	41.6 ± 6.2

It is planned to introduce a pre-defined flexibility into the system, with some interlocks being configurable. To limit the risks when masking interlocks, a ‘Safe Beam Flag’ allows masking only when operating with ‘safe beam’ (beam below damage threshold). The information on disabled interlock channels will be maintained in a centralized database. This will facilitate bootstrapping of the LHC and allow optimizing the protection systems and interlocks during operation.

In addition, many of the systems which provide interlock signals need a degree of configurability, which must be carefully controlled. For the most critical systems like the interlock system and the beam dumping system, the reconfiguration of the system can only be made at a hardware level. For other devices a software configuration will be possible, but only by authorized users and making use of digital encryption of the data and electronic signature to ensure that the reconfiguration data has not been corrupted.

In safety systems, redundancy is widely used. For the LHC protection, there is redundancy within some of the systems, and redundancy across systems. For redundancy within a system, several channels have the same functionality. An example is the protection from beam losses: in general particles will be lost at different locations around the LHC, and several loss monitors will detect the loss. In addition, for many failures, equipment monitoring should also detect certain classes of failure and request a beam dump, while in a redundant channel the failed equipment produces a beam position drift which will also be detected by the fast position monitoring system or the beam loss monitoring system.

Although it increases the safety, redundancy leads to an increased complexity and possibly to a reduction of the availability: having two channels in parallel might increase the number of false beam dumps. A voting scheme is implemented in some systems to optimize the system availability [47]. This strategy is however not always applicable (for example, it was not conceivable to have more than one beam dumping system per beam).

The variety of protection systems has led to a coherent quantification of risks across systems using industrial standards [28]. This has enabled a comparison of the safety and availability contributions of the different systems to be compared using the same criteria, with modelling of the functional architecture of the individual systems and their interdependencies, and detailed Failure Mode Effect and Criticality Analysis with reliability prediction at the component level. Results for some key protection systems and the overall system are shown in table 5, for conservative assumptions regarding the operational scenario, the maintenance policy and the

detection rate of redundant branches [48]. Sensitivity analyses were also made, including the effects of different operational scenarios and different component failure rates, allowing the reliability of the systems to be optimized by design changes in some areas.

To accurately understand what caused a failure and if the protection systems functioned correctly, complete and coherent diagnostics data are required [49]. Most systems include transient recorders that are triggered in case of a beam dump. A timing system ensures that the recording of the data is synchronized. Analysing data from the various systems has several objectives:

1. Reconstruction of the event sequence that leads to a beam or power abort.
2. Demonstrate that the protection systems worked correctly, and that there is no loss of redundancy.
3. Fault diagnostics must help improve the operational efficiency of the LHC.

11. Conclusions

The energy stored in the LHC beams is two orders of magnitude larger than for other colliders, and operating with a stored beam energy of some 100 MJ in the presence of magnets that would quench if some millijoules of energy are released is one of the major challenges. The LHC is also the first accelerator with beams that potentially could damage a large fraction of the machine equipment. Safe operation with high intensity beams should always have priority and relies on the correct functioning of several complex and mutually dependent protection systems.

Protection of the accelerator starts already at extraction from the SPS where fast beam and hardware monitoring ensures extraction with correct parameters. Procedures will be enforced to ensure that the injection process is carried out safely, including the requirement that a low intensity beam be circulating in the LHC before a high intensity injection can be allowed. Collimators and beam absorbers must define the aperture in the transfer line and in the LHC during the entire magnetic cycle. For circulating beam in the LHC, equipment and beam monitoring will generate a beam dump request in case of failures. Novel ideas for detecting failures on the timescale of less than a millisecond are pursued, by hardware monitoring as well as by beam monitoring. Beam dump requests are transmitted with high reliability by the distributed interlock system to the beam dumping system, which must extract the beam from the accelerator and safely dispose of it within a few turns. Passive absorbers have been designed to intercept mis-steered beams in the event of single-turn failures.

Redundant systems are required to cover the spectrum of possible failures and to be able to react accordingly and in time. This results in a highly complex machine protection system, and availability and maintainability of the machine due to the complex protection are issues deserving much attention. For the first time in the development of high energy accelerators, comprehensive studies of reliability and availability for the main protection systems have been performed and the results applied to the design of the equipment.

Commissioning of the machine protection systems to ensure correct and reliable functioning at the start of LHC operation will be very challenging. The LHC Machine Advisory Committee, chaired by Professor M Tigner, June 2005, concluded for the machine protection systems: ‘we

recognize that the planned schedule is very aggressive, given the complexity and potential for damage involved in the initial phases of operation. It will be important to understand the performance of the machine protection system, the collimation system and the orbit feedback system as well as cycle repeatability and adequate beta-beat control before proceeding to run with significant stored beam energy. Pressure to take shortcuts must be resisted.'

