LHC Main Dipole Magnet Circuits: Sustaining Near-Nominal Beam Energies

Scott M. Rowan, MEng

Submitted in fulfilment of the requirements for the Degree of Doctor of Philosophy

November 2016

School of Engineering
College of Science and Engineering
University of Glasgow

Declaration of Authorship

I, Scott M. Rowan, hereby declare that, unless explicitly stated via reference or citation, the following thesis is composed entirely of my own work and that this thesis has not been submitted elsewhere in consideration for a higher degree. The copyright of this thesis therefore belongs to the author under the terms and conditions of the United Kingdom Copyright, Designs and Patents Act 1988 and its subsequent amendments to date.

Scott M. Rowan
November 2016

I would like to dedicate this thesis to my father, Dougal Rowan, who, living with advanced heart failure, remains as optimistic, loving and as charismatic as ever. A true inspiration.

Abstract

Crossing the Franco-Swiss border, the Large Hadron Collider (LHC), designed to collide 7 TeV proton beams, is the world's largest and most powerful particle accelerator – the operation of which was originally intended to commence in 2008. Unfortunately, due to an interconnect discontinuity in one of the main dipole circuit's 13 kA superconducting busbars, a catastrophic quench event occurred during initial magnet training, causing significant physical system damage. Furthermore, investigation into the cause found that such discontinuities were not only present in the circuit in question, but throughout the entire LHC. This prevented further magnet training and ultimately resulted in the maximum sustainable beam energy being limited to approximately half that of the design nominal, 3.5-4 TeV, for the first three years of operation (Run 1, 2009-2012) and a major consolidation campaign being scheduled for the first long shutdown (LS 1, 2012-2014).

Throughout Run 1, a series of studies attempted to predict the amount of post-installation training quenches still required to qualify each circuit to nominal-energy current levels. With predictions in excess of 80 quenches (each having a recovery time of 8-12+ hours) just to achieve 6.5 TeV and close to 1000 quenches for 7 TeV, it was decided that for Run 2, all systems be at least qualified for 6.5 TeV operation. However, even with all interconnect discontinuities scheduled to be repaired during LS 1, numerous other concerns regarding circuit stability arose. In particular, observations of an erratic behaviour of magnet bypass diodes and the degradation of other potentially weak busbar sections, as well as observations of seemingly random millisecond spikes in beam losses, known as unidentified falling object (UFO) events, which, if persist at 6.5 TeV, may eventually deposit sufficient energy to quench adjacent magnets.

In light of the above, the thesis hypothesis states that, even with the observed issues, the LHC main dipole circuits can safely support and sustain near-nominal proton beam energies of at least 6.5 TeV.

Research into minimising the risk of magnet training led to the development and implementation of a new qualification method, capable of providing conclusive evidence that all aspects of all circuits, other than the magnets and their internal joints, can safely withstand a quench event at near-nominal current levels, allowing for magnet training to be carried out both systematically and without risk. This method has become known as the Copper Stabiliser Continuity Measurement (CSCM). Results were a success, with all

circuits eventually being subject to a full current decay from 6.5 TeV equivalent current levels, with no measurable damage occurring.

Research into UFO events led to the development of a numerical model capable of simulating typical UFO events, reproducing entire Run 1 measured event data sets and extrapolating to 6.5 TeV, predicting the likelihood of UFO-induced magnet quenches. Results provided interesting insights into the involved phenomena as well as confirming the possibility of UFO-induced magnet quenches. The model was also capable of predicting that such events, if left unaccounted for, are likely to be commonplace or not, resulting in significant long-term issues for 6.5+ TeV operation.

Addressing the thesis hypothesis, the following written works detail the development and results of all CSCM qualification tests and subsequent magnet training as well as the development and simulation results of both 4 TeV and 6.5 TeV UFO event modelling. The thesis concludes, post-LS 1, with the LHC successfully sustaining 6.5 TeV proton beams, but with UFO events, as predicted, resulting in otherwise uninitiated magnet quenches and being at the forefront of system availability issues.

Preface

This thesis provides an account of the doctoral research carried out by the author during his time working on-site at the CERN as part of the Technology Department's Machine Protection and Electrical Integrity Group (TE-MPE), in particular, the Performance Evaluation Section (TE-MPE-PE). Research was carried out as part of the CERN Doctoral Student Programme under supervision of Hon.-Prof. Rüdiger Schmidt and Dr. Bernhard Auchmann of CERN as well as Prof. Andrew Knox and Prof. Valentine O'Shea of the University of Glasgow. The original scientific contributions presented in this thesis are as follows:

Research surrounding the development and implementation of the copper stabiliser continuity measurement (Chapter 4) – a new global and conclusive method of superconducting magnet string continuity qualification. Original research included: the development of an electro-thermal analytical model of main busbar interconnects; development and optimisation of the qualification procedure; risk assessment and feasibility study by means of a single circuit type test; LHC-wide full qualification and outcome analysis.

Research surrounding the development of a semi-analytical modelling of so called unidentified falling object (UFO) events (Chapter 5), that is, events in which dust or debris particles interact with the beam and the resultant beam losses may quench an adjacent dipole magnet. Original research included: the development of a representative semi-analytical UFO event model; model predictions of the threat which high-energy UFO events pose on LHC availability; initial high-energy UFO event statistics and analysis.

Author Involvement

The following describes the thesis author's direct involvement in relation to the presented thesis works:

Development of the Copper Stabiliser Continuity Measurement:

The project (Chapter 4) as a whole was a large scale multi-departmental endeavour, including major contributions from cryogenics, powering and LHC operation specialists. As part of the core working group, personal contributions included assisting in the development and continued reform of the test procedure, numerical modelling of circuit performance for the purposes of calculating accurate protection system thresholds, and test outcome analysis. Being a designated expert, the author was posted in the control room throughout both the type test and LHC-wide campaigns. As an expert, the author was responsible for the on-going performance evaluation of each circuit, for the debugging and operation of protection systems and in particular, online calculation and setting of protection system thresholds.

UFO Event Modelling & Analysis:

This study (Chapter 5) involved investigating and taking into account the origins of previous works, importantly questioning the validity of the fundamental equations used therein. As the sole researcher apart from the listed supervisors, the author was responsible for the literature review, development and construction of the model itself as well as the simulation and reproduction of 4 TeV UFO events and the extrapolation and analysis of such events at 6.5 TeV beam energies.

Acknowledgements

First and foremost, it gives me great pleasure to extend my deepest gratitude to my colleagues at CERN, in particular, my mentors Rüdiger Schmidt, Arjan Verweij, and Bernhard Auchmann, to whom I am almost certainly indebted for life. To Rüdiger for providing the opportunity to carry out the entirety of my doctoral studies on-site at CERN, for his impressionable enthusiasm and encouragement and for opening my eyes to the world of accelerator sciences and large scale engineering. To Arjan for greatly improving my professional capabilities, having me work alongside long-standing members of the LHC machine protection panel, eventually becoming a designated electrical qualifications expert and regularly working from the LHC control centre. And to Bernhard for essentially being the ideal supervisor, i.e. knowledgeable, enthusiastic and always willing to help and or give tuition. Bernhard also saw me through the most troubling times of my doctoral studies and was unwavering in support, both professionally and, arguably more important, personally. The latter half of this thesis would not have come to fruition without him.

Furthermore, I would like to thank my University of Glasgow supervisors, Prof. Andrew Knox and Prof. Valentine O'Shea, for their continual guidance and support, for helping me secure my Graduate School Studentship and for their crucial feedback on my thesis topics and written works. I would also like to acknowledge Dr. Tony Kelly and Mrs. Elaine M^{c}Namara for their guidance and assistance in year-to-year progression formalities.

More personally, I would like to thank my office co-workers Michał Maciejewski, Ondřej Pícha and Deepak Paudel for making my time at CERN as fun as it was rewarding.

Finally, I thank my family and friends for their continued love and support, in particular, my fiancé Swati Chopra, who, even with her own doctoral thesis to burden her, was always there for me.

Related Publications

Journal Articles

[p1] **S. Rowan**, B. Auchmann, A. Lechner, W. Riegler, H. Schindler, R. Schmidt, and F. Zimmermann, "LHC dust-to-beam interactions: modeling, parametrics and impact projections," *Phys. Rev. ST Accel. Beams*, (final redaction).

Conference Proceedings

[p2] G. P. Willering, C. Giloux, M. Bajko, M. Bednarek, L. Bottura, Z. Charifoulline, K. Dahlerup-Petersen, G. Dib, G. D'Angelo, A. Gharib, L. Grand-Clement, S. I. Bermudez, H. Prin, V. Roger, **S Rowan**, F. Savary, J.-P. Tock, and A. Verweij, "Performance of the cold powered diodes and diode leads in the main magnets of the LHC," in *IOP Conference Series: Materials Science and Engineering*, vol. 101, June 2015, p. 012 076. DOI: 10.1088/1757-899X/101/1/012076.

[p3] **S. Rowan**, B. Auchmann, K. Brodzinski, Z. Charifoulline, B. Panev, H. Pfeffer, F. Rodriguez-Mateos, I. Romera, R. Schmidt, A. Siemko, J. Steckert, H. Thiesen, A. Verweij, D. Willering G. Wollmann, and M. Zerlauth, "Qualification of the Bypass Continuity of the Main Dipole Magnet Circuits of the LHC," in *Proceedings of IPAC2015*, May 2015, pp. 3141–3144, ISBN: 978-3-95450-168-7. [Online]. Available: http://inspirehep.net/record/1418047.

[p4] **S. Rowan**, A. Apollonio, B. Auchmann, A. Lechner, O. Picha, W. Riegler, H. Schindler, R. Schmidt, and F. Zimmermann, "Interactions between Macroparticles and High-Energy Proton Beams," in *Proceedings of IPAC2015*, May 2015, pp. 2112–2115, ISBN: 978-3-95450-168-7. [Online]. Available: http://inspirehep.net/record/1417783.

[p5] **S. Rowan**, B. Auchmann, K. Brodzinski, Z. Charifoulline, R. Denz, H. Pfeffer, I. Romera, V. Roger, R. Schmidt, A. Siemko, J. Steckert, H. Thiesen, A. Verweij, D. Willering G. Wollmann, and M. Zerlauth, "CSCM: Experimental and Simulation Results," in *Proceedings of IPAC2014*, June 2014, pp. 3988–3990, ISBN: 978-3-95450-132-8. [Online]. Available: https://inspirehep.net/record/1314585.

Related Awards

6^{th} International Particle Accelerator Conference (IPAC) Student Poster Award.
Sponsor: Institute of Electrical and Electronics Engineers (IEEE)
Awarded: Richmond, VA, USA, May 2015.
Related Works: "Interactions between Macroparticles and High-Energy Proton Beams".
Published Acknowledgements:

- Summary article of works published in the American Physical Society: Division of Physics of Beams Summer 2015 Newsletter.
- Award acknowledgement published in the IEEE Nuclear & Plasma Sciences Society (NPSS) newsletter, NPSS News Issue 3: Sept 2015.

Contents

Nomenclature

α	Magneto-resistivity	$\Omega m T^{-1}$
β	Beta Function *or* Relativistic Beta	m *or* -
ϵ	Beam Emittance	m
Γ	Quench Integral	$10^6 A^2 s$
γ	Lorentz Factor *or* Gamma Function	- *or* m
κ	Heat Transfer Coefficient	W/K
$\omega(z)$	Fadeeva Function	-
ρ	Resistivity *or* Density	Ωm *or* kg/m^3
σ	Beam Dimensions *or* Standard Deviation	m *or* -
Σ_{iel}	Macroscopic Cross-Section of Inelastic Interaction	m^2
σ_{iel}	Microscopic Cross-Section of Inelastic Interaction	m^2
τ	Time constant *or* Line-Charge Density	s *or* e/m
φ	Coulomb Potential	V
A	Atomic Mass	-
a	Area	m^2
B	Magnetic Field Strength	T
C	Circumference *or* Volumetric Thermal Capacitance	m *or* J/m^3K
c	Speed of light in a Vacuum *or* Specific Heat Capacitance	m/s *or* J/kg.K
E	Energy *or* Electric Field Strength	J *or* V/m
e	Eletron Charge	C
F	Force *or* Magnet Transfer Function	N *or* -
f	Frequency	Hz
H	Enthalpy	J
J	Current Density	A/m^2
L	Inductance *or* Half-Cell Length	H *or* m
l	Length	m
M	Threshold Margin	%
m	Mass	kg
n	Electron Density	m^{-3}
N_{A}	Avogadro's Constant	mol^{-1}
N_{u}	Molar Mass Constant	g/mol

N_x	Number of ...	-
Q	Charge	e
R	Resistance *or* Macroparticle Radius	Ω *or* m
r_x	Radius of ...	m
s	Longitudinal Coordinate *or* Pentration Depth	m *or* m
T	Temperature *or* Energy *or* Energy Density	K *or* J *or* J/m^2
t	Time	s
U	Voltage	V
V	Volume	m^3
v	Velocity	m/s
x, y	Transverse Offsets	m
Z	Atomic Number	-
Al	Aluminium	
ALICE	A Large Ion Collider Experiment (Collider Experiment)	
ATLAS	A Toroidal LHC ApparatuS (Collider Experiment)	
BCS	Bardeen Cooper Schrieffer	
BLM	Beam Loss Monitor	
BS	QPS Busbar System	
CERN	European Organization for Nuclear Research	
CMS	Compact Muon Solenoid (Collider Experiment)	
CSCM	Copper Stabiliser Continuity Measurement	
Cu	Copper	
DS	QPS Diode System	
EE	Energy Extraction System	
eV	Electronvolts	
He	Helium	
He-I	Helium (liquid)	
He-II	Helium (superfluid)	
LHC	Large Hadron Collider	
LHCb	Large Hadron Collider beauty (Collider Experiment)	
LINAC	Linear Accelerator	
LS	Long Shutdown	
MB	Main Dipole Magnet	
mBS	modified QPS Busbar System	
MQ	Main Quadrupole Magnet	
Nb–Ti	Niobium-Titanium Alloy	
nQPS	new Quench Protection System	
PC	Power Converter	

PS	Proton Synchrotron *or* Protection System
PSB	Proton Synchrotron Booster
QLU	Quench Level Uncertainty
QPS	Quench Protection System
RF	Radio Frequency
RRR	Residual Resistivity Ratio
RS	Running Sum
Run	Stable operation period of the LHC
SPS	Super Proton Synchrotron
TM	Transverse Magnetic
UFO	Unidentified Falling Object

Chapter 1

Introduction

Today, physicists and engineers are pushing technological boundaries in both the scale and complexity of particle accelerators in an attempt to answer some of particle physics' most important questions. With the construction of the largest and most powerful particle accelerator, the Large Hadron Collider (LHC), scientists at the European Organization for Nuclear Research (CERN) are at the forefront of this endeavour. Unfortunately, significant performance limitations were discovered during initial operation and, as a result, the maximum beam energy was restricted to approximately half that of the design nominal for the first three years of operation (Run 1). The issue was shown to result from discontinuities in the superconducting interconnects, which can be found either between magnets or connecting long busbar sections. To rectify this, the LHC was scheduled to be shutdown for a two year consolidation campaign such that each interconnect, over 10,000 in number, could be systematically repaired. However, several other limitations were identified throughout Run 1 which still have the potential to restrict operation at near-nominal energy levels - such is the topic of this thesis.

To provide perspective to the following thesis works, this chapter will briefly touch on modern day particle physics investigations as well as introduce the fundamentals of particle acceleration and superconductivity, the phenomenon which allows for such high-energy accelerators to exist. The chapter will then continue to give an account of the LHC, briefly describing the project mission, all encompassed collider experiments and the particle accelerator itself.

1.1 Beyond the Standard Model

To explore the origins and characteristics of elementary particles, physicists the world over utilise particle accelerators to collide charged particles at near the speed of light. The resultant energy density levels are of such magnitude that it can lead to the production of particles normally only present in some of the universe's most grandiose events [1], such as those predicted by the Standard Model [2]. Furthermore, detection and analysis of the

characteristics of these particles can provide crucial empirical backing for yet unproven particle physics theories.

The Standard Model, developed between 1961 and 1967 [3]–[7], is the most empirically validated theory for subatomic particles and forces. The model provides a plausible explanation for three of the four fundamental forces of nature, incorporating the electromagnetic force as well as the strong and weak nuclear forces [2]. Furthermore, the model predicts the existence of 17 elementary particles, see Figure 1.1, which make up all of known matter [2]. Note that, as of July 2012, with the announcement of the discovery of the Higgs Boson at CERN [8], [9], high-energy physics has now provided empirical evidence for the existence of all 17 predicted particles.

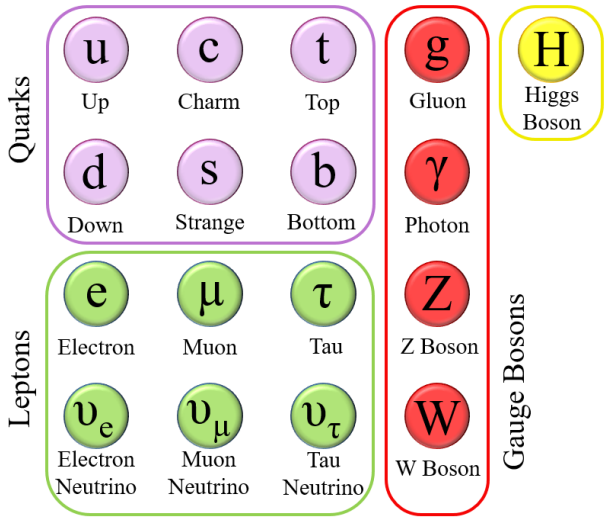

Figure 1.1: Image listing all 17 known elementary particles, as predicted by the Standard Model. The Standard Model particles include various forms (flavours) of quarks, leptons and gauge bosons (force carriers) as well as the recently discovered Higgs Boson. Note that as it is not the focus of this thesis, the properties and characteristics of each particle will not be discussed.

The Standard Model does not, however, incorporate an explanation for the full theory of gravitation [10], nor does it explain relatively recent discoveries that the universe is both more massive than previously thought as a result of Dark Matter [11], and that the expansion of the universe is in fact accelerating due to Dark Energy [12]. There are some potentially conclusive theories which give plausible explanations for each of the above phenomena, such as the existence of supersymmetric particles (more massive equivalents of all Standard Model particles) [13], however, all such theories require higher-energy particle collisions than ever before to be validated. It is for this reason that high-energy physics must continue to push the boundaries of the energy frontier, operating particle accelerators with the highest beam energies possible.

1.2 Particle Acceleration

Following Einstein's relativistic mass-energy equivalence equation,

$$E = \gamma m_0 c^2, \tag{1.1}$$

the particle energy, E, is proportional to the particle rest mass, m_0, and the speed of light in a vacuum, c, given the relativistic Lorentz factor, γ [14]. Following this, in an attempt to produce the most massive elementary particles predicted by particle physics theories, particle colliders have been pushing technological boundaries over the course of the last century to reach ever increasing collision energies by accelerating particles as close to the speed of light as possible, as the higher energy they obtain, the more likely it is for more massive or previously unknown particles to be produced [15].

Accelerating beams of charged particles requires an electric field. In practice, this is implemented by use of radio-frequency (RF) resonant cavities which act as an electromagnetic waveguide operating in a transverse-magnetic (TM) mode, *i.e.* no magnetic field in the direction of propagation [15]. Depending on the type of accelerator, energy is then transferred to the beam via the on-axis electric field at either a fixed frequency and an ever increasing distance (linear accelerator) or, as is the case for the LHC, at a fixed distance and an ever increasing frequency (circular accelerator) [15]. The force applied to the beams at a given point follows the Lorentz Force equation,

$$\vec{F} = e\big[\vec{E} + \vec{v} \times \vec{B}\big], \tag{1.2}$$

where $\vec{F}$ is resultant force, e is the particle charge, $\vec{v}$ is the velocity, $\vec{E}$ the electric field and $\vec{B}$ is magnetic flux density. For circular accelerators, the accelerating electric fields work in conjunction with ramped magnetic fields, produced by dipole electromagnets, which provide bending forces perpendicular to the beam axis, allowing the particles to circulate with a fixed orbit – this specific type of accelerator is known as a synchrotron [15]. Looking solely at the magnetic field component of the above equation one can show that the sustainable bunch momentum, p, when expressed in electronvolts (eV) divided by the speed of light, *i.e.* GeV/c, follows

$$p = \gamma m_0 v \approxeq 0.3Br, \tag{1.3}$$

where r is the accelerator radius. It follows that increasing the maximum sustainable beam energy requires either an increase in the size of the accelerator or, for a defined tunnel size,

as was the case with the LHC [16], an increase the bending forces via greater magnetic fields [15]. Note that the expression of particle momentum in eV/c is commonplace in accelerator sciences, with accelerators commonly having sustainable beam energies in the order of GeV or TeV.

1.3 Superconductivity & Quench Phenomenon

To achieve the magnetic fields capable of bending such high energy beams around the LHC, the main magnet circuits are almost entirely superconducting [16]. Superconductivity is a well-known phenomenon in which, at cryogenic temperatures, some highly resistive conductors may enter an altered state in which they have zero electrical resistance. According to the Bardeen Cooper Schrieffer (BCS) theory of superconductivity, this is due to low-energy electron-phonon interactions displacing a conductor's lattice such that following electrons are attracted to the resultant increase in positive charge density [17]. Furthermore, the resultant highly correlated motion between electrons (known as Cooper Pairs) cannot be broken by lattice vibrations due to being at such low temperatures, *i.e.* the lattice does not have sufficient energy to do so. This results in the electrons never actually colliding with the surrounding atoms whilst below a material dependent critical temperature, resulting in zero resistivity. Superconductivity therefore allows for low voltage-high current systems and when applied to electromagnets allows for the production of very high magnetic fields.

To maintain superconductivity, one must operate within the used superconductor's critical surface – temperature, current density and magnetic field thresholds [18]. Exceeding any of these thresholds will result in the transition of the superconductor to its normal resistive state. This process is commonly referred to as a quench. Figure 1.2 shows the critical surface of Niobium-Titanium, Nb–Ti – the superconducting alloy used throughout the LHC.

1.3.1 Typical Quench Scenarios

In a cryogenic electromagnetic system, such as the main systems of the LHC, there are numerous scenarios which may result in a quench. For example, a quench may arise if the circuit is powered beyond the maximum current or rate of change of current, resulting in excessive changes in the magnetic fields, inducing eddy current losses and resulting in excessive localised heating [19]. Another electrical scenario may be that a defect in the coil or busbar results in an area of non-zero electrical resistance. The resultant voltage drop and energy deposition across the resistive segment results in the formation of a hot-spot which can exceed the critical temperature. Insufficient cooling may also be the cause of

a quench, wherein the cryogenics systems fail to maintain the system below the critical temperature [20]. Such a failure may even propagate as the warmer coolant diffuses along the cryostat, quenching adjacent sections or systems. Finally, an example of a rarer quench scenario is when there are sufficient localised beam losses which deposit enough energy in the adjacent subsystems to create a hot-spot and initiate a quench [21]. Note that for higher operating currents, *i.e.* operating the system closer to the critical values for current and magnetic flux density, the system is inherently more susceptible to a quench, due to the reduced temperature margins.

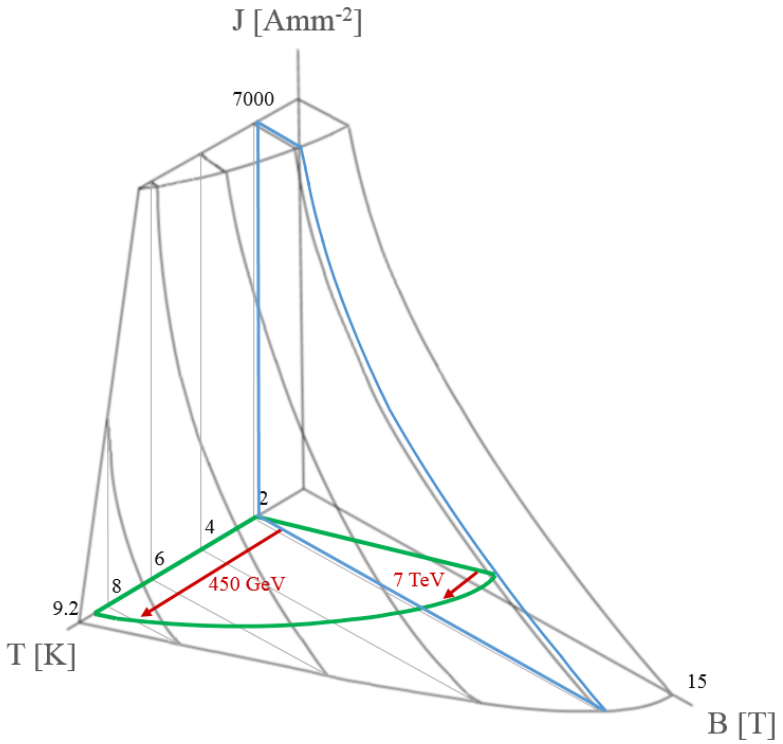

Figure 1.2: Plot of the superconducting critical surface of Nb–Ti. Also shown is the nominal operating conditions for the main dipole magnets (blue plane) and the corresponding temperature margins for various current densities (green plane). Note that as the operating energy increases, temperature margins decrease (indicated in red).

1.4 The Large Hadron Collider

The LHC is the largest circular accelerator in the world, spanning ∼26.66 km in circumference. It is designed to accelerate two proton beams to a nominal energy of 7 TeV, the velocity equivalent of ∼99.9999991% the speed of light. Beams are accelerated in opposite directions and are made of up to 2808 particle bunches with $\sim 1.15 \times 10^{11}$ protons per bunch. There are then four specific regions where the beams are made to cross each other for collisions – the detector systems that comprise the LHC collider experiments are

present at these locations [16].

The LHC is the latest and largest of the accelerators at CERN, however, previous accelerators are still used to gradually increase beam energy prior to injection into the LHC at an energy of 450 GeV – denoted the injection energy. Figure 1.3 presents a simplified layout of the complete accelerator complex. Initial acceleration is performed by a linear accelerator, LINAC2, followed by a series of increasingly larger circular accelerators – the proton synchrotron booster (PSB), the proton synchrotron (PS) and the super proton synchrotron (SPS) [16].

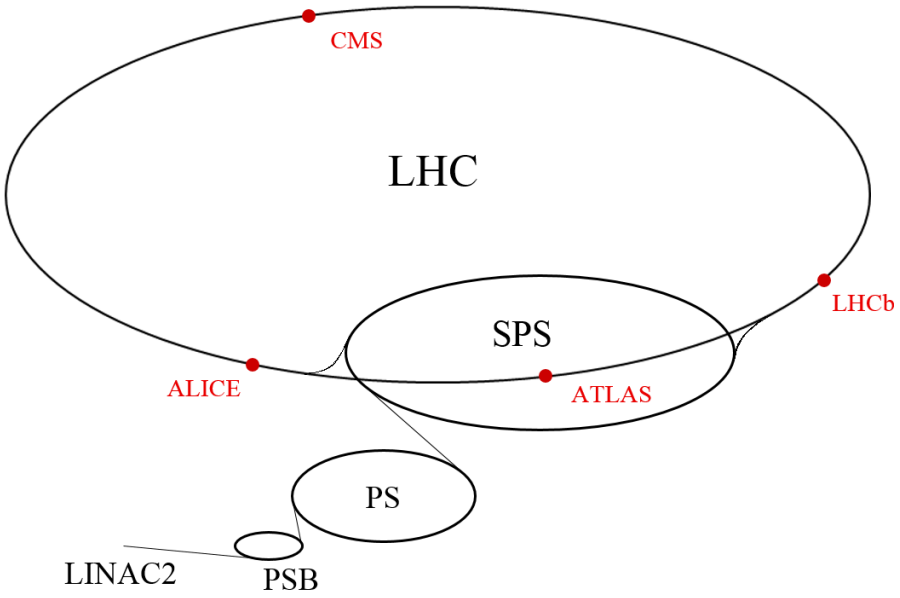

Figure 1.3: Illustration of the CERN accelerator complex. Protons are gradually accelerated prior to injection in the LHC, initially by a LINAC and subsequently by the PSB, PS and SPS to reach the injection energy of 450 GeV.

Focusing on the LHC, Figure 1.4 illustrates a simplified layout of the ring and main facilities. For engineering purposes, the LHC is divided into eight powering sectors, with the ends of each sector denoted as points. Naming conventions follow such that Sector 12 is found between points 1 and 2. In addition to the four collision regions, encompassing detectors ATLAS, ALICE, CMS and LHCb (discussed in Section 1.5) at points 1, 2, 5 and 8, respectively, the LHC also consists of beam cleaning regions (points 3 and 7) as well as an acceleration region (point 4) and a beam dump region (point 6) (not discussed in this thesis, for more information see [16], [22]). Beams are injected into the LHC at points 2 and 8 for beam 1 and 2, respectively, with beam 1 travelling in the clockwise direction and beam 2 anti-clockwise. Any reference to being left or right of a detector or point is from the perspective of standing in the centre of the circle.

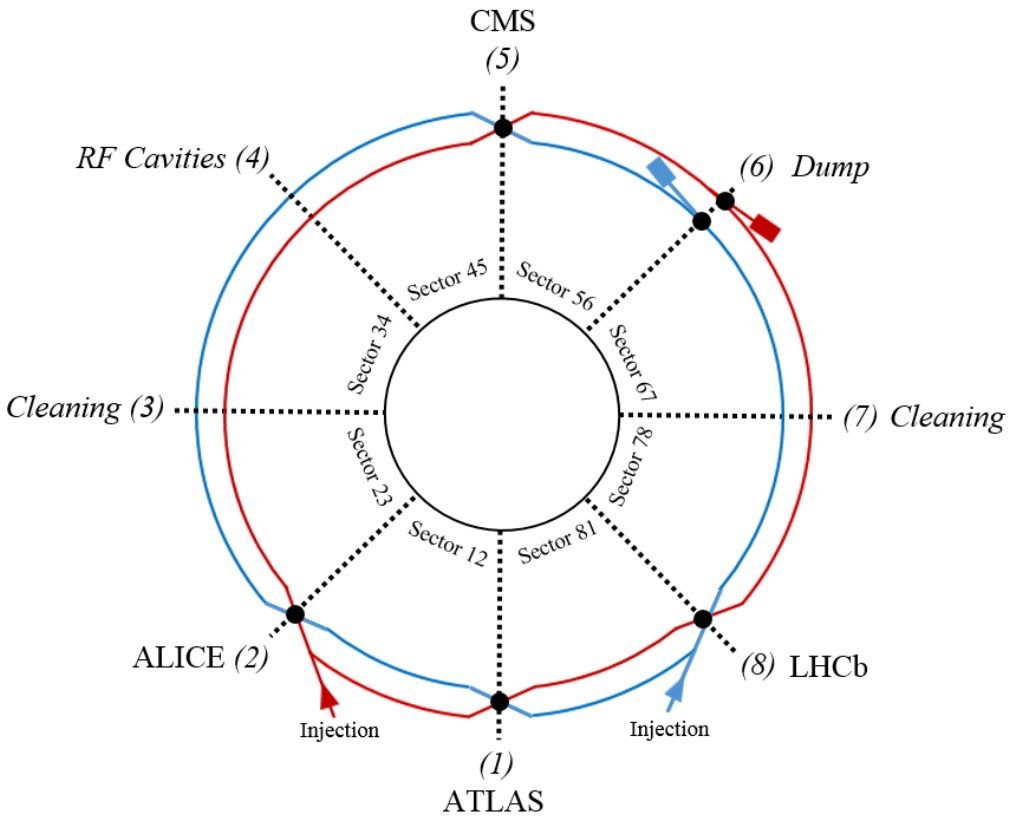

Figure 1.4: Illustration of the LHC layout, showing the various facilities at each point (numbered). ATLAS, ALICE, CMS and LHCb are the main detectors, located where the beams are made to cross each other for collisions. Beams are accelerated, cleaned and dumped at points 4, 3 and 7, and 6, respectively.

1.5 The LHC Collider Experiments

The LHC was designed and constructed to push the boundaries of the energy frontier of physics, reaching collision energies much higher than have ever been achieved. The justification for such a large scientific endeavour was that numerous fundamental physics theories, such as the Higgs Theory [5] and the existence of Dark Matter [11], had, to date, no empirical evidence. Furthermore, due to energy limitations, no evidence could be produced with existing systems, such as the 1 TeV Tevatron at Fermilab National Laboratory in North America [23]. The following section briefly describes the basic principles of particle collision analysis and discusses each of the four LHC collider experiments, in particular, why each would benefit from an increase in maximum sustainable beam energy. Note that detector specifics, *i.e.* design and detector technologies, are not the focus of this thesis and will therefore not be discussed. For technical information on each detector see [26], [27], [28] and [29] for ATLAS, ALICE, CMS and LHCb, respectively.

1.5.1 Particle Collision Outcome Analysis

Although each detector utilises different magnet and particle tracking technologies, the underlying principle of how each detector works remains the same. In a typical proton-proton collision event, if the resultant energy density is high enough, the quarks and gluons within the protons will interact to form various low-energy elementary particles [1]. On occasion, rarer heavier or more energetic particles can be produced along with their anti-matter equivalents, which can then decay into specific groups of elementary particles. All particles, having sufficient energy, then scatter in all possible directions and, if charged, can have their momentum and energy calculated. To determine the momentum of charged particles, an encompassing magnetic field, produced by a large superconducting magnet system, bends the trajectory of all charged particles, allowing their relative position to be tracked and their momenta can be calculated [26]. The energy of a particle is then calculated from the deposited energy and penetration depth measured by surrounding calorimeters and spectrometers in accordance with the energy-momentum equation

$$E^2 = (pc)^2 + (m_0c^2)^2. \tag{1.4}$$

A basic 2-D illustration of a particle tracking and detection system is shown in Figure 1.5. Ejected particles are detected by an encompassing layered array of particle specific detector structures (hadron calorimeters for protons and neutrons, electromagnetic calorimeters for electrons and photons and muon spectrometers for muons) which span several metres before and after the collision point [26].

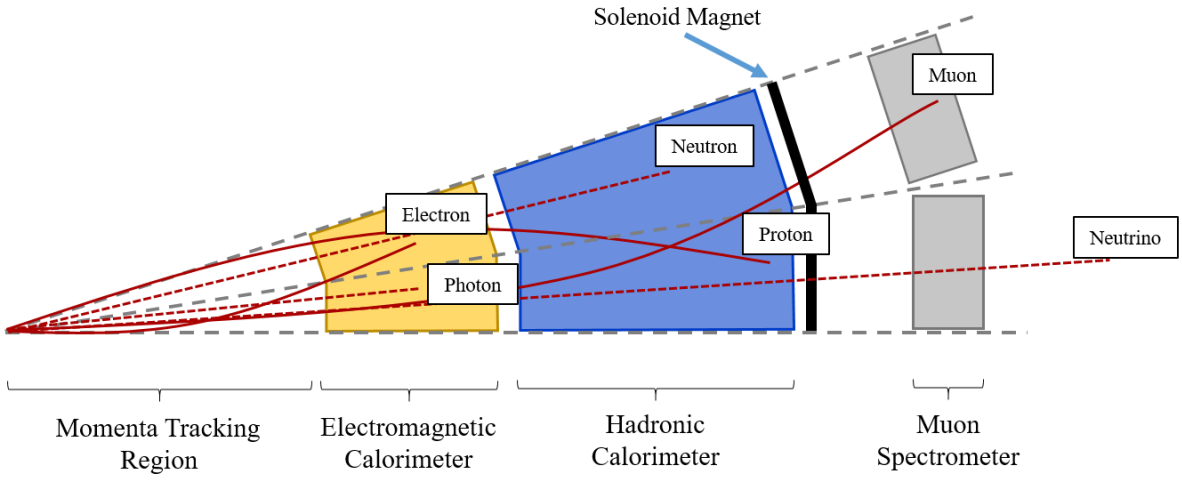

Figure 1.5: 2-D Illustration of scattered-particle detection. Shown are example particle trajectories as well as the various detector structures. Note that this is a cross-section arc transverse to the beam-axis and that the detector arrays span 360 degrees around the collision centre.

Particle Identification

With sufficient theoretical knowledge regarding plausible rare particle decays, it is possible to identify which rare particles were produced in a given collision event. Due to the law of conservation of energy, it follows that, having established the energy of the expected resultant elementary particles, it can be assumed that an undetected particle of equivalent mass or energy must have been present. For example, a Z-boson may decay into a fermion, such as an electron, and its antimatter equivalent, in this case, a positron ($Z^0 \rightarrow e^+e^-$) [24]. To establish the mass of the Z-boson, investigations can be limited to events which are found to have had an electron and a positron originating from the same location. The mass of the original particle can then be calculated from the sum of the decay particle masses, derived from their measured energies and momenta, respectively. Errors in the calculation of a rare particle's mass, such as unascertained energy losses in the beam pipe, are mitigated by the accumulation of sufficient event statistics. The accepted statistical significance for a particle to be formally recognised is 5 σ, *i.e.* 5 standard deviations above the background and a one in 3.5 million chance of being a statistical fluctuation [25].

1.5.2 ATLAS & CMS

The ATLAS (A Toroidal LHC ApparatuS) and CMS (Compact Muon Solenoid) experiments are large (25 m long, 45 m in diameter and 15 m long, 28.7 m in diameter, respectively) general-purpose particle detectors capable of investigating a wide range of physics. Both ATLAS and CMS experiments primarily focus on the detection and understanding of the properties and characteristics of the Higgs Boson [26], [28]. However, more specific to the requirement of an increase in maximum sustainable beam energy, both experiments are also involved in the search for super-symmetric particles, which could further understanding of the relatively unexplored phenomenon of Dark Matter. Note that with the LHC being the only accelerator capable of such high-energy collisions, having two technologically different and independently operated detectors both investigating the same physics is important as it allows for immediate comparison and validation of collected data. A 3-D graphical reproduction of a 7 TeV collision event (3.5 TeV individual beam energies), *circa* 2010, as measured by the ATLAS experiment is shown in Figure 1.6.

1.5.3 ALICE

The LHC is not only designed to accelerate and collide protons, but is also capable of heavy-ion collisions, lead-lead or lead-proton [27]. ALICE (A Large Ion Collider Experiment), 26 m long and 16 m in diameter, is a heavy-ion detector optimised to study the

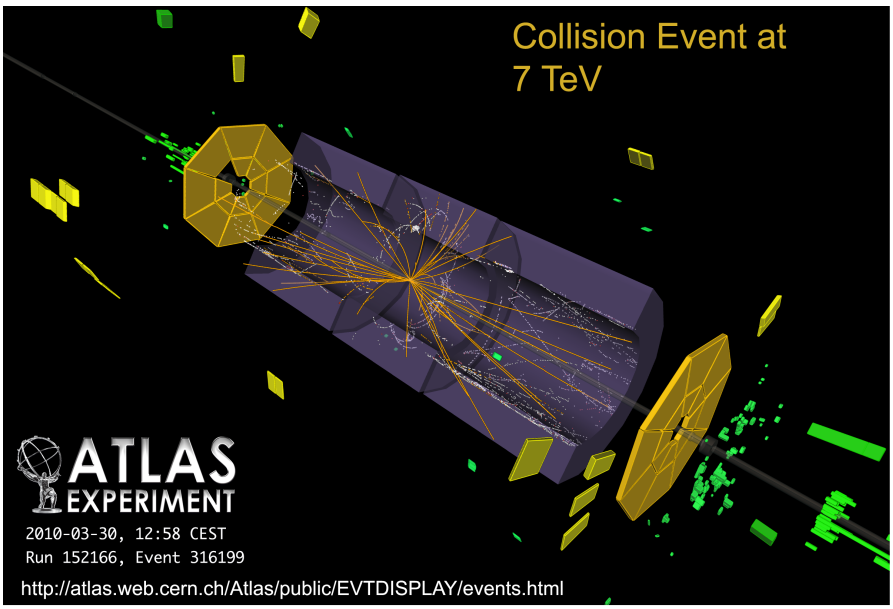

Figure 1.6: 3-D graphical reproduction of a 7 TeV collision event showing the ejected particle trajectories and relevant detector structures as measured by the ATLAS experiment, circa 2010. Image courtesy of CERN [30].

outcome of these collisions, in which energy densities and temperatures are high enough to produce a quark-gluon plasma [31], a state of matter similar to that which is only theorised to have existed in the first seconds after the Big Bang [32]. In both cases, either lead-lead or lead-proton collisions, the energy density levels are key to production of this state of matter and the experiment can therefore only benefit from an increase in maximum sustainable beam energy.

1.5.4 LHCb

The LHCb (Large Hadron Collider beauty) particle detector is designed to investigate observed differences between matter and antimatter [29]. In particular, investigations study so called charge-parity (CP) violations [33], in which the rate of decay of some elementary particles and their anti-matter equivalents has been shown to contain slight differences. The LHCb experiment focuses on the precise measurement of the properties and decay of the bottom (beauty) quark and its anti-matter equivalent, in the hope of furthering knowledge as to why nature seemingly prefers matter over anti-matter [33]. Furthermore, with the probability of a collision producing bottom quarks very much relating to the collision energy, once again, the LHCb experiment can only benefit from an increase in maximum sustainable beam energy.