Acknowledgments

The paper has been written on behalf of all CERN staff participating to the protection systems of the LHC, as well as all those involved from external collaborations. In particular, we thank all members of the Machine Protection Working Group, where most of the ideas have been presented and discussed.

References

- [1] Sibley C 2003 *Machine Protection Strategies for High Power Accelerators*, *Particle Accelerator Conf. 2003 (Portland, OR, 12–16 May 2003)*
- [2] Schmelzbach P 2006 Experience with high-power operation of the PSI proton accelerator facility *ICFA-HB 2006 Workshop (Tsukuba, May 2006)*
- [3] Holtkamp N 2006 *Commissioning Highlights of the Spallation Neutron Source*, *10th European Particle Accelerator Conf. EPAC 2006 (Edinburgh, 26–30 June 2006)*
- [4] Findlay D 2006 ISIS upgrade *ICFA-HB 2006 Workshop (Tsukuba, May 2006)*
- [5] The LHC Study Group 1995 *The Large Hadron Collider: Conceptual Design*, CERN/AC/95-05 (LHC)
- [6] Assmann R *et al* 2006 *The Final Collimation System for the LHC*, *10th European Particle Accelerator Conf. EPAC 2006 (Edinburgh, 26–30 June 2006)*
- [7] Benedikt M *et al* 2001 Performance of the LHC injectors *18th Int. Conf. High Energy Accelerators, HEACC 2001 (Tsukuba)*
- [8] Wikipedia keyword: fuel efficiency: http://en.wikipedia.org/wiki/Fuel_efficiency
- [9] Werner M and Wittenburg K 2006 *Very Fast Beam Losses At HERA, And What Has Been Done About It*, *ICFA-HB 2006 Workshop (Tsukuba, May 2006)*
- [10] Mokhov N 2006 *Very Fast Beam Losses At the TEVATRON, And What Has Been Done About It*, *ICFA-HB 2006 Workshop (Tsukuba, May 2006)*
- [11] Goddard B 2006 *Very Fast Beam Losses At the SPS, And What Has Been Done About It*, *ICFA-HB 2006 Workshop (Tsukuba, May 2006)*
- [12] Bottura L *et al* 1998 LHC main dipoles proposed baseline current ramping *CERN LHC-Project-Report 172*, CERN, 23 March 1998
- [13] Dahlerup-Petersen K *et al* 1999 *The Protection System for the Superconducting Elements of the Large Hadron Collider at CERN*, *18th Particle Accelerator Conf. PAC 1999 (New York City, NY, 29 March–April 1999)*
- [14] Dahlerup-Petersen K *et al* 2001 Energy extraction in the CERN LHC: a project overview *13th Int. Pulsed Power Conf. PPS2001 (NV, USA)*
- [15] Jeanneret J B *et al* 1996 Quench levels and transient beam losses in LHC magnets *CERN LHC-Project Report 44*
- [16] Fasso A *et al* 2003 The physics models of FLUKA: status and recent development *CHEP 2003 (LA Jolla, CA)*
- [17] Kain V *et al* 2005 *Material Damage Test with 450 GeV LHC-type beam*, *Proc. Particle Accelerator Conf. PAC 2005 (Knoxville, TN, May 2005)*