1.6 The LHC Main Ring & Beam

Spanning the majority of the LHC, ∼24 km, are the Arcs. The LHC Arcs house, amongst others, the main superconducting dipole and quadrupole magnet circuits [16]. Technical information on the main dipole circuit can be found in Chapter 2. The main quadrupole circuit is not the focus of this thesis and will not be discussed in detail. For technical information on this circuit, see [16].

1.6.1 LHC Arc Cryostat

As discussed in Section 1.3, the LHC utilises superconducting magnet technologies. This requires cooling the magnets to cryogenic temperatures, sufficiently within critical boundaries. The LHC main dipole magnets have design-nominal fields of 8.3 T, which with an nominal operating current of 12.85 kA results in a critical temperature of ∼3-4 K. Superfluid He (He-II), having zero viscosity [34], is used as it can flow through micrometre gaps in the magnet coils (not the case with liquid He-I), providing additional cooling. As a result, each of the LHC Arc cryostats, which span the length of an Arc sector, ∼3 km in length, maintains a superfluid helium (He-II) bath at a temperature of 1.9 K (sufficiently below the maximum temperature for superfluidity (the Lamba point), 2.2 K) with a pressure of 1.3 bar [16]. A picture of the cryostat is shown in Figure 1.7.

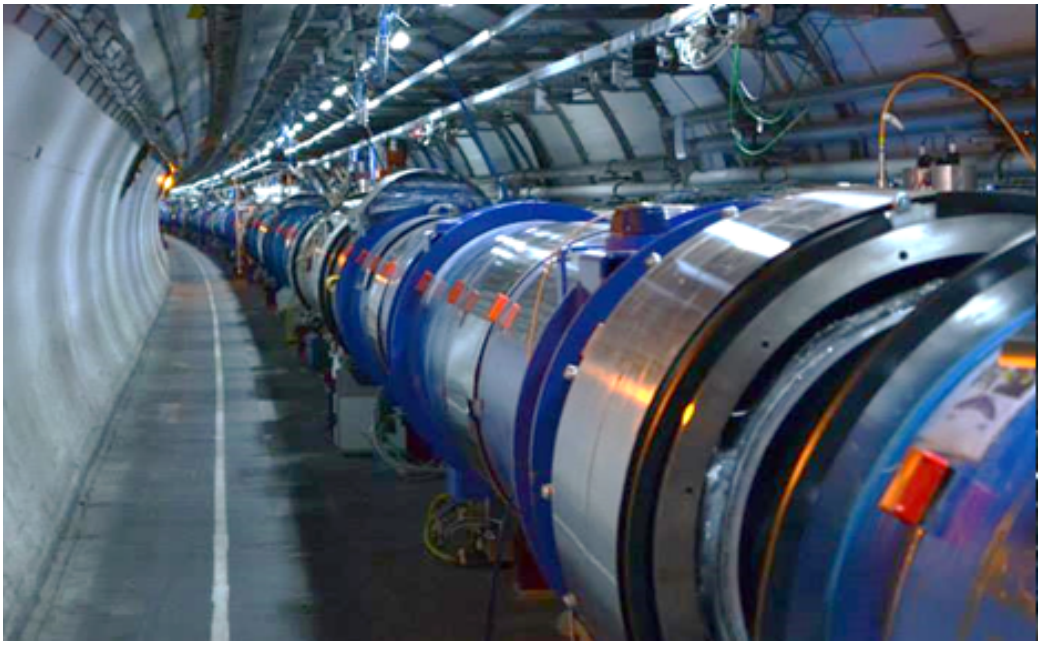

Figure 1.7: Picture of the LHC Arc cryostat which encases, amongst others, the main dipole circuits. Openings can be seen along the cryostat as this picture was taken during a repair campaign (discussed in Chapter 3).

1.6.2 Beam Pipe & Beam Screen

To maintain a vacuum in which the beam itself can travel, a steel beam pipe, commonly referred to as the cold bore, encircles the full circumference of the LHC [16]. To protect the beam pipe from light emissions that occur in the process of bending the beam, known as synchrotron radiation [15], a Cu-lined stainless steel beam screen is installed on the inside of the beam pipe, see Figure 1.8 [35].

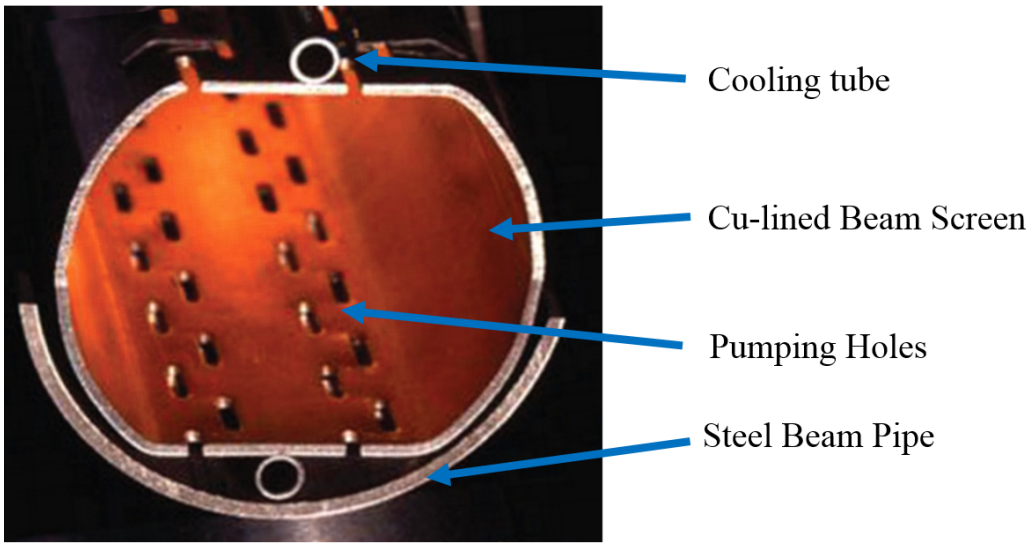

Figure 1.8: Picture of the cross-section and internals of the LHC beam pipe and beam screen. The beam itself is guided such that it is always centred with respect to the beam screen. Photograph courtesy of CERN [36].

The Cu lining provides a low resistance path for image currents, while the stainless steel frame provides sufficient structural integrity to withstand large electromagnetic forces in the event of a magnet quench. The beam screen is cooled by two cooling tubes which contain gaseous He maintained at a temperature between 4.6-20 K, depending on beam activity. Furthermore, pumping slots are present along the flat surfaces of the beam screen to allow any gas build-up due to photo-desorption to escape, minimising vacuum pressures.

The orientation of the beam screen also varies, with the smallest aperture being orientated either horizontally or vertically, see Figure 1.9. Orientation of the beam screen in the Arcs is typically vertical, with the exception of sectors prior to the two main detectors, ATLAS and CMS, in the direction of the beam, *e.g.* beam 1 for Sector 81 (left of ATLAS) and beam 2 pipe for Sector 12 (right of ATLAS) [16]. This is to increase the aperture size, allowing for a larger beam size and longer beta (discussed in Section 1.6.4) prior to focusing for collision.

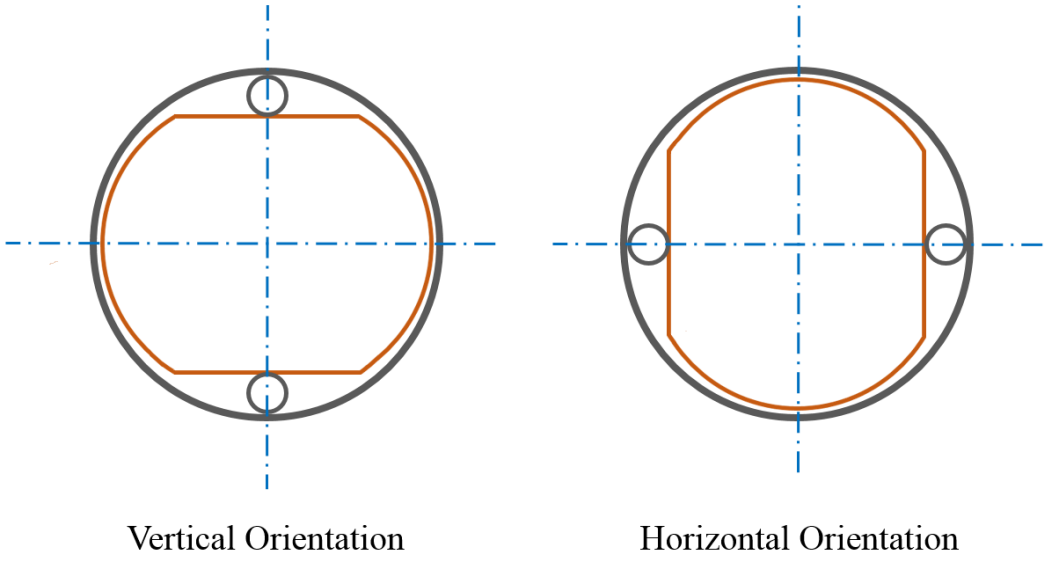

Figure 1.9: Orientation variants of the beam screen within the beam pipe.

1.6.3 Magnetic Lattice

Fundamental to having particle bunches guided around a circular accelerator is the magnetic lattice. An accelerator magnetic lattice is a fixed arrangement of dipoles and quadrupoles, known as the lattice cell, which, assuming no straight sections or collision regions, is repeated for the length of the accelerator [16]. The LHC Arc lattice cell is shown in Figure 1.10. As shown, a typical Arc cell encompasses 6 dipole (MB) and 2 quadrupole magnets, one focusing (MQF) and one defocusing (MQD), and is ∼110 m in length. Do note that not shown are higher-order corrector magnets which are installed alongside each main magnet to account for chromatic aberration, *i.e.* variations in focusing angles due to intra-bunch variations in proton energies, as well as to compensate for manufacturing flaws or alignment issues. Corrector magnet systems will not be discussed in this thesis, for information see [16].

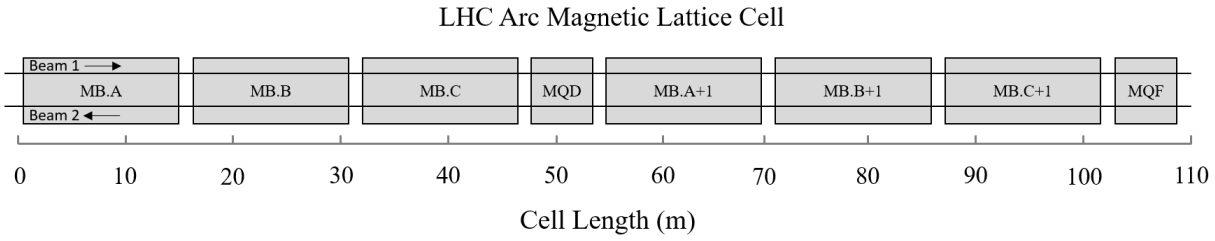

Figure 1.10: Illustration of a typical LHC magnetic lattice cell. At ∼110 m long, each cell consists of 6 dipole magnets with a quadrupole magnet following every group of 3 dipoles.

1.6.4 Beam Profile & Dynamics

This section discusses the main longitudinal and transverse characteristics of the LHC beam under normal operation, in particular, detailing beam bunching and the how the magnetic lattice affects the transverse beam profile.

Beam Bunching & Bunch Spacing

The nominal LHC beam comprises of up to 2808 particle bunches with $\sim 1.15 \times 10^{11}$ protons per bunch and a bunch spacing [37] (the time between each bunch) of 25 ns [16]. Particle bunches are the controllable outcome of using RF acceleration technologies [37]. The use of a continuous oscillating TM mode electromagnetic field allows the system to both marginally accelerate or decelerate particles that arrive relatively early or late with respect to a synchronous particle (the particle which is subject to zero field), such that they oscillate longitudinally in position within the particle bunch. This is a desired outcome as stable operation requires that all particles within a bunch have the same average velocity [37].

While ramping up in momentum, higher-energy particles will circulate the ring faster and are therefore decelerated by being subject to a negative field to compensate, *vice-versa* for lower-energy particles [37]. However, as the particle bunches near the speed of light, *i.e.* exceed the so called transition energy, the velocity remains fairly consistent between all particles within a bunch and the bending angle of individual particles dominates, that is, the trajectories of higher-energy particles are not bent to the same extent by the dipoles as those with lower-energy, and higher-energy particles therefore have a longer orbit circumference and a lower revolution frequency [37]. It follows that lower-energy particles within a bunch, with an inherently higher revolution frequency, will arrive earlier than the synchronous particle and therefore need to be accelerated to compensate. Similarly, above transition, higher-energy particles, with a lower revolution frequency, arrive late with respect to the synchronous particle and are therefore decelerated. Figure 1.11 illustrates this process for a particle bunch with an energy level above the transition energy ($v \approx c$).

Note that, throughout Run 1, the bunch spacing was set to 50 ns (1380 bunches) [38] to minimise the heat load in the beam pipe as a result of electron cloud build-up, which is significantly greater with 25 ns bunch spacing [38]. The electron cloud is a build-up of backscattered and secondary electrons, as a result of beam-induced multipactoring [39] which fills the beam pipe and surrounds the beam as a proton bunch passes by. Note that, following successful tests towards the end of Run 1, it is expected that the LHC will operate with 25 ns bunch spacing throughout Run 2 [40].

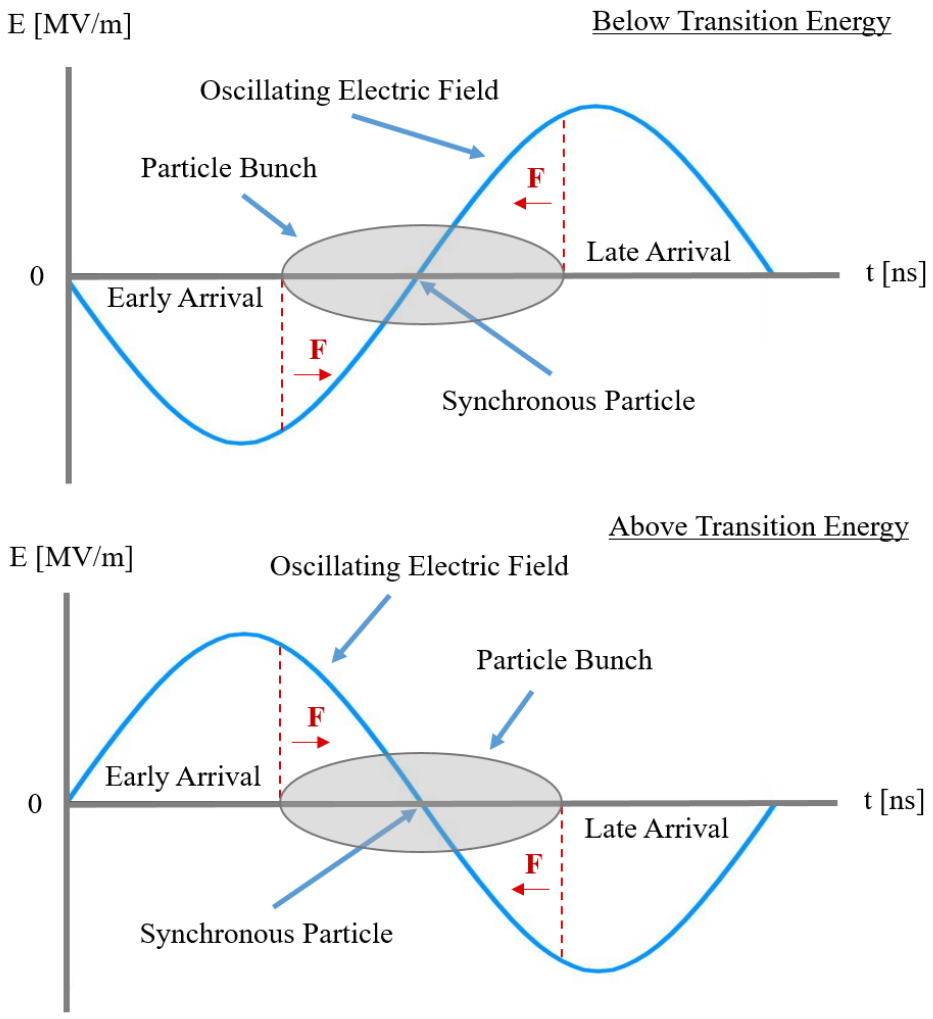

Figure 1.11: Illustration of how the oscillating TM mode electromagnetic field affects an individual particle bunch with an energy level below (top) and above (bottom) the transition energy, accelerating or decelerating particles that arrive early or late with respect to the synchronous particle.

Transverse Shape & Dynamics

The transverse particle spread is controlled predominately by quadrupoles and higher-order corrector magnets. Each quadrupole, however, can only focus the beam in one plane, *i.e.* either horizontally or vertically [37]. Consequently, every second quadrupole must either be focusing or defocusing (focusing in the other plane), as was shown for the magnetic lattice (Section 1.6.3). The result is that the transverse beam profile continuously varies from predominately vertical to predominately horizontal, while remaining centred around the beam axis. Figure 1.12 depicts how the transverse beam profile varies along a typical Arc cell.

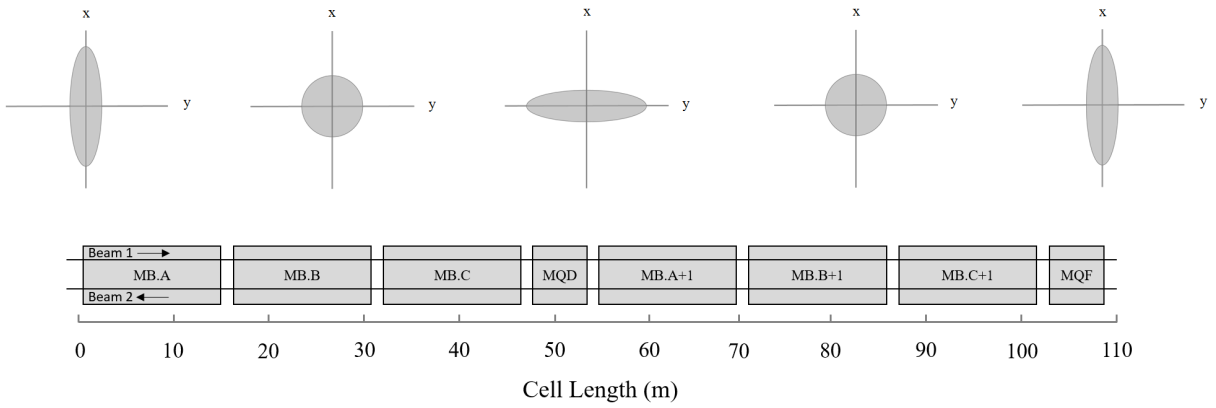

Figure 1.12: Depiction of how the transverse beam profile varies along a typical Arc cell with respect to the magnetic lattice. The beam notably varies from predominately vertical to predominately horizontal and back again every second quadrupole.

The transverse motion of the beam at any given location can be quantified by the beam emittance, ϵ, in conjunction with the corresponding Twiss parameters [37]. The Twiss parameters are beam-related Alpha, α, Beta, β, and Gamma, γ, functions which relate the emittance to the beam size and its transverse dynamics. The emittance is the area of the particle distribution in phase space and is a quantity which expresses the transverse motion of all involved particles. Figure 1.13 depicts how, in horizontal phase space (x by x'), the emittance relates the horizontal beam size and motion to the corresponding Twiss parameters.

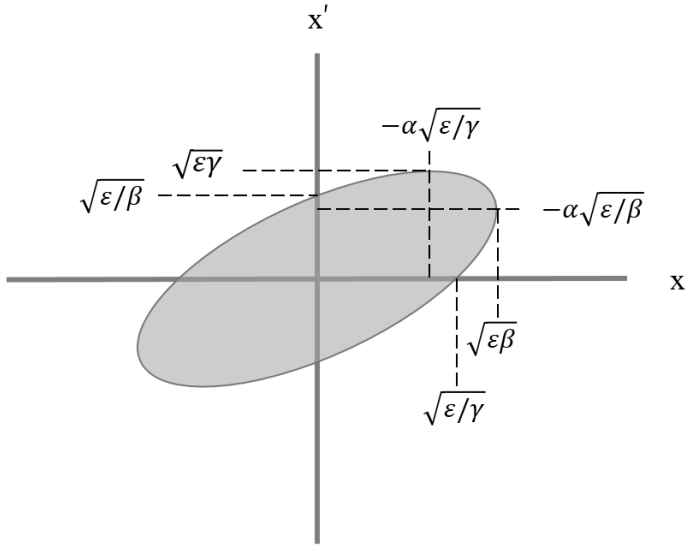

Figure 1.13: Horizontal phase space plot (x by x') of how the Twiss parameters, α, β and γ, relate the horizontal beam size and particle motion to the beam emittance, ϵ.

Chapter 2

The Main Dipole Magnet Circuits

The LHC is designed to circulate 7 TeV proton beams around a 27 km tunnel. This requires a magnetic field of ∼8.3 T. As a result, the main ring of the LHC consists of 1232 14.3-m-long superconducting dipole electromagnets, which produce a continuous magnetic field of 8.3 T when powered at a peak current of 11.85 kA. Such a magnet system requires extraordinary infrastructure, circuitry and powering systems. This chapter provides technical details on the main dipole circuits – the main magnet systems of the LHC and the subject of this thesis.

2.1 Circuit Overview

There are eight main dipole circuits, one for each sector. Each circuit contains one hundred and fifty four 100 mH 14.3-m-long twin-aperture superconducting dipole magnets and can store in excess of 1.3 GJ in magnetic energy [41]. All magnets are electrically connected in series and powered by a 13 kA/±190 V power converter (PC). Each magnet has individual bypass diodes to redirect the current in event of a quench event. Furthermore, two active energy extraction (EE) systems, consisting of high energy 75 mΩ dump resistors, provide safe means of dissipating the stored circuit energy in the event of a power abort. A simplified circuit schematic is shown in Figure 2.1.

2.1.1 13 kA/±190 V Power Converters

Each circuit is powered with a 13 kA high precision power converter, capable of stable operation across a large dynamic range of 750 A/±1 V to 12.80 kA/±190 V [42]. Each PC is constructed from two 6.5 kA/±190 V sub-converters connected in parallel, see Figure 2.2, which have water-cooled bypass thyristors capable of handling a full current decay from peak current levels. The sub-converters are connected via a series of transformers to CERN's own 400 kV powering substation [43]. An RLC filter is included to dampen any large mains voltage transients or spikes [16].

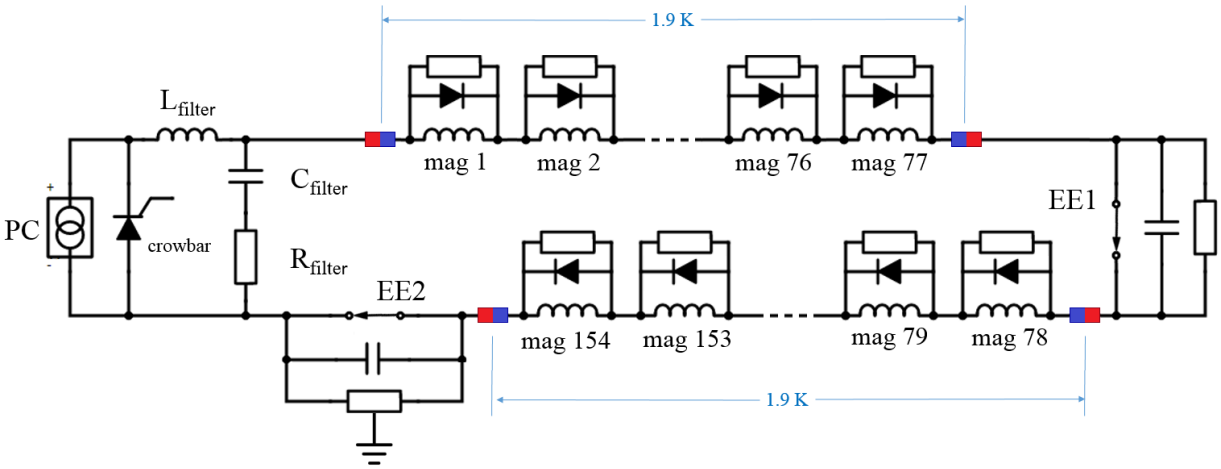

Figure 2.1: Simplified schematic of a typical main dipole circuit. Red/blue blocks represent the current leads and the transition from room temperature to cryogenic operating temperatures, 1.9 K.

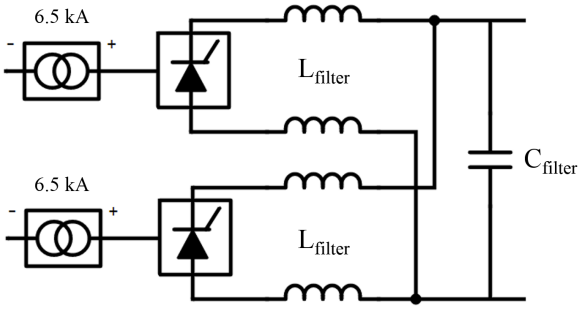

Figure 2.2: Simplified schematic of a typical LHC 13 kA/±190 V power converter. Each PC consists of two 6.5 kA/±190 V sub-converters connected in parallel and an RLC filter to suppress mains voltage spikes.

2.1.2 Main Dipole Magnets

To guide the beam around the circumference of the accelerator, the LHC utilises superconducting twin-aperture dipole magnets. A cross section is depicted in Figure 2.3. As shown, the magnet consists of a twin-aperture quasi-circular coil structure, detailed in Section 2.1.2, itself encased within a magnet collar which is itself encased within a large iron yoke structure. The magnet collar is a block of non-magnetic material, designed to compress and maintain the structure and alignment of the magnet coils [44]. The iron yoke is the main flux guide for the magnetic field. It is designed to ensure that the flux density does not saturate and that it fits within the helium cooling vessel whilst incorporating the magnet collar, heat exchanger pipes and magnet circuit busbars [44].

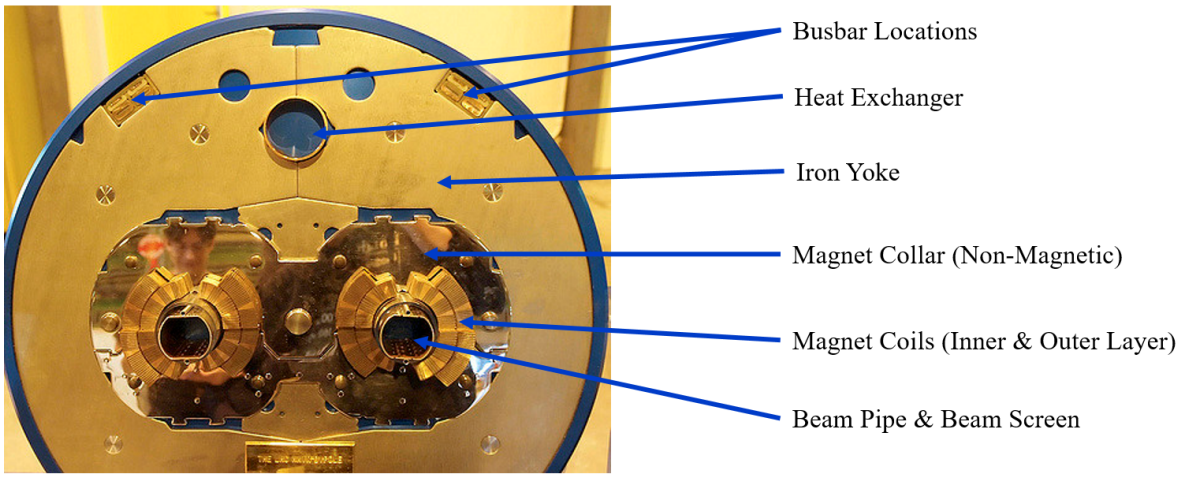

Figure 2.3: Picture of the cross section of a twin-aperture dipole magnet. As shown, the magnet structure consists of a symmetrical pair of magnet coils encased within a magnet collar, itself encased within an iron yoke. Image courtesy of CERN [46].

The LHC main dipole magnets were produced by three different manufacturers [45]. For this purposes of thesis, each manufacture, if referenced, will be referred to as Firm 1, 2 or 3, respectively.

Magnet Coils

Magnet coils are wound from Nb-Ti Rutherford cables. The cumulative total cable length of a twin-aperture dipole coil is ∼2.5 km. Coils are wound such they that surround the beam pipes in a quasi-circular geometry and consist of inner and outer layers (wound separately). Figure 2.4 depicts the middle cross section of a typical twin-aperture dipole magnet pair, showing current direction, magnetic field and Lorentz forces. Note that, with opposite current flow and proton velocities, both magnets produce Lorentz forces in the same direction, which allows for opposing beams to be bent around the ring simultaneously. Also note that coils are wound with a quasi-circular geometry, azimuthally separated by copper wedges. This is to optimise the area of the homogeneous field (good-field region) in the aperture [47].

Figure 2.5 gives a 3–D perspective of typical dipole magnet pair (for both beams), showing current and beam directions (marked in green and red, respectively) as well as the resultant magnetic field produced during operation. The coil end design can also be seen. Coils ends are optimised in order to minimise fringe fields as well as reduce quench susceptibility by ensuring that there are no sharp bends in the coil winding [48].

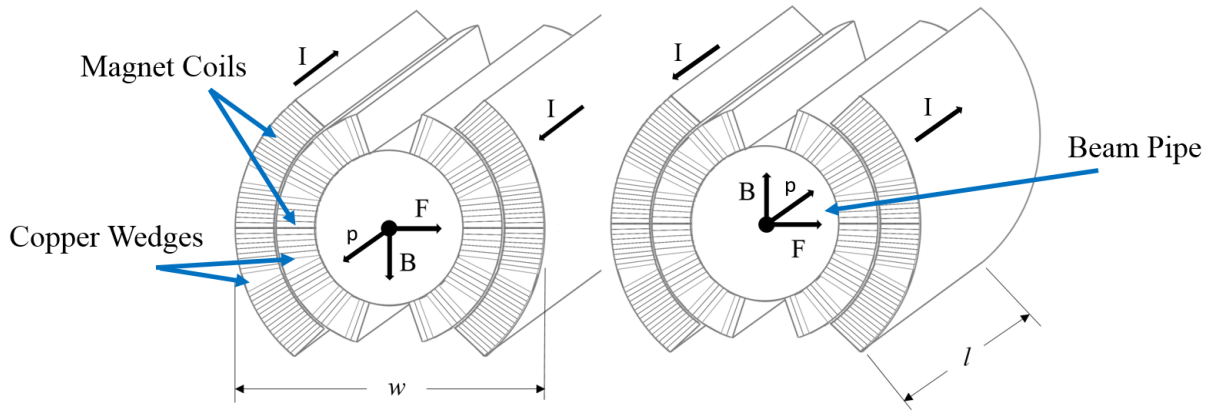

Figure 2.4: Illustration of the cross section of the coil structure of a twin-aperture dipole magnet, showing current direction, $\vec{I}$, beam direction, $\vec{p}$, and the directions of resultant magnetic fields, $\vec{B}$, and Lorentz forces, $\vec{F}$.

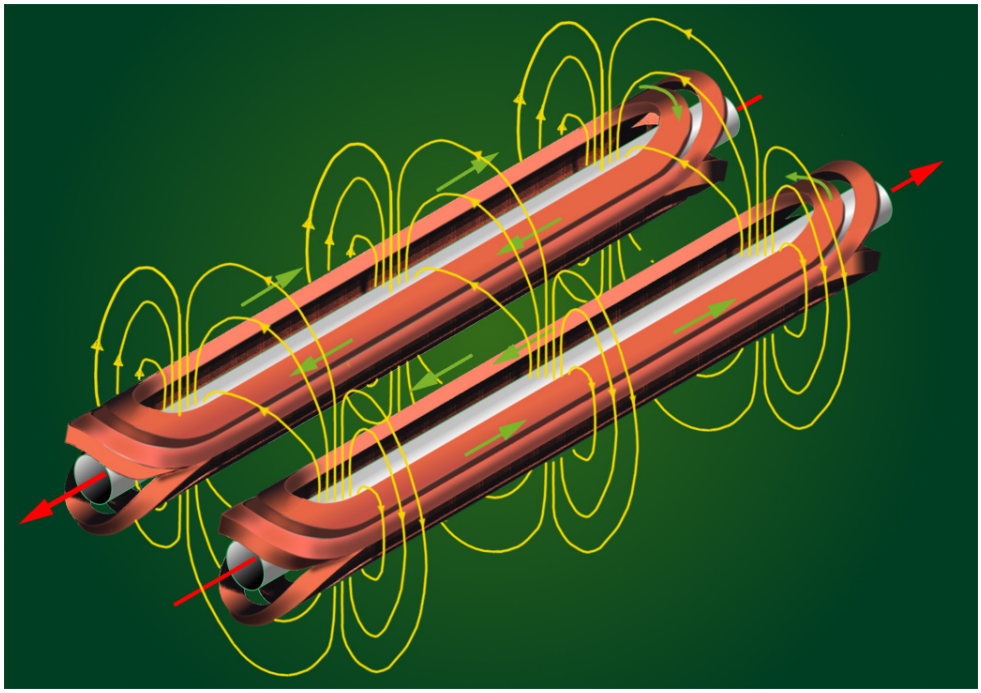

Figure 2.5: 3–D model of a typical twin-aperture magnet, including coil end geometry. Graphic shows the beam direction (red), current flow (green) and magnetic field lines (yellow). Longitudinal length is not to scale. Image courtesy of CERN [49].

Superconducting Magnet Cables

Magnet coils are wound from specially designed superconductor cables, known as Rutherford cables [50]. Each cable consists of 36 strands of superconducting wire, housing ∼6000 Nb–Ti filaments encased in high-purity Cu, see Figure 2.6. Strands are wound and shaped such that the resultant cable forms a two-strand-thick rectangle, which is then wound to create the magnet coils.

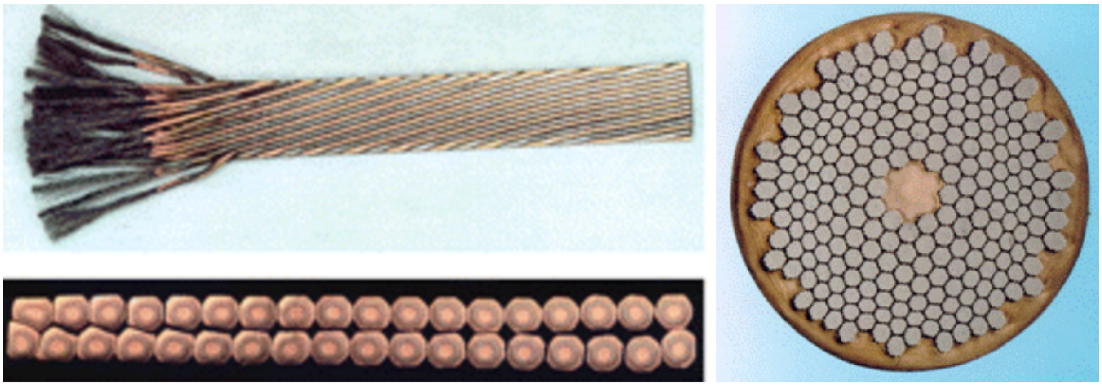

Figure 2.6: Picture set showing the internal structure of a typical 13 kA superconducting Rutherford cable. The cable (top left) of cross section (bottom left) is wound from 36 strands (right) which consist of an array of superconducting Nb–Ti filaments encased in high-purity Cu. Photographs courtesy of CERN [51].

2.1.3 13 kA Busbars

Connecting the PCs to the Arc cryostat (maintained at room temperature, 293 K), are standard water-cooled high-purity Cu cables [16], however, connecting the magnets together within the cryostat are specially designed superconducting busbars [52]. A typical busbar segment consists of a hollow Cu profile which houses a Nb–Ti Rutherford cable – an illustration of the busbar cross section can be seen in Figure 2.7.

Busbar Interconnects

Due to the physical scale of the circuit, busbar-to-busbar interconnects are located either between magnets or in long busbar sections. All interconnects are welded and machined post-installation, allowing for a maximum resistance across the interconnect of <0.6 nΩ [16]. An illustration of a typical busbar interconnect is shown in Figure 2.8.

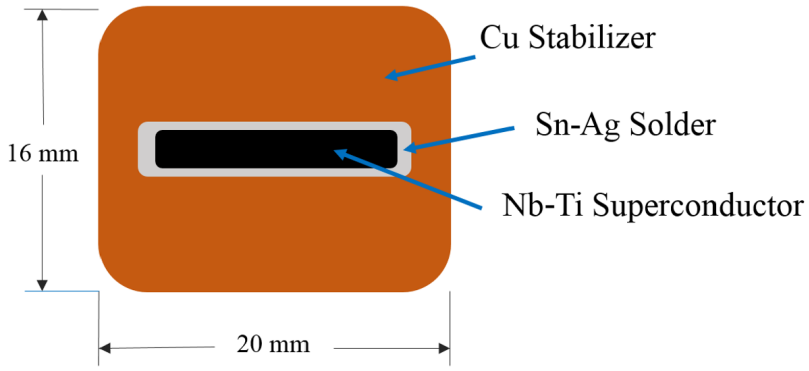

Figure 2.7: Illustration of the cross section of an LHC superconducting busbar showing the Nb–Ti superconductor, i.e. Rutherford Cable, the Cu stabiliser and Sn–Ag solder.

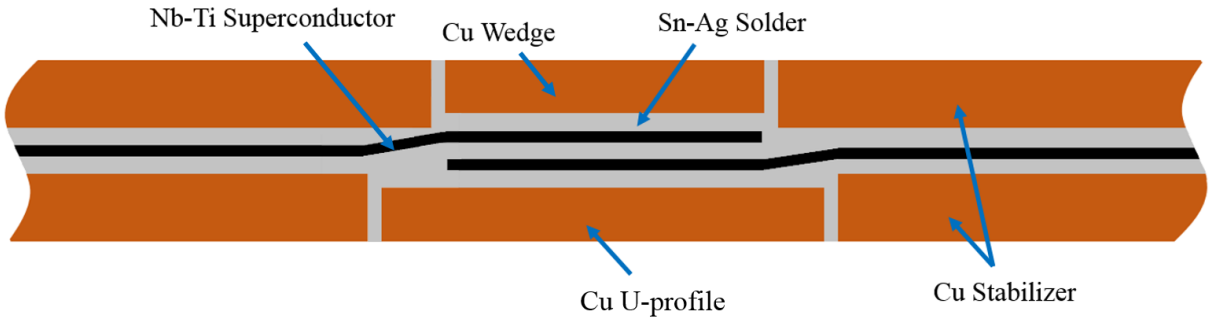

Figure 2.8: Illustration of the longitudinal cross-section of typical busbar interconnect, showing the overlapping joint style, encasing Cu wedges and where Sn–Ag solder is present.

Lyres & Pigtails

Lyres and pigtails, named after their shape, are curved segments of busbar, designed to account for thermal contraction and expansion during system warm-up and cool-down procedures [52]. These segments differ from the typical busbar in that the Cu surrounding the superconductor is constructed from layers of Cu sheets, allowing it to be curved and hence expand and contract without excessive mechanical stress. Lyres are found before and after each magnet and pigtails are found below the current leads at both ends of the Arc. Figure 2.9 depicts typical lyre and pigtail busbar segments.

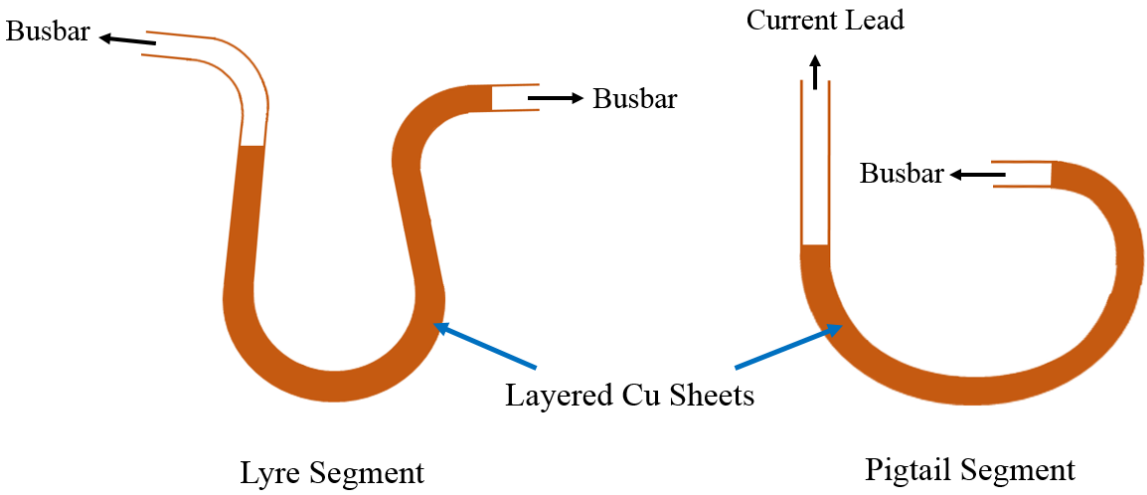

Figure 2.9: Illustration of a typical lyre busbar segment (left) and a typical pigtail busbar segment (right).