- [18] Tahir N A *et al* 2005 Impact of 7-TeV/c Large Hadron Collider proton beam on a copper target *J. Appl. Phys.* **97** 83532-1-8 124
- [19] Schmidt R *et al* 2003 *Beam Loss Scenarios and Strategies for Machine Protection at the LHC, HALO 2003 (Montauk, Long Island, NY, 19–23 May 2003)*
- [20] Schmidt R and Wenninger J 2000 LHC Injection Scenarios, *CERN LHC-Project-Note 287* March
- [21] Werner M, Zerlauth M, Schmidt R, Kain V and Goddard B 2005 *A Fast Magnet Current Change Monitor for Machine Protection in HERA and the LHC, 10th ICALEPCS (Geneva, 10–14 Oct. 2005)*
- [22] Holzer E B *et al* 2006 *Beam Loss Monitoring System for the LHC, Nuclear Science Symp. Conf. Record, 2005 IEEE 23–29 October 2005* Vol 2, pp. 1052–56 and *CERN-AB-2006-009*
- [23] Zamantzas C 2006 The Real-Time Data Analysis and Decision System for Particle Flux Detection in the LHC Accelerator at CERN, Brunel University, *CERN-THESIS-2006-037*
- [24] Zamantzas C *et al* 2005 The LHC beam loss monitoring system's real-time data analysis card *7th European Workshop on Beam Diagnostics and Instrumentation for Particle Accelerators DIPAC'05 (Lyon)*
- [25] Guaglio G 2005 Reliability of the Beam Loss Monitors System for the Large Hadron Collider at CERN *PhD Thesis* CERN and Université Clermont Ferrand Blaise Pascal, *CERN-THESIS-2006-012*
- [26] Jones R *et al* 2003 *The LHC Orbit and Trajectory System, DIPAC 2003 (Mainz)*
- [27] Bordry F *et al* Machine Protection for the LHC: Architecture of the Beam and Powering Interlock Systems *CERN LHC-Project-Report 52*
- [28] IEC 61508 1998 Functional safety of electrical/electronic/programmable electronic safety-related systems, International Electrotechnical Commission
- [29] Todd B, Dinius A, Nouchi P, Puccio B and Schmidt R 2005 The architecture, design and realisation of the LHC beam interlock system *ICALEPCS 2005 (Geneva, 10–14 Oct. 2005) Europhys. Conf. Abstracts* **29J**
- [30] Bruno L and Peraire S 1999 Design studies of the LHC beam dump, *CERN LHC-Project-Note 196*
- [31] Zazula J and Peraire S 1997 Design studies of the LHC beam dump *CERN LHC-Project-Report 112*
- [32] Filippini R *et al* 2004 *Reliability Issues of the LHC Beam Dumping System, 9th European Particle Accelerator Conf. EPAC 2004 (Lucerne, 5–9 July 2004)*
- [33] Filippini R *et al* 2006 *Update and Summary of the Dependability Assessment of the LHC Beam Dumping System, 10th European Particle Accelerator Conf. EPAC 2006 (Edinburgh, 26–30 June 2006)*
- [34] Shaposhnikova E, Fartoukh S and Jeanneret B 2004 LHC abort gap filling by proton beam *9th European Particle Accelerator Conf. EPAC 2004 (Lucerne, 5–9 July 2004)*
- [35] Hofle W 2004 *Experience in the SPS with the Future LHC abort Gap Cleaning, 9th European Particle Accelerator Conf. EPAC 2004 (Lucerne, 5–9 July 2004)*
- [36] Kain V 2005 Machine Protection and Beam Quality during the LHC Injection Process, *CERN-THESIS-2005-047*, Vienna University
- [37] Burkhardt H *et al* 2005 *Collimation in the Transfer Lines to the LHC, Particle Accelerator Conf. PAC 2005 (Knoxville, TN, 16–20 May 2005)*
- [38] Kain V *et al* 2005 *Protection Level during Extraction, Transfer and Injection into the LHC, Particle Accelerator Conf. PAC 2005 (Knoxville, TN, 16–20 May 2005)*
- [39] Goddard B *et al* 2006 *Protection of the LHC against Unsynchronised Beam Aborts, 10th European Particle Accelerator Conf. EPAC 2006 (Edinburgh, 26–30 June 2006)*
- [40] Goddard B *et al* 2006 *Dynamic Stresses in the LHC TCDS Diluter from 7 TeV Beam Loading, 10th European Particle Accelerator Conf. EPAC 2006 (Edinburgh, 26–30 June 2006)*
- [41] Assmann R A *et al* 2004 The New Layout of the LHC Cleaning Insertions *CERN LHC-Project-Report 774*
- [42] Assmann R A *et al* 2002 Requirements for the LHC collimation system *CERN LHC-Project-Report 599*
- [43] Bertarelli A *et al* 2005 *Mechanical Design for Robustness of the LHC Collimators, Particle Accelerator Conf. PAC 2005 (Knoxville, TN, 16–20 May 2005)*
- [44] Burkhardt H *et al* 2005 *Measurements of the LHC Collimator Impedance with Beam in the SPS, Particle Accelerator Conf. PAC 2005 (Knoxville, TN, 16–20 May 2005)*

- [45] Abele O *et al* 2005 *LHC Collimation: Design and Results from Prototyping and Beam Tests*, *Particle Accelerator Conf. PAC 2005 (Knoxville, TN, 16–20 May 2005)* and *LHC-Project-Report 850*
- [46] Kain V 2003 Power converter failure of the normal conducting D1 magnet at experiment insertions IR1 and IR5, *CERN LHC-Project-Note 322*
- [47] Vergara-Fernandez A 2003 Reliability of the Quench Protection System for the LHC Superconducting Elements, *CERN-THESIS-2004-019*
- [48] Filippini R *et al* 2005 *Reliability Assessment of the LHC Machine Protection System*, *Particle Accelerator Conf. PAC 2005 (Knoxville, TN, 16–20 May 2005)*
- [49] Ciapala E 2002 The LHC Post-mortem System, *CERN LHC-Project-Note 303*