2.1.4 Current Leads & Helium Chamber

13 kA high-temperature superconductor (HTS) current leads are located at points where the magnet circuit transitions from 293 K to 1.9 K or *vice-versa* [53]. Each current lead consist primarily of two stages, see Figure 2.10, firstly gaseous-He-cooled Cu heat-sinks, which guides the initial transition from room temperature down to 20 K, followed by a section of HTS coils, which guides the transition from 20 K down to 4.5 K, He-I. The 4.5 K He-I is used to cool adjacent superconducting systems which do not require superfluid He cooling for nominal operation. Below the current lead is the liquid He chamber, see Figure 2.11, which maintains a He bath at 4.5 K as well as guides the final transition from 4.5 K to superfluid He-II at 1.9 K [16]. This transition is separated by an aptly named Lambda plug. The chamber also contains a feedback loop which provides the current lead with 20 K gaseous He.

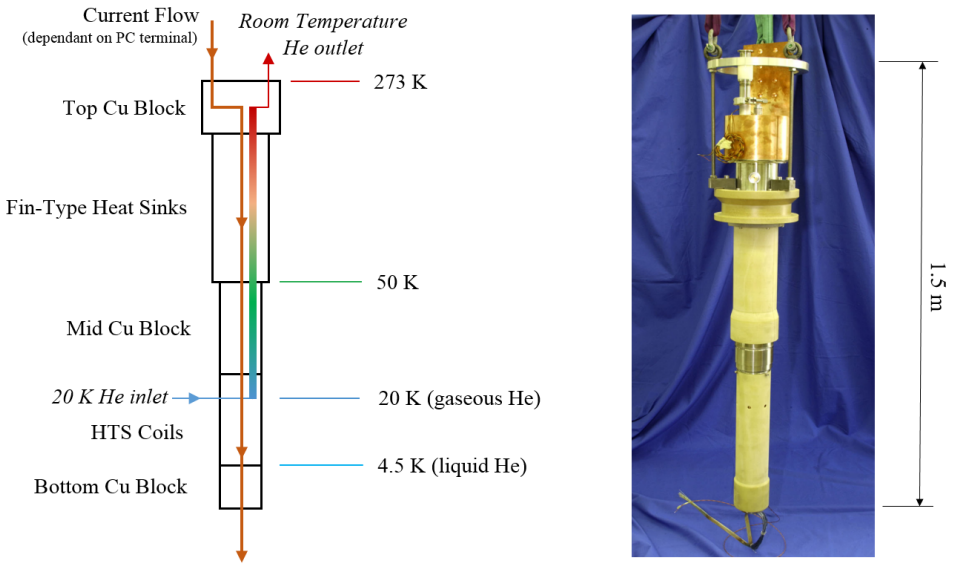

Figure 2.10: Simplified illustration (left) and picture (right) of the 13 kA HTS current leads installed in LHC for, among others, the main dipole circuits. Photograph courtesy of CERN [54].

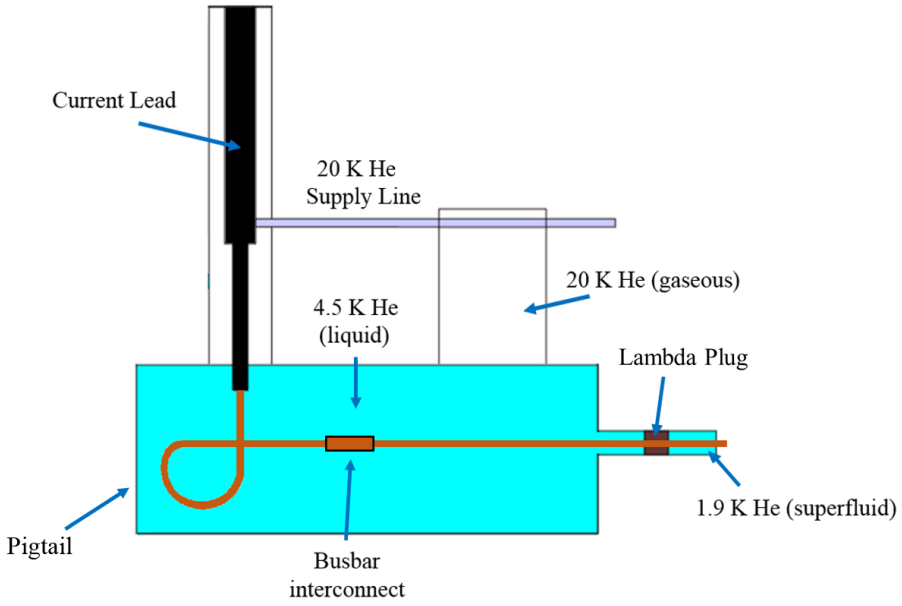

Figure 2.11: Simplified illustration of the He chamber, showing the feedback loop and position of the Lambda plug separating 4.5 K He-I and 1.9 K He-II sections.

2.2 Circuit Protection Systems

With each dipole chain storing up to 1.3 GJ of magnetic energy, it is important that the energy can be dissipated safely. This is especially the case during a magnet quench or failure event which initiates a power abort. This section introduces the installed circuit protection systems for the main dipole circuits.

2.2.1 Magnet Bypass Diodes

In the event of a quench in one of the superconducting dipole magnets, the current must be redirected to prevent significant hot-spot temperatures arising in the now resistive, normal-conducting zone. For this reason, a passive bypass is installed alongside each magnet, rated to withstand up to 13 kA with a 120 s decay constant [16]. Diodes are installed within the cryostat such that they also operate at an initial temperature of 1.9 K. The conducting voltage at 1.9 K is ∼6 V. For additional heat dissipation, each diode is connected to two large Cu heat-sinks, clamped together with a force of 40 kN. Each bypass also encompasses 6 connections, 2x busbar-to-diode-busbar, 2x diode-busbar-to-heatsink and 2x diode-to-heatsink contacts. Figure 2.12 shows a picture of the diode stack as installed in the LHC [p2].

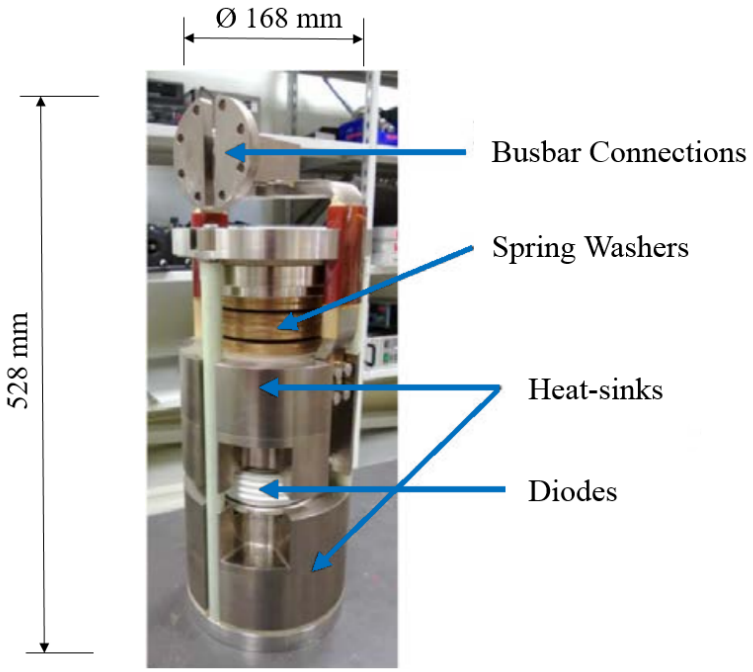

Figure 2.12: Picture of a main dipole magnet bypass diode stack, showing the diode wafer, heatsinks, springs washers to ensure constant pressure in the diode-to-heatsink contacts, and the main busbar connections. Photograph courtesy of G. Willering [p2].

2.2.2 Energy Extraction System

To ensure a relatively fast and controlled current decay in the event that a power abort is initiated, energy extraction (EE) systems are installed at both ends of the Arcs – two per circuit [55].

Figure 2.13: Picture of a single dump resistor configuration of a main dipole circuit. Each resistor, consisting of three larger resistors connected in parallel, is built to withstand the full stored circuit energy of 1.3 GJ. Together with a mechanical switch, these resistors comprise the energy extraction system. Photograph courtesy of CERN [56].

Each EE system consists of three 225 mΩ/230 MJ high-energy capacity dump resistors connected in parallel such that the total is 75 mΩ/690 MJ for each system and 150 mΩ/1380 MJ for each circuit. Break switches are included in the main bus line which open when triggered, forcing the current to flow through the dump resistors. At nominal current, this results is a decay constant of ∼104 s [16]. A typical EE dump resistor stack is shown in Figure 2.13.

2.2.3 Magnet Protection System

It is widely known that, in the event of a quench, dangerously high hot-spot temperatures can result in cable damage from melting [57]. It is therefore important that in such

circumstances, not only is the system energy extracted quickly and safely, but that the normal-conducting zone volume is deliberately increased via so called quench heaters such that the stored magnet energy becomes more homogeneously distributed, reducing peak temperatures. It is also important that this is done as soon after the initiation of a quench as possible. All magnets are built with quench heaters incorporated into the design as well as connections for voltage monitoring electronics (external to the cryostat). Collectively these components make up the magnet quench protection system (QPS) [58].

Quench Detection Electronics

To reliably and rapidly detect the initiation of a quench, high-precision voltage monitors are connected across all magnets and busbar sections [59]. The voltage across a magnet will increase as a normal-conducting zone arises and begins to propagate. This system will initiate a power abort and fire the quench heaters in the event that the voltage detected across a magnet exceeds that of a defined threshold.

Quench Heaters

Quench heaters for the main dipoles consist of pairs of ∼0.025-mm-thick stainless steel strips which are insulated and then installed in direct contact with the magnet coils, running the length of the magnets [16]. These strips are individually powered by a fast discharge (decay constant ∼75 ms) 900 V capacitor bank such that the strips heat up and that the heat is transferred to the adjacent magnet coils as quickly as possible. Quench heater temperatures can reach up to ∼280 K and therefore quickly quench the remaining superconducting regions of the magnet, distributing the hot-spot energy and minimising peak temperatures [60].

2.2.4 Beam Loss Monitoring

To prevent excessive beam losses which may result in a beam-induced magnet quench and or significant system damage, ionisation-chamber-type beam loss monitors (BLMs) are installed along the Arc [61], [67]. The structure of a BLM ionisation chamber and how it is installed in the LHC can be seen in Figure 2.14.

BLMs are typically installed on the outside of the cryostat at locations where high beams losses are expected – typically adjacent to quadrupole magnets. A BLM consists of a 60-cm-long nitrogen filled ionisation chamber with an array of Aluminium (Al) discs held at a potential of either + or −500 V, alternating every second disc [16]. The result is that incoming particles are ionised in the nitrogen and ions flow to one side of the chamber while electrons flow to other. This results in a current flow which can be read

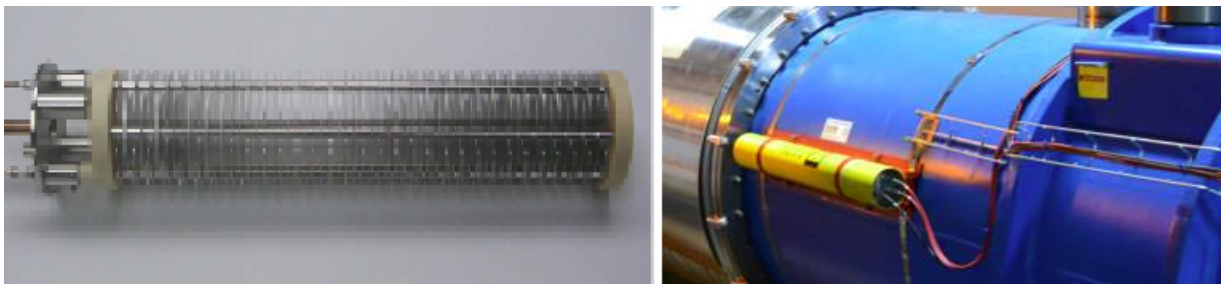

Figure 2.14: Pictures showing the structure of an ionisation-chamber type BLM (left) and how it is installed on the outside of the cryostat (right). Left image courtesy of CERN [63].

as a voltage across an output impedance and a beam dump can be initiated in the event that the detected voltage crosses a programmed threshold [62].

Chapter 3

Sustainable Beam Energies & Thesis Hypothesis

During initial commissioning of the LHC, an unexpected quench event due to the presence of busbar interconnect discontinuities led to significant system damage. Investigation into the source of the issue found that such busbar discontinuities were not only present in the circuit in question, but in all main dipole and quadrupole circuits. The result was that the LHC would only be capable of operating at approximately half the design energy throughout Run 1 until each discontinuity could be repaired. This chapter briefly discusses the damaging quench event as well as investigating possible remaining issues with achieving and sustaining near-nominal energy levels. The chapter concludes with the defining of a thesis hypothesis and corresponding prerequisites.

3.1 Run 1 Beam Energy Restrictions

3.1.1 The 2008 Incident

On September 19$^{\text{th}}$ 2008, during a commissioning powering cycle to 11.85 kA, a quench event occurred at a current of 9.3 kA in the Sector 34 main dipole circuit [64]. As the quench was not detected quickly enough, an electrical arc developed which subsequently punctured the helium enclosure, leaking evaporated helium into the vacuum. Both the normal zone and heated helium propagated along the Arc, inducing numerous other quenches and significantly increasing system pressure. All pressure release valves successfully maintained the pressure below critical levels, 1.9 MPa, with the exception of the subsector covering cells 19-21, in which pressures reached 2.1 MPa. The excessive resultant forces physically displaced several dipoles magnets from their external support jacks at cells 23, 27 and 31, example shown in Figure 3.1. In total, $\sim$2 tons of helium evaporated and escaped into the tunnel.

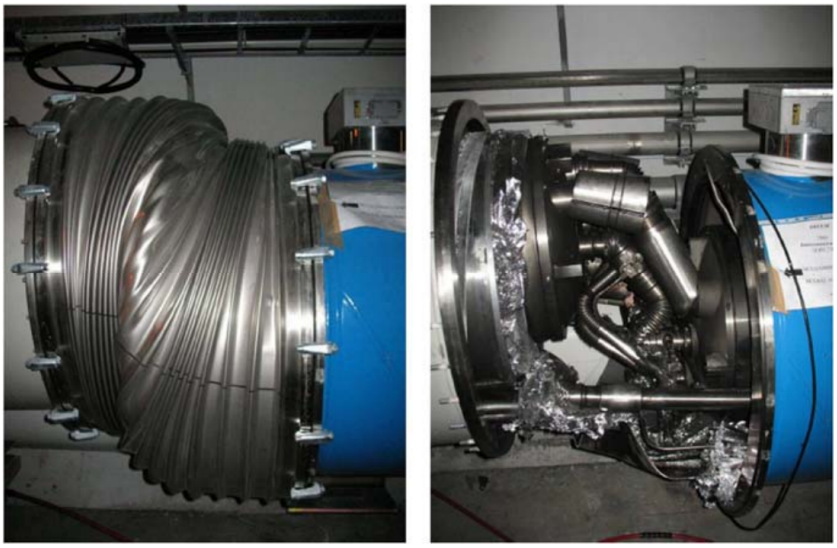

Figure 3.1: Photograph showing the displacement of a dipole magnet following the 2008 quench event, external view (left), internal view (right). Pictures courtesy of CERN [67].

The 2008 incident resulted in a yearlong repair campaign, following which the maximum beam energy was still restricted to just above half that of the design nominal, 3.5-4 TeV instead of 7 TeV [64]. Investigation into the issue led to the discovery that it was as a result of a discontinuity in one of the superconducting busbar interconnects – thought to be due to a lack of solder between the superconductor and its adjacent stabiliser, see Figure 3.2 [65]. Further investigation found that this was not a localised issue but rather such discontinuities were widespread in similar interconnects throughout all main circuits. A new QPS system, the nQPS, was installed to monitor, among other things, voltages across all busbar segments [66] and a two-year-long interconnect repair campaign was scheduled, Long Shutdown 1 (LS1).

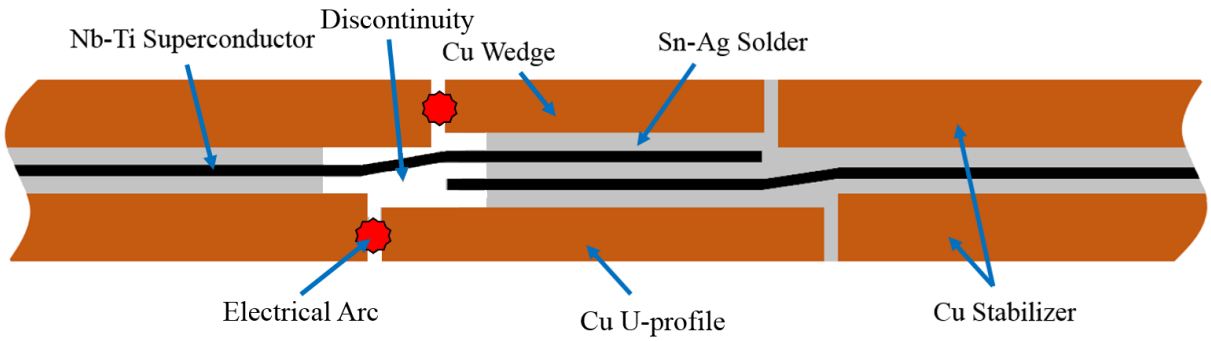

Figure 3.2: Illustration of the longitudinal cross-section of a busbar interconnect with a typical discontinuity. Note that without solder an electrical arc would form either across the Cu stabiliser gaps or even between the Cu stabiliser and external systems.

3.1.2 Interconnect Consolidation

To ensure each sector was capable of withstanding nominal current levels, every single interconnect (over 10,000) was scheduled to be opened, re-machined and have additional Cu shunts applied which bridge the gaps between the main segments of the Cu stabiliser, see Figure 3.3 [68]. Such shunts ensure that a low resistive path would remain in the event that discontinuities either remain post-consolidation or form over time.

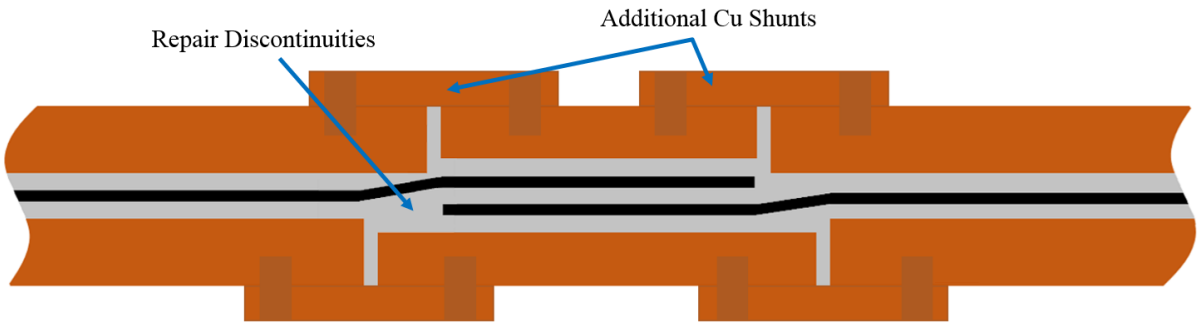

Figure 3.3: Illustration of the proposed interconnect shunts. Such shunts provide a low resistance current path in the event that solder discontinuities remain or form over time.

3.2 Achieving Near-Nominal Beam Energies

3.2.1 Magnet Training

In the field of applied superconductivity, it is well known that there is a significant disparity between theoretical (extrapolated from short-sample coil equivalents) and physical performance, *i.e.* quench limits, of full length superconducting magnets [20], [69]. Data collected from LHC magnets would find this disparity to be greatest during initial powering, *i.e.* virgin magnet conditions, as well as after significant periods without powering, *e.g.* storage prior to installation [70]. Reducing this disparity involves a process of repeatedly powering the magnet to the point of quenching. This process gradually increases the magnet's quench limits and is commonly known as magnet training. The amount of training required varies significantly even between magnets of the same design constructed by the same manufacturer, see Figure 3.4 [71]. There is also a tendency for training progress to decrease as quench limits approach the design nominal. This is due to the training mechanism itself, which involves the 'relaxing' of coil imperfections caused by the present Lorentz forces (localised cable or strand bending or misalignment) and a quench event is initiated as a result of the involved friction. Larger imperfections have a tendency to quench first and therefore relax first. As such, the largest increases in maximum current levels occurs during the first few training cycles, eventually saturating [71]. Furthermore,

note that there is a potential for de-training, *i.e.* training quenches either worsen existing imperfections and or result in imperfections being created elsewhere in the magnet coil. Magnet training to 12.85 kA was carried out on all dipoles prior to installation and was scheduled to be re-done post-installation [72].

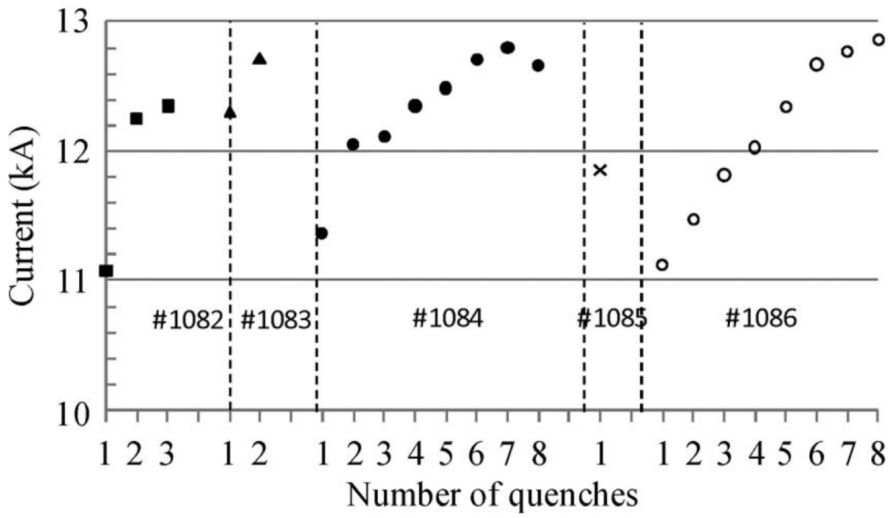

Figure 3.4: Magnet training of five LHC dipoles prior to installation. Note the significant variation in number of quenches required from one magnet to the next. Graph courtesy of E. Todesco [71].

Post-Installation Training Predictions

Following initial training, all magnets were stored for up to 2 years prior to installation. This led to a reduction of the maximum sustainable current levels for some of the magnets and, following initial system commissioning, global magnet re-training was scheduled to be carried out via systematic circuit powering. Unfortunately, as a result of the 2008 incident (Section 3.1.1), the training campaign was cancelled having not even re-trained the first scheduled circuit, Sector 56. However, with initial re-training quenches occurring at between 6 and 6.5 TeV equivalent current levels, the majority of which were Firm 3 magnets, see Figure 3.5, several studies were carried out in attempt to predict the number of training quenches required to qualify all 8 circuits, [73] and [74]. Table 3.1 lists the main prediction results, showing that the number of quenches required increases significantly with energy, requiring in excess of 80 quenches to achieve 6.5 TeV and close to 1000 quenches for 7 TeV. As a result, the studies concluded that 7 TeV operation was an unrealistic goal for Run 2.

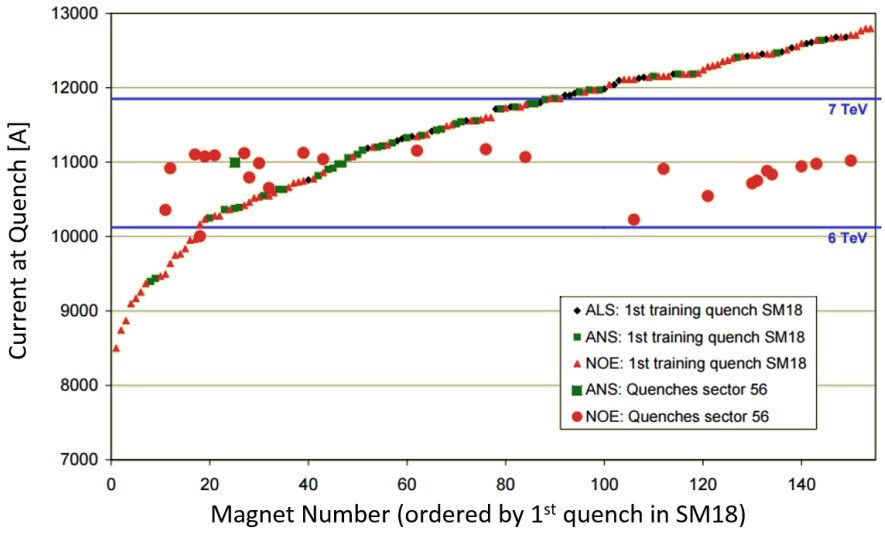

Figure 3.5: Plot of all Sector 56 dipole re-training quenches alongside the first initial training quench pre-installation of each magnet. Most post-installation quenches are from Firm 3 (NOE) magnets and lie between 6-6.5 TeV equivalent currents, far from the 7 TeV design nominal. Graph courtesy of A. Verweij [74].

Table 3.1: Post-installation dipole training predictions, circa 2009. Data courtesy of A. Verweij [74].

	No. of Magnets			No. of Quenches			
Sector	Firm 1	Firm 2	Firm 3	6 TeV	6.5 TeV	7 TeV[1]	7 TeV[2]
12	49	96	9	0	4	41	40
23	56	60	38	1	8	97	92
34	56	65	33	1	8	87	83
45	46	46	62	2	12	145	136
56	28	42	84	1	15	190	178
67	57	36	61	2	12	142	133
78	54	40	60	2	12	140	132
81	64	24	66	2	13	151	142
Total	410	409	413	11	84	993	936

[1] Estimates based on extrapolation of Sector 56 data
[2] Estimates assuming 2 quenches per Firm 3 magnet

3.2.2 Remaining Uncertainties & Project Proposal

Although an interconnect discontinuity was found to be the cause of the 2008 incident, it may not be the only limitation preventing the successful achievement of nominal energy levels. With 6.5 TeV training predictions foreseeing in excess of 80 quenches, it is important that all other potential weak points in the system, such as the connections in the magnet bypass path (Section 2.2.1) or in the curved busbar sections, the lyres and pigtails (Section 2.1.3), are identified and also fully qualified. At this point, a new method of qualification for superconducting magnet systems was proposed, see Chapter 4.

Throughout LS1, further justification for such a procedure arose as the accumulated resistance of each magnet bypass' connections were measured. Most notably, even having been measured following installation, *circa* 2008, a number of bypass paths far exceeded the specified maximum resistance of 15.5 $\mu\Omega$ – one even exceeding 200 $\mu\Omega$, see Figure 3.6. Furthermore, the original 2008 measurement of the bypass in question was missing from the original database. This observation was not only an issue in that the LHC had being operating for 3 years with a potentially catastrophic flaw, but was also an important reminder that even with systematic measurements, human error is inevitable in systems with such numerous components. All bypasses exceeding the threshold were re-machined accordingly.

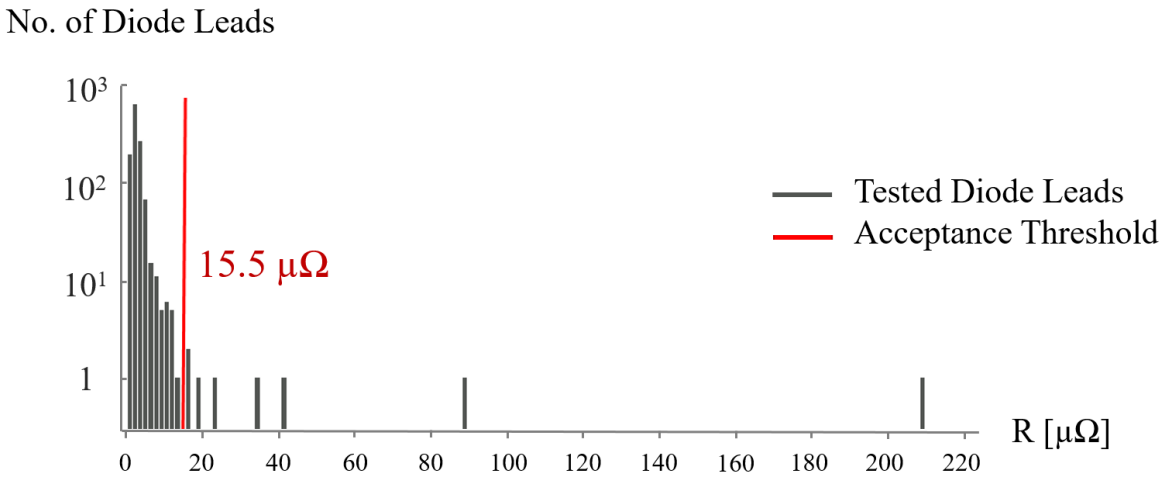

Figure 3.6: Diode bypass resistance data from 4 of the 8 main dipole circuit. Note that several bypasses exceed the resistance thresholds with one even being measured at 14 times the limit. Data courtesy of M. Bednarek [75].

To add to the this, during LS 1, all pigtail busbar segments were checked and a number of them, in fact, showed significant deterioration and or fraying. The pigtails in question were re-machined and repaired, however, the more numerous lyre busbar segments (made

from the same layered design) were neither scheduled for measurement nor repair and without proper qualification, there would be no means of asserting the individual capability of such components.

The development, type-test and full qualification results of the devised method are detailed in Chapter 4.

3.3 Sustaining Near-Nominal Beam Energies

With the end-goal being the collection of valid collision data, it is important that near-nominal beam energies cannot only be achieved but can be sustained. However, numerous issues arise in sustaining a stable beam [76], the most troubling of which is the resultant beam losses that follow particles of dust or debris interacting with the beam. Observed during Run 1, see [77] and [78], the so called unidentified falling object (UFO) events may pose significant threat to system availability at 6.5 TeV operation.

3.3.1 Unidentified Falling Objects

Dust and or debris is well documented to have had significant detrimental effects in previous electron and anti-proton accelerators, see [79]–[82]. In such cases, positively charged macroparticles would become trapped within the beam, causing beam losses due to inelastic collisions over time, reducing beam lifetime and affecting availability. The LHC is, however, the first proton accelerator where similar events have been observed.

Run 1 measured BLM signals show that UFO events typically produce loss signals of $\sim$1 ms in duration which follow asymmetric Gaussian profiles with faster fall times than rise times, an example is shown in Figure 3.7 [83]. An accepted interpretation of an LHC UFO event is as follows:

- A macroparticle falls from the ceiling of the beam screen or vacuum chamber.
- Inelastic collisions between incident protons and macroparticle nuclei result in significant beam losses. These losses result in an initial rise in BLM signal.
- Concurrently, the macroparticle is ionised due to elastic collisions with incident protons releasing electrons from the macroparticle itself.
- Subsequent to ionisation, the macroparticle is repelled by the beam's electric field. This results in a subsequent fall in the recorded BLM signal.
- The overall BLM signal produced is asymmetric due to the repulsion forces and resultant macroparticle speed exceeding that which is solely due to gravity.

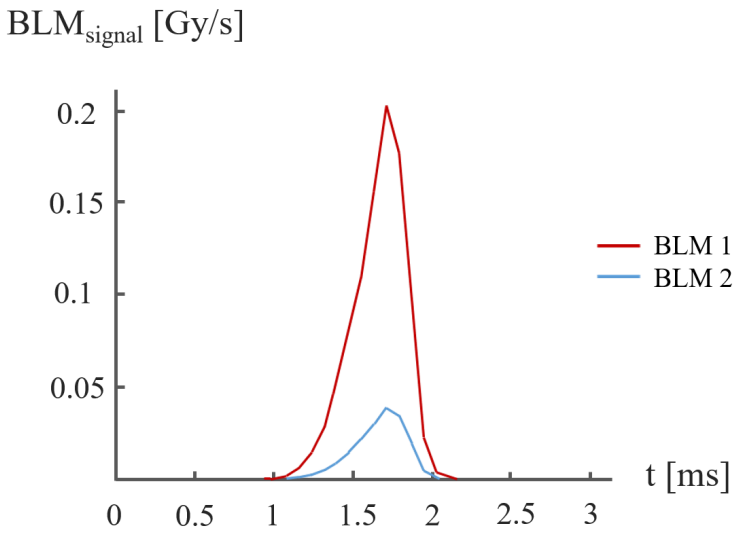

Figure 3.7: Example of a UFO event BLM signal measured during Run 1 (recorded 04/07/2012). The loss signal is ∼1 ms in duration and follows an asymmetric Gaussian profile, with a faster fall time than rise time. Data courtesy of T. Baer [83].

Dust/Debris Size

To investigate the size of a typical macroparticle, in October 2011, dust and debris was collected from a replaced main injection kicker (MKI) magnet (the magnets most susceptible to UFO events during Run 1) [83]. Figure 3.8 shows a microscopic image of a typical macroparticle found in this study.

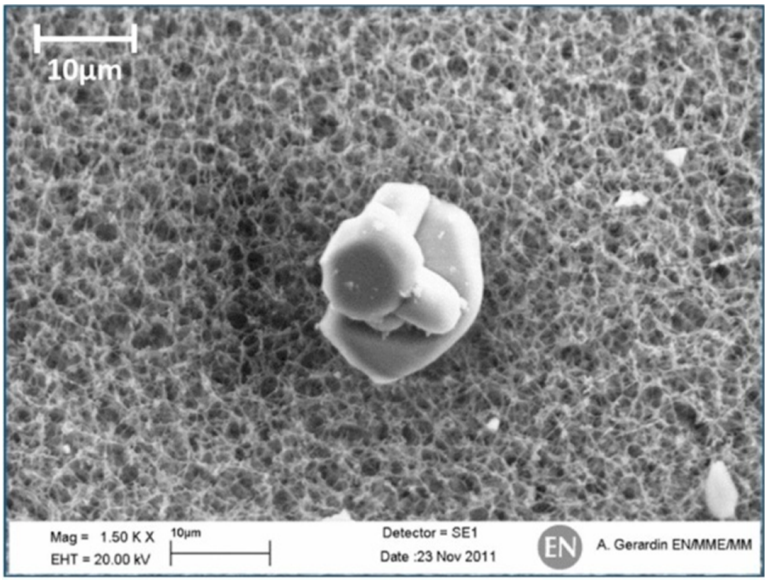

Figure 3.8: Microscopic zoom of a typical dust particle found in an main LHC kicker magnet, circa 2011. Image courtesy of T. Baer [83].

The study found that, compared to a level 100 clean room reference value of 10,000 macroparticles for an unused MKI beam pipe, flushing out the used kicker magnet resulted in a filter count of over 5 million macroparticles [83]. The majority of the macroparticles were found to have a diameter in the order of ∼1 μm, with some being notably much larger, surpassing ∼100 μm. However, as the majority of the particles were also Al, and that the aluminium-oxide (Al_2O_3) ceramic tubing is unique to the MKI beam screen, the extent of correlation between the dust collected in this study and that which is present in the LHC Arc dipoles is unknown. This is especially the case in considering which macroparticles are most likely to detach from the beam screen surface and fall – another remaining unknown.

3.3.2 Run 1 MKI UFO Event Statistics

There are two beam injection regions in the LHC, one at point 2 for beam 1 and one at point 8 for beam 2, respectively. There are 8 MKI magnets, 4 at each region. Of the ∼4000 UFO events recorded in these regions, which comprise only ∼0.06% of the LHC length, 21 UFO events produced sufficient losses to initiate a beam dump [83]. However, is was discovered that the MKI UFO rate was directly correlated with pulsing the MKIs and the resultant surface electric fields. Furthermore, pulsing similarly designed quadrupoles with metalized ceramic tubing, the MKQs, resulted in not a single UFO event being recorded [83]. It was therefore concluded that the source of the MKI UFOs was the ceramic tubing.

Investigating MKI UFO event mitigation strategies led to very successful improvements. Improvements included: a more rigorous cleaning procedure to reduce the number of existing dust particles; the ceramic tubes were upgraded to have 24 screen conductors instead of 15, reducing the maximum surface electric field by ∼40%; a Non-Evaporable-Getter (NEG) coating on the magnet interconnects and bypass tube to minimise electron multipacting and electron-cloud build-up; and, for the Hi-Luminosity LHC upgrade, a chromium-oxide (Cr_2O_3) coating of the ceramic tubing is being investigated [83]. Preliminary testing of these improvements have concluded that MKI UFOs should not be a threat to LHC availability throughout Run 2.

3.3.3 Run 1 Arc UFO Event Statistics

Approximately 12,500 Arc UFO events were recorded throughout Run 1. Figure 3.9 shows the recorded Run 1 Arc UFO event statistics, in UFOs/hour, with beam energies ranging from 3.5 to 4 TeV [83]. When looking at the shown statistics, it is important to note that UFO event monitoring only began after the phenomenon was first witnessed in 2011, while the LHC had already been operating at 3.5 TeV for ∼7 months between April and

November 2010. Furthermore, UFO events are only recorded if the resultant BLM signal is high enough to exceed a defined threshold (discussed in Chapter 5).

There are several notable trends. Firstly, at the start of 2011, the UFO rate was significantly higher, ∼8-12 UFOs/hour, with the average rate reducing over the course of the first year to ∼2 UFOs/hour – this effect has become known as UFO 'conditioning'. It is, however, assumed that a potentially significant amount of conditioning may have occurred throughout 2010 prior to event logging *i.e.* initial UFO rates may have been higher. It is also of note that, following a technical shutdown over the Christmas period of 2011, the UFO rate notably rose again prior to reconditioning. Neither of the phenomena involved with worsening or conditioning are understood. Lastly, there is a notable spike in UFO rate towards the end of 2012, where the LHC underwent a test period of operating with 25 ns bunch spacing. Throughout this period, the UFO rates increased significantly, even exceeding initial figures. The current hypothesis is that the use of 25 ns bunch spacing has a direct effect on a UFOs release mechanism, *i.e.* causes more macroparticles to detach from the beam screen. This has been thought plausible as, if the dust particles themselves were attached to the beam-screen as a result of unintended cryo-condensation or cryo-sorption [84], the increase in heat may reduce the surrounding adhesion forces allowing the particles to fall. Note, however, that there is no evidence to back up this theory. 25 ns bunch spacing is expected for 6.5+ TeV operation [85].

3.3.4 UFO-Induced Dipole Quenches & Project Proposal

It is well known that beam losses in high-energy accelerators, either from magnet optics errors or system failures, can produce sufficient localised heating to initiate the quench of adjacent superconducting magnets, *e.g.* HERA [86], Tevatron [87] and, more recently, the LHC [21]. However, throughout Run 1 of the LHC, with operation energies of 3.5-4 TeV, UFO events were not regarded as an issue for the main dipole circuits as the deposited energy from resultant particle showers was far below the quench level of adjacent magnets. It is thought that with 6.5 TeV beams, especially with magnet quench margins also being 2-3 times lower (Figure 1.2), UFO-induced dipole quenches may be a possibility and, if left unaccounted for, could have a significant impact on overall system availability. It was therefore proposed that a numerical model be constructed capable of simulating a typical LHC UFO event, giving accurate energy deposition estimates from the resultant particle showers.

The development of the numerical model and simulation results of both 4 TeV and 6.5 TeV UFO events are detailed in Chapter 5.

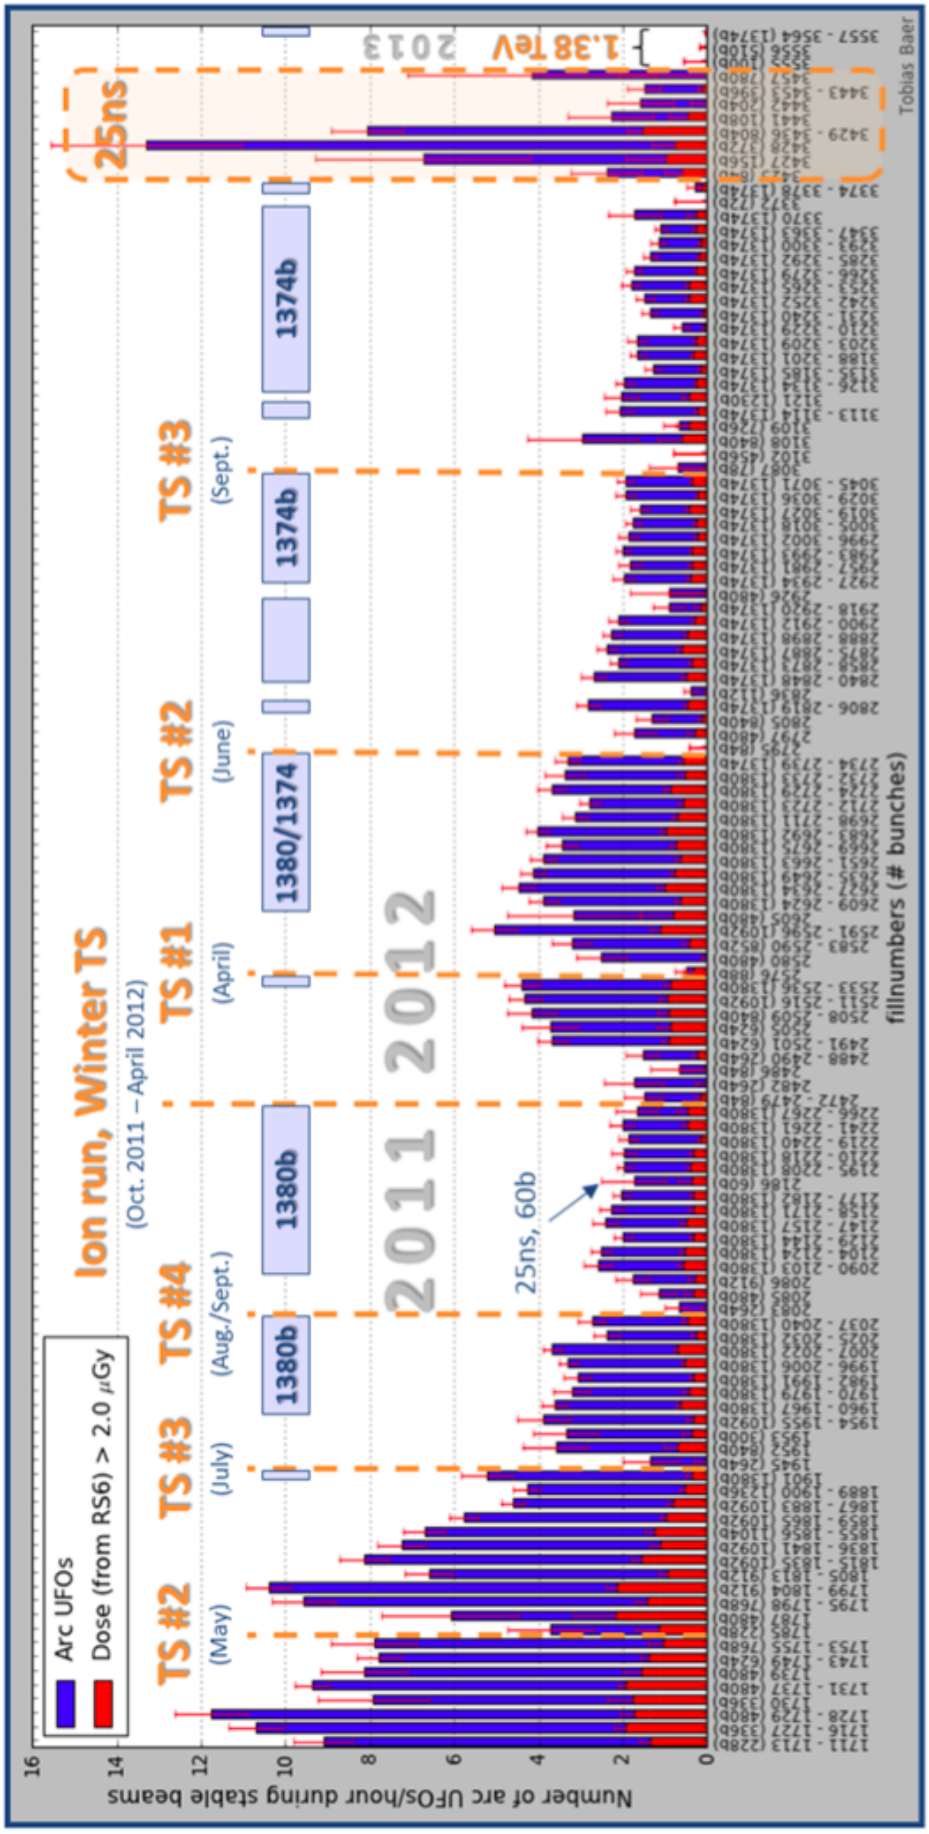

Figure 3.9: Plot of Run 1 Arc UFO rate statistics from 2011 and 2012. The bottom axis lists the operation period in 'fill' numbers (every beam injection attempt is given a corresponding fill number) as well as the corresponding number of bunches – periods with a near-nominal number of bunches (1380b) are also highlighted in pale blue. All technical shutdown (TS) periods are marked for reference. A test period in which the LHC was operated with 25 ns bunch spacing towards the end of 2012 is also highlighted. Image courtesy of T. Baer [83].

3.4 Thesis Hypothesis & Prerequisites

3.4.1 Thesis Hypothesis

Following successful circuit qualification and dipole magnet training, the LHC's main dipole magnet circuits can safely support and sustain near-nominal proton beam energies of at least 6.5 TeV.

3.4.2 Prerequisites

1. To ensure magnet training can be carried out both systematically and without risk, the development and subsequent implementation of a new method of superconducting magnet bypass qualification is required.
2. Magnet quenches following dust-to-beam interaction scattering are thought to be a possibility at 6.5 TeV operation. Further understanding of the involved phenomena and research into potential mitigation strategies is therefore required.

Addressing 1st Prerequisite

- Develop and investigate both the risk and feasibility of a global superconducting magnet circuit continuity qualification method.
- Carry out the developed qualification method on all main dipole circuits prior to magnet training.
- Upon successful qualification of all circuits, carry out magnet training as planned (up to 6.5 TeV equivalent current levels).

Addressing 2nd Prerequisite

- Develop a numerical model capable of simulating dust-to-beam interactions as well as quantifying the resultant particle shower energy deposition in adjacent magnets.
- Use the model to determine the probability of a 6.5 TeV UFO event resulting in a magnet quench, quantifying probable losses to system availability.
- Investigate and propose mitigation strategies.

The following chapters will detail the work carried out to address the defined prerequisites. Chapter 4 details the development and implementation of the copper stabiliser continuity measurement, a new method of superconducting magnet circuit qualification. Chapter 5 will detail the development of a UFO event model and subsequent predictions of plausible UFO-induced quenches and Chapter 6 will discuss initial Run 2 observations and statistics. The thesis will conclude with Chapter 7, recalling the above hypothesis and discussing project outcomes.

Chapter 4

Copper Stabiliser Continuity Measurement

To address risk factors involved with resuming magnet training, a new method of qualification was developed to provide conclusive evidence that all the potentially weak components of all dipole circuits, i.e. the busbar's Cu stabiliser and magnet bypass circuitry, were safe to operate at a given current level. This chapter, based on works presented in listed publications [p2]*,* [p3] *and* [p5]*, details the procedure development, thermal runaway simulations and type test and LHC-wide qualification results.*

4.1 Background & Concept

As discussed in Section 3.1.1, due to the presence of busbar interconnect discontinuities, an unforeseen quench event occurred during initial commissioning of the main dipole magnets, *circa* 2008. To address this issue, the LHC was scheduled to have all corresponding interconnects, over 10,000, be consolidated, *i.e.* repaired or re-machined, during LS 1. Moreover, subsequent to consolidation, magnet training to 6.5 TeV would, following predictions, take at least 80 training quenches. As a result, to compliment the scheduled consolidation campaign and to eliminate any residual risk posed by magnet bypass circuitry and lyres, it was proposed that a method of stabiliser continuity qualification be developed and carried out to provide conclusive evidence that all aspects of a given circuit, with the exception of the magnets themselves, can safely sustain near-nominal current levels. Focused on continuity qualification, the method became commonly known as the Copper Stabiliser Continuity Measurement (CSCM).

CSCM was originally thought of as a means of determining the maximum safe operating current [88]. In principle, it involves operating the circuit at just above the critical temperature of Nb–Ti of 9.2 K at ∼20 K – gaseous helium. Operating above the critical temperature would see the circuit's magnets and busbars operate in a resistive state, forcing all diodes to conduct upon powering with sufficiently high voltage as well as having the current flow through the busbar's Cu stabiliser and magnet bypass circuitry. The reason

for remaining at ∼20 K was such that the circuit would still have a relatively low circuit resistance whilst leaving sufficient thermal margins to account for potential cryo-instability, *i.e.* variations in temperature along the Arc. The resultant current flow is illustrated in Figure 4.1. Qualification of the circuit follows as if significant defects or discontinuities remain present in the busbar or magnet bypasses, the excess resistance would induce a thermal runaway if current densities exceed critical values. As such, a CSCM inherently tests all busbar sections, including the interconnects, and all magnet bypasses, including diodes and diode lead connections.

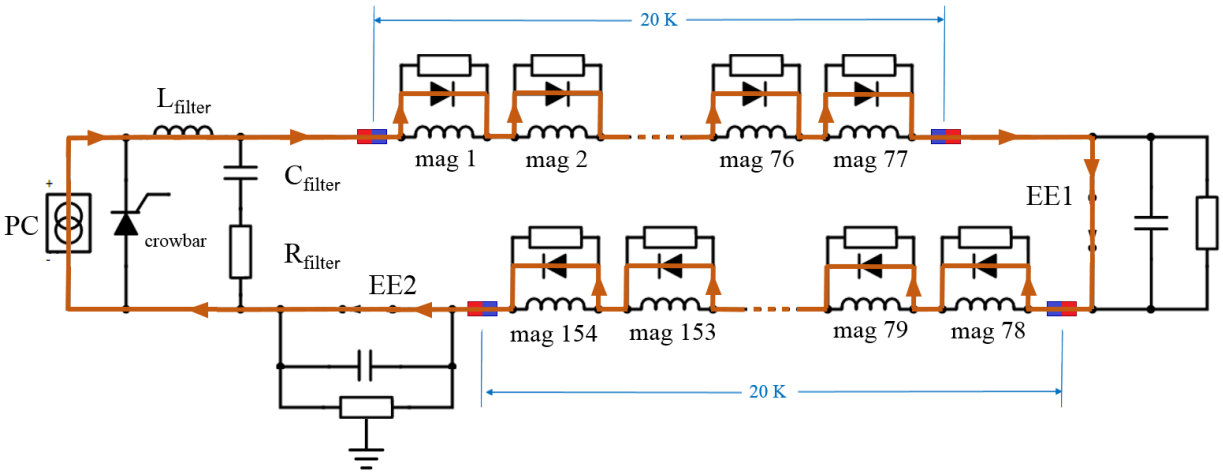

Figure 4.1: Illustration of the current flow during a typical CSCM test on a main dipole circuit (marked in brown). Note that all current flows through the bypass diodes instead of the magnets.

4.2 Engineering Challenges

Consequential to the non-standard operating conditions required for a CSCM test, a number of engineering challenges arose. This section summarises the main issues as well as their respective solutions or resultant test constraints.

4.2.1 Operating at 20 K

CSCM requires an operating temperature of ∼20 K. This posed the problem of stabilising the, now gaseous, helium bath along the 2.5 km Arc cryostat. The defined specifications were to maintain 20±4 K at a pressure of 5±0.5 bar, giving sufficient margin for the maximum expected energy transfer to the helium, ∼250 kJ. The Arc cryostat pressure limit is 15 bar [16]. Figure 4.2 shows both measured temperature and pressure profiles along the Arc, during the Sector 23 type test in 2013.

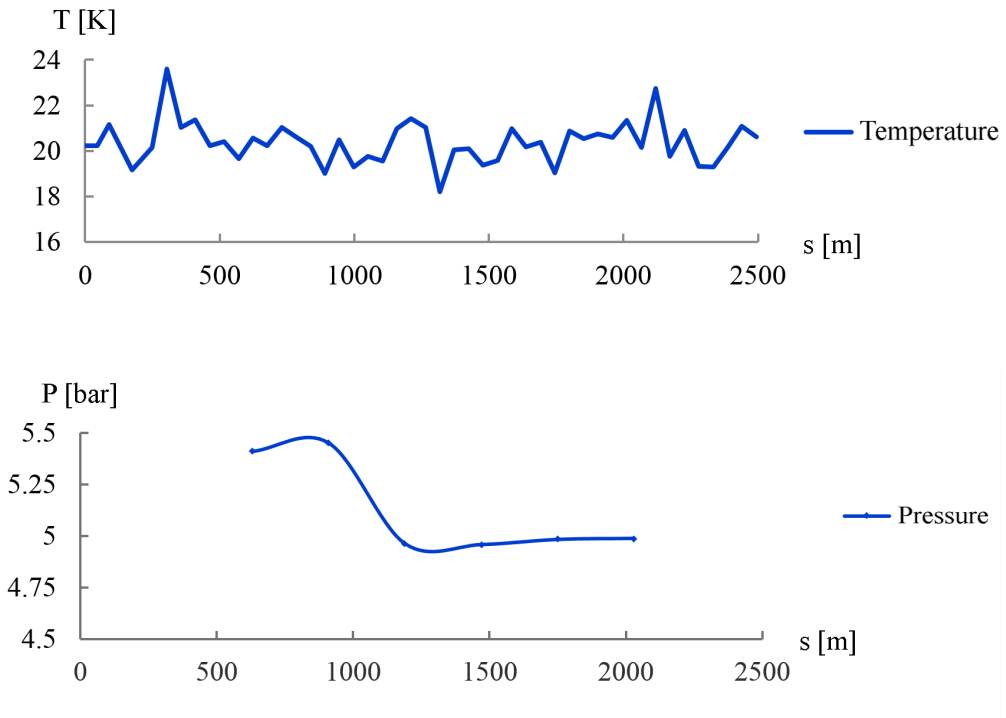

Figure 4.2: Temperature (top) and pressure (bottom) profiles along the Arc measured during the Sector 23 CSCM type test in 2013. s is the longitudinal coordinate along the Arc at which a given monitor is located. Test specifications were 20±4 K and 5±0.5 bar, respectively.

Current-Lead Restrictions

In view of a full qualification CSCM, with the intention of qualifying all aspects of all circuits, a current-lead specific test was carried out to investigate whether or not a CSCM test could be carried out with the current lead He chamber (Section 2.1.4) also stabilised at 20 K – normally 4.5 K. The test, however, concluded that the voltage across the current-lead HTS coils would exceed the quench protection threshold of 3 mV. Note that, in the event of a quench in the HTS coils, due to the nature of its layer tape design, rapid expansion of any trapped helium may result in permanent damage, even if the current is made to decay quickly [16]. As a result, increasing the protection thresholds was not an option and the He chambers were to remain at 4.5 K. Due to this, the encompassed busbar pigtails as well as the first 1-2 meters of busbar after the Lambda plug would remain in a superconducting state and, as such, would *not* be qualified by the test. This was deemed acceptable as all pigtails were individually inspected and repaired throughout LS 1.

4.2.2 Power Converter Configuration

As the main dipole circuit is several kilometres long, operating it in a resistive state, especially with all diodes initially conducting at the same time, requires significantly more voltage. It was expected that operating at 20 K would require up to ∼350 V. The required

voltage would be much higher if the diodes were not also also maintained at 20 K, as this reduces the individual diode conducting voltage from ∼6 V down to ∼2.1 V. Nonetheless, the PCs standard configuration, 13 kA/±190 V, would be incapable of supplying the required voltage. The solution was to connect the PC's two 6.5 kA/±190 V sub-converters in series, rather than in parallel – illustrated in Figure 4.3. The resultant 6.5 kA/±380 V configuration was tested by the PC team and could successfully provide the required voltage with sufficient precision, ±1 mV. Note that a full qualification test would, however, require up to 12 kA and further tests were carried out to determine if the PC could safely operate in an 'over-current' state for the required, relatively short, duration of ∼60 s. All tests proved to be successful [89].

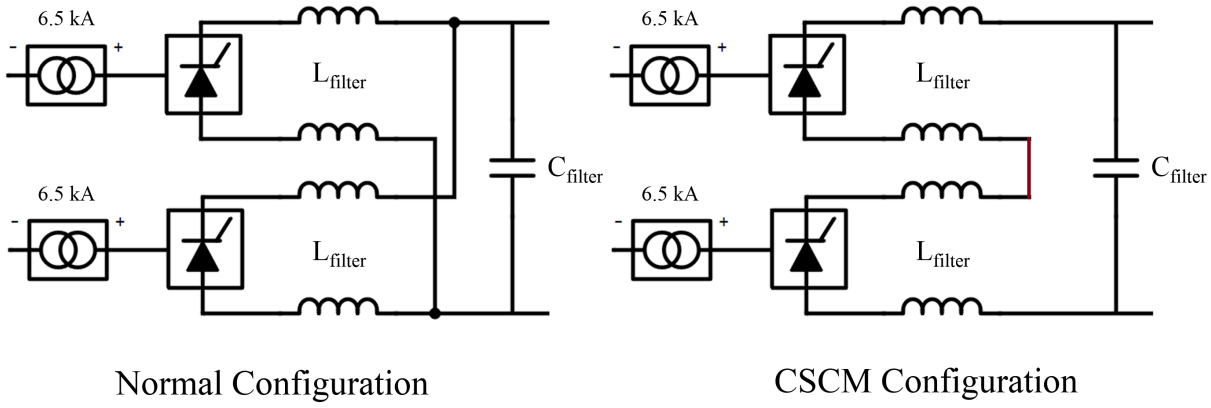

Figure 4.3: PC configurations – parallel (normal operation) (left) and series (CSCM) (right), providing 13 kA/±190 V and 6.5 kA/±380 V, respectively.

As all magnets are bypassed, the circuit has a significantly reduced inductance, 3 mH instead of 15 H. With a decay constant of only 3 ms, reduced from 104 s, the PC also had to be programmed to deliberately perform an exponential decay mimicking that of normal operation. Furthermore, both EE systems, being unnecessary due to the low circuit inductance, were shorted to minimise system risk [89].

4.2.3 Circuit Protection

With the possibility of a thermal runaway occurring during a CSCM test, minimising system risk was of paramount importance. A CSCM test would therefore only be deemed safe if a developed protection system could successfully detect and protect against a thermal runaway by means of initiating a power abort. The CSCM protection was developed by the QPS team, implementing modifications to the already installed nQPS system. The utilised electronics allowed for both monitoring and logging of the voltage, U, as well as the first and second derivatives, $\mathrm{d}U/\mathrm{d}t$ and $\mathrm{d}^2U/\mathrm{d}t^2$, respectively, across individual bus-

bar segments (board type (m)BS_A), diodes (board type DS_A) and diode-leads (board type (m)BS_B and DS_B). Note that d^2U/dt^2 monitoring was included for long busbar segments, >100 m, such as those present in the main quadrupole circuits and were therefore not used for any main dipole circuit tests. The final board configuration and main specifications are shown in Figure 4.4 and Table 4.1, respectively.

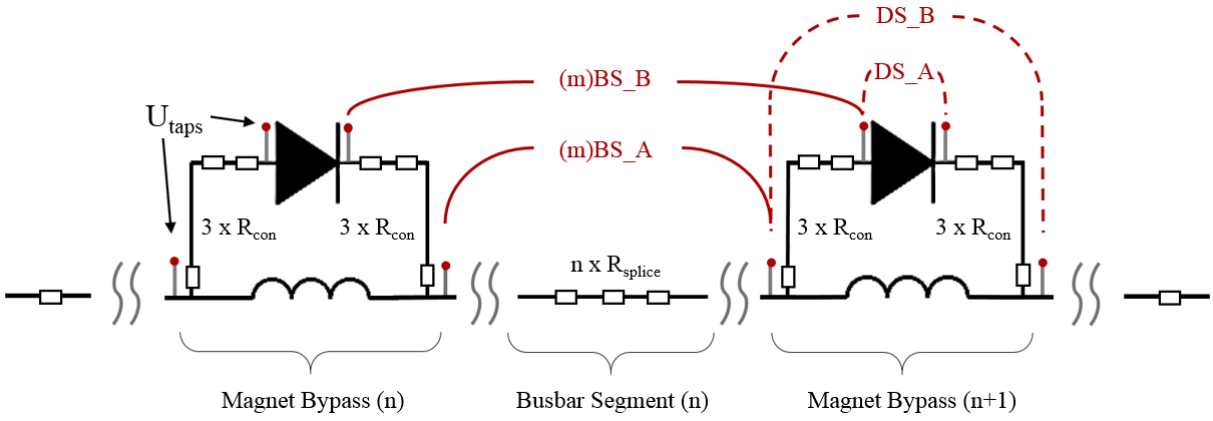

Figure 4.4: Illustration of the PS board locations as implemented for CSCM to monitor and log the busbar (board type BS_A), diode (board type DS_A) and diode-lead voltages (BS_B and DS_B).

Table 4.1: CSCM PS board specifications.

Board Type	Detection Frequency [Hz]	Logging Frequency [Hz]	Input Range [V]
mBS_A	33	33	±2.56
mBS_B	33	33	±2.56
BS_A	33	0.8	±2.56
BS_B	33	0.8	±2.56
DS_A	6410	6410	±10
DS_B	6410	6410	±10

For the type test, CSCM specific busbar monitoring electronics were used, denoted 'modified' mBS boards. However, for the LHC-wide campaign, the BS boards used for normal operation were altered such that they incorporate a 'CSCM mode', preventing the QPS operatives from having to exchange all boards both before and after testing. This was, however, at the expense of data logging frequency.

4.3 CSCM Procedure Development

4.3.1 Risk Assessment

Similar to a quench event, if a thermal runaway was to occur, excessive hot-spot temperatures would arise and may result in significant system damage – cable melting, electrical arcing, *etc.* As such, a risk assessment was carried out to investigate plausible failure scenarios as well as to incorporate any relevant mitigation strategies. Following the principle that risk is the product of probability and consequence [90], one can determine the severity of a failure scenario in accordance with the following risk matrix, Figure 4.5.

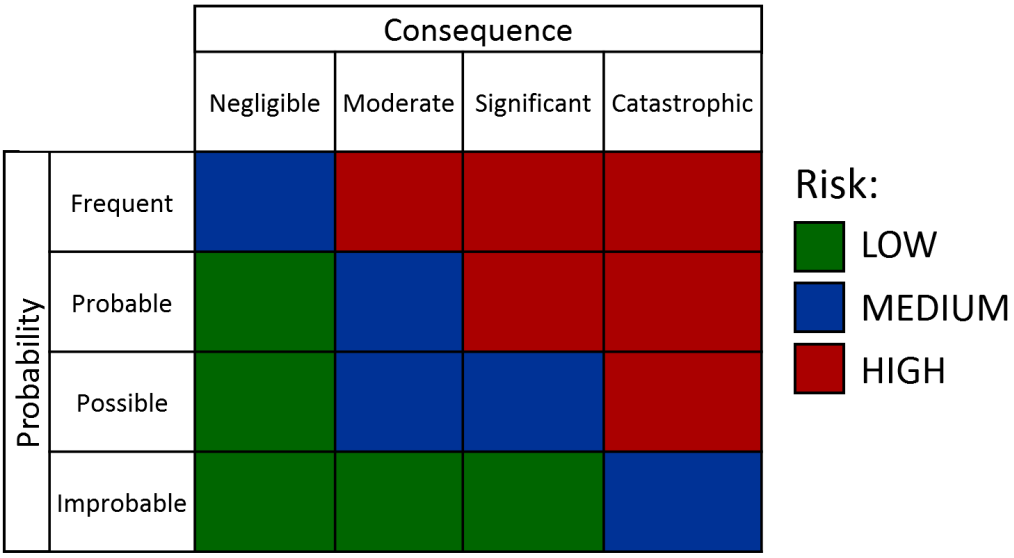

Figure 4.5: Risk assessment matrix. Risk = Probability × Consequence.

Focusing on a dipole circuit's main systems and components, a risk assessment was carried out for the proposed qualification test, both in assuming no mitigation steps would be taken and in assuming that mitigation proposals would be implemented. Results can be found in Table 4.2 and Table 4.3, respectively.

It is clear that, due to the likelihood of thermal runaways, with plausible catastrophic outcomes, CSCM is inherently a high risk test. To minimise risk, it was decided that a type test should be carried out on one of the circuits prior to consolidation to safe-test the powering and protection systems. It was also concluded that the test would involve incremental powering stages and that all systems must be fully (re)commissioned both before and after a CSCM.

Table 4.2: CSCM risk assessment without mitigation.

System/Component	Scenario	Outcome	Consequence	Probability	Risk
Power Converter	System Failure	Power Abort	Moderate	Possible	**Medium**
Protection System	Detection Failure	Irreparable Damage	Catastrophic	Possible	**High**
Busbars (inc. lyres)	Thermal Runaway	Irreparable Damage	Catastrophic	Probable	**High**
Diodes	Breakdown	Irreparable Damage	Catastrophic	Possible	**High**
Diode-leads	Thermal Runaway	Irreparable Damage	Catastrophic	Probable	**High**
Current-leads	Thermal Runaway	Irreparable Damage	Catastrophic	Improbable	**Medium**

Table 4.3: CSCM risk assessment with mitigation proposals.

System/Component	Scenario	Risk	Mitigation	New Risk
Power Converter	System Failure	**Medium**	System (Re)commissioning	**Low**
Protection System	Detection Failure	**High**	System (Re)commissioning	**Low**
Busbars (inc. lyres)	Thermal Runaway	**High**	Incremental Powering	**Medium**
Diodes	Breakdown	**High**	Incremental Powering	**Low**
Diode-leads	Thermal Runaway	**High**	Incremental Powering	**Medium**
Current-leads	Thermal Runaway	**Medium**	Incremental Powering	**Low**

4.3.2 Powering Profiles

In developing the powering procedure it was important to distinguish between the aim of the full qualification test and the aim of the type test. This section details the developed powering profile for both test types.

Type Test Powering Procedure

The aim of the type test was to qualify the stability of the PC as well as the reliability of the CSCM protection system. In view of the full qualification test, it was important to determine the PC's ramp-rate capabilities as well as its ability to perform a full programmed decay with a time constant equivalent to that of normal operation. The protection system's reliability would be tested by ramping up the current until a thermal runaway was induced for the purposes of detection. With machine protection being a priority, it was decided also that long current plateaus be included, particularly for current levels exceeding that of 4 TeV operation, 4.85 kA. This would allow for a gradual increase in circuit temperature due to Joule heating, only inducing thermal runaways at the lowest possible current. Furthermore, due to the low inductance, fast ramp-rates were possible, allowing for a constant ramp-time of 12 s to be set and ramp-rates of up to 900 A/s. An example powering cycle as well as the full devised test sequence are shown in Figure 4.6 and Table 4.4, respectively.

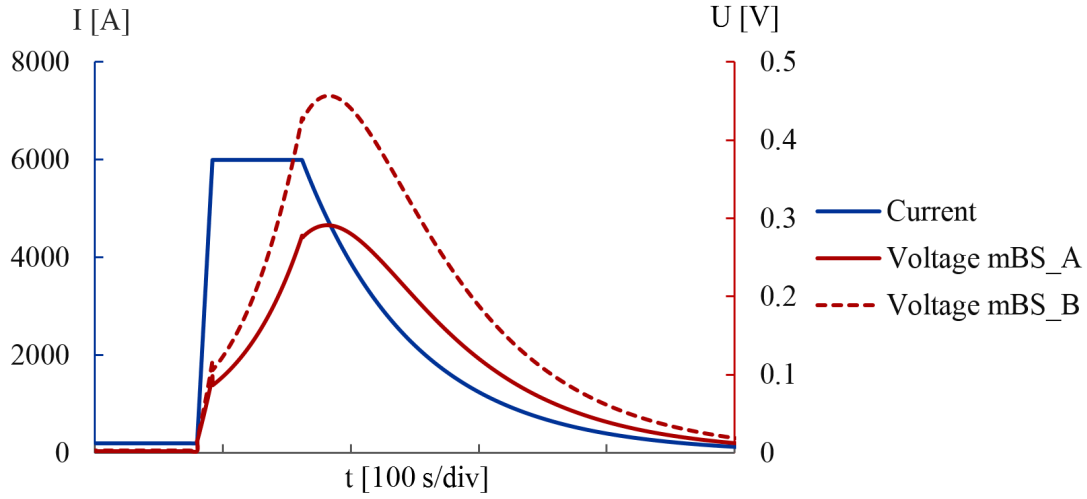

Figure 4.6: Example of a typical CSCM type test powering cycle. Also shown are the corresponding simulated voltages as would be measured by the installed protection system board types mBS_A (busbar) and mBS_B (busbar and diode-leads).

Table 4.4: Planned powering profiles for the CSCM type test.

Peak Current [A]	Ramp Rate [A/s]	Plateau Duration [s]	Quench Integral [10^6A^{2}s]
1000	66.66	40	95
2000	150	40	365
4000	316.66	20	1060
4000	316.66	70	1933
6000	483.33	24	2675
6000	483.33	46	3470
6000	483.33	70	4335
6000	483.33	94	5200
8000	650	42	5904
8000	650	54	6672
12000	983.33	2	7507

Full Qualification Powering Procedure

The aim of a full qualification test is to prove that a circuit can be subject to the full decay from 6.5 TeV equivalent current levels, ultimately mimicking the current decay and accumulated quench integral as would be present in a quench event during normal operation. The quench integral,

$$\Gamma_{\text{quench}} \approx I_0^2 \int_{t_0}^{t_{\text{quench}}} \mathrm{d}t + I_0^2 \int_{t_{\text{quench}}}^{\infty} \mathrm{e}^{-\frac{2t}{\tau}} \mathrm{d}t, \tag{4.1}$$

is the time-integral of the current squared from the onset of a quench [91], where I_0 is the current level at the onset of the quench, t and t_0 are time and t_{quench} is the time of initial quench detection, respectively. τ is the current decay constant – derived from the circuit's total inductance, L_{tot} and resistance, R_{tot}, following

$$\tau = \frac{L_{\text{tot}}}{R_{\text{tot}}}. \tag{4.2}$$

Related to dissipated power, the quench integral is a quantity commonly used in supercon-

ducting magnet systems design to calculate adiabatic estimates for hot-spot temperatures. For a CSCM, the quench integral is used as a reference value with which to design the final current profile. To qualify the circuit for 6.5 TeV operation, it follows that

$$\Gamma_{\mathrm{CSCM}} = \int_0^{\infty} I^2 \mathrm{d}t \equiv \Gamma_{\mathrm{quench}}(I_0 = I_{\mathrm{nom}} + I_{\mathrm{delta}} = 11.1 \text{ kA}). \tag{4.3}$$

where I_0 is the sum of the 6.5 TeV operation current, I_{nom}, 10.98 kA, and an additional 120 A safety margin, I_{delta}.

For simplicity, the resultant decay constant, current plateau length and ramp-time were fixed across all tests at 92 s (slightly faster than nominal to account for diode-opening, ramp-up and plateau phases), 2 s and 8 s, respectively – final numbers were established taking into account type test performance. An example of a full qualification powering cycle as well as the full devised test sequence are shown in Figure 4.7 and Table 4.5, respectively. Note that to monitor any variation in circuit resistances, resistance measurements were carried out at 400 A before and after each test.

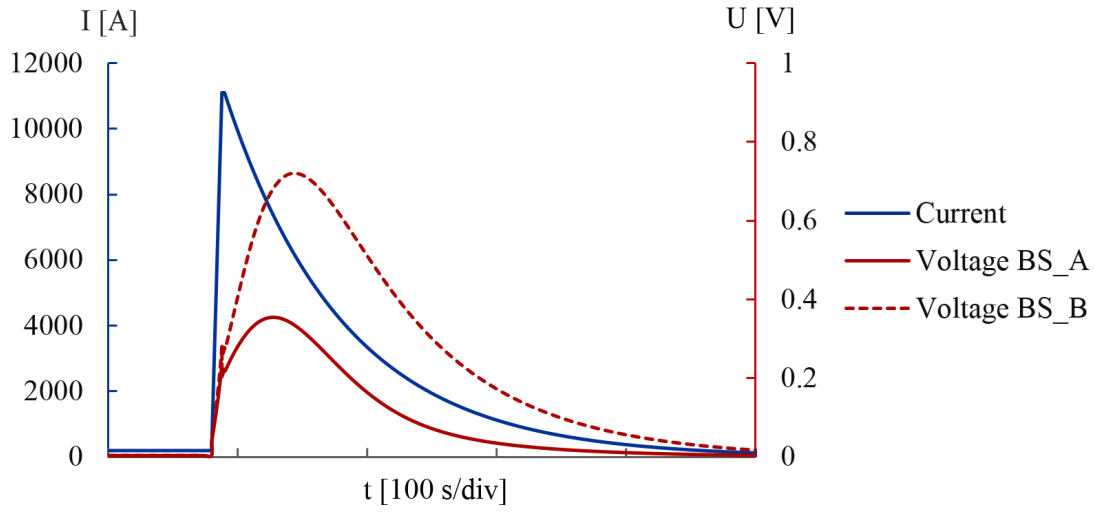

Figure 4.7: Example of a typical full qualification CSCM powering cycle. Also shown are the corresponding simulated voltages as would be measured by the installed protection system board types mBS_A (busbar) and mBS_B (busbar and diode-leads).

Table 4.5: Planned powering profiles for the full qualification CSCM.

Peak Current [A]	Ramp Rate [A/s]	Plateau Duration [s]	Quench Tax [10^6A^{2}s]
2000	225	2	195
5000	600	2	1261
7000	850	2	2478
8600	1050	2	3743
10000	1225	2	5064
11100	1360	2	6240
11100	1360	2	6240

4.3.3 CSCM Simulator

Protection system thresholds had to be set high enough to prevent false trips, yet low enough to prevent system damage in the event of a thermal runaway. To ensure relative accuracy, thresholds were calculated via simulation.

For simplicity, it was determined that a 0–dimensional solver of the heat equation, determining the voltage and temperature in a typical 41-m-long busbar as well as in a single magnet bypass, would be sufficient.

The numerical solver was constructed in accordance with the following:

$$U = U_{\text{res}} + U_{\text{ind}}, \tag{4.4}$$

where

$$U_{\text{res}} = I \cdot \left(R'_{\text{Cu}} l + R_{\text{cont}}\right) \tag{4.5}$$

and

$$U_{\text{ind}} = L \cdot \frac{\mathrm{d}I}{\mathrm{d}t}. \tag{4.6}$$

The total voltage, U, is determined by the sum of the resistive and inductive voltages, U_{res} and U_{ind}, respectively. I is the circuit current, L is the circuit inductance with all diodes in a conducting state, $R'_{\text{Cu}} l$ is the voltage across a specific length of busbar and R_{cont} is

the sum of all encompassed contact resistances.

The accuracy of R'_{Cu} relies heavily on the empirical relationship between the Cu stabiliser's resistivity, ρ_{Cu}, and its Residual Resistivity Ratio (RRR). Where

$$R'_{\mathrm{Cu}} = \frac{\rho_{\mathrm{Cu}}}{A_{\mathrm{Cu}}} \tag{4.7}$$

and

$$\rho_{\mathrm{Cu}} = \frac{c_0}{\left(\mathrm{RRR} + \frac{1}{\frac{c_1}{T_{\mathrm{Cu}}^5} + \frac{c_2}{T_{\mathrm{Cu}}^3} + \frac{c_3}{T_{\mathrm{Cu}}}}\right)} + \alpha B, \tag{4.8}$$

where

$$\mathrm{RRR} = \frac{\rho_{\mathrm{Cu,\ \sim 293K}}}{\rho_{\mathrm{Cu,\ \sim 0K}}}. \tag{4.9}$$

The RRR is defined as the ratio between a metal's resistivity at room temperature and just above absolute zero (or as close as possible) [92]. Note that for RRR values in the expected range of 200-400 [93], the RRR must be calculated in accordance with Equation 4.8, as the Cu resistivity-temperature relation becomes non-linear above ∼15 K. Equation 4.8 shows that ρ_{Cu} can be expressed as a function of empirical constants c_0, c_1, c_2 and c_3, the busbar temperature, T_{Cu}, magneto-resistivity, α, and magnetic flux density, B. B is calculated from the product of the current in the magnet and the corresponding dipole magnetic transfer function [94], F,

$$B = I \cdot F. \tag{4.10}$$

The busbar and helium temperature, T_{Cu} and T_{He}, respectively, are calculated by using the explicit method of solving the following 0–D heat balance equations (it is assumed that all Joule heating occurs in the Cu stabiliser), where

$$C_P \cdot \frac{\mathrm{d}T}{\mathrm{d}t} = P_{\mathrm{Cu}} - P_{\mathrm{He}} \tag{4.11}$$

and

$$C_{V,\mathrm{He}} \cdot \frac{\mathrm{d}T_{\mathrm{He}}}{\mathrm{d}t} = P_{\mathrm{He}}, \tag{4.12}$$

where

$$C_P = \left(\frac{\partial H}{\partial T}\right)_p = \sum \rho c_p V \tag{4.13}$$

and

$$C_{V,\mathrm{He}} = \left(\frac{\partial H_{\mathrm{He}}}{\partial T}\right)_v = m_{\mathrm{He}} c_{v,\mathrm{He}}. \tag{4.14}$$

P_{Cu} and P_{He} are the power dissipated in copper stabiliser and helium bath, respectively. H is the system enthalpy and ρ, c_p, V and m are the density, specific heat capacity, volume and mass of a given material, respectively. C_P is the heat capacity of the entire busbar (Cu stabiliser, Nb–Ti strand and Sn–Ag solder included) for a constant pressure. $C_{V,\mathrm{He}}$ is the heat capacity of the helium bath for a constant volume. The power dissipated in the copper is calculated following

$$P_{\mathrm{Cu}} = I \cdot U_{\mathrm{res}} \tag{4.15}$$

and the power dissipated in the helium is

$$P_{\mathrm{He}} = \kappa \cdot (T_{\mathrm{Cu}} - T_{\mathrm{He}}). \tag{4.16}$$

The heat transfer coefficient, κ, is an unspecified parameter representing the thermal conductivity from the copper to the helium. This parameter is tuned following each powering cycle, improving simulation accuracy for a given current level.

A final note is that, to manage simulation convergence, manually adaptive time-stepping was utilised, with dt eventually being set to 1.8, 0.03, 0.379 and 0.45 s for the diode opening, ramp, plateau and decay times, respectively.

4.3.4 Calculation and Setting of Protection Thresholds

To minimise risk, it was decided that unique thresholds be calculated for all individual busbar segments. It was concluded that this would best be done by first calculating the RRR of individual segments during CSCM system commissioning and subsequently establishing the maximum U and dU/dt as a function of RRR. Busbar RRR values were calculated in accordance with Equation 4.8, with segment specific temperatures being determined from temperature probes installed at various locations along the Arc. Calculated RRR values averaged 260, lying between 200-320 – in-line with previous measurement attempts [93]. Thresholds were calculated prior to each run, adding operational margins to simulated maximum values for U and dU/dt as well as taking into account variations in temperature and cooling.

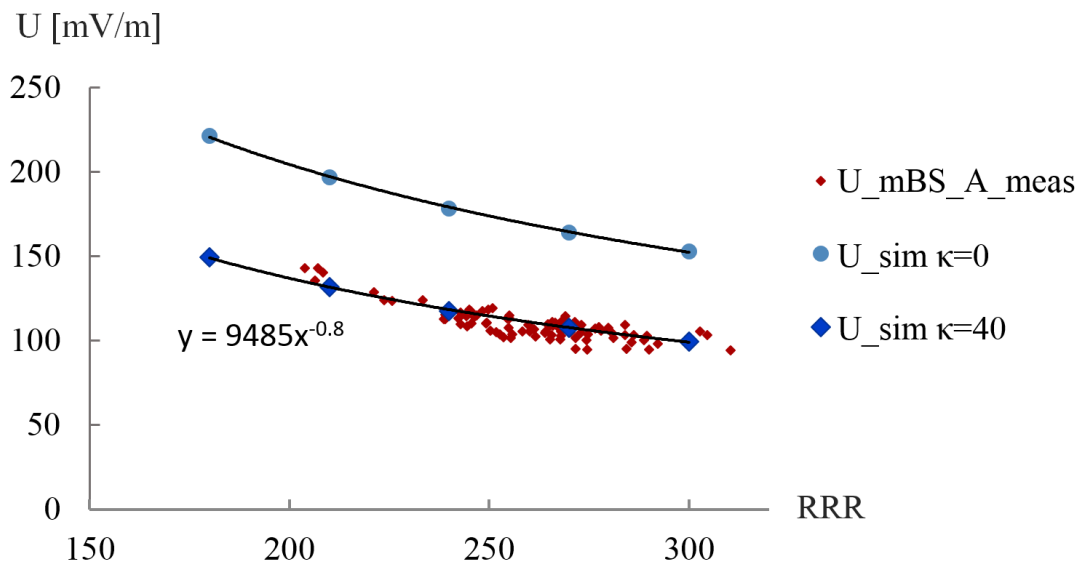

Figure 4.8: Plot of maximum individual busbar voltages as measured during a 6 kA CSCM cycle. Also shown are the corresponding simulated values, both with and without tuning the heat transfer coefficient, κ.

Using the simulator, U and $\mathrm{d}U/\mathrm{d}t$ were calculated for each powering cycle with various RRR values, 180-300, using the mean temperature at the time. Simulated values were then plotted against measured values from the previous cycle and the heat transfer coefficient, κ, was tuned accordingly, see Figure 4.8. Fit functions were determined and subsequently applied to individual busbar segments, taking into account their respective lengths and RRR, such that

$$U(\mathrm{RRR}, l) = a_0 \cdot \mathrm{RRR}^{b_0} \cdot \frac{l}{l_{\mathrm{mode}}}, \tag{4.17}$$

where a_0 and b_0 are the fit function constants and l and l_{mode} are the specific and mode busbar lengths, respectively. Note that the calculation of $\mathrm{d}U/\mathrm{d}t$ (not shown) followed the same principle.

Following simulation of segment specific voltages, additional margins were added to reduce the probability of a false trip. Firstly, taking into account the ratio of measured to tuned values, such that

$$U_{\mathrm{sim+err}} = U_{\mathrm{sim}} \cdot \frac{U_{\mathrm{meas,prev}}}{U_{\mathrm{sim,prev}}}. \tag{4.18}$$

Secondly, to account for excess interconnect resistances as well as simulation errors and potential unknowns, an additional 10% margin (100 mV minimum) was added. The maximum additional margins were set to 500 mV and 50 mV/s for U and $\mathrm{d}U/\mathrm{d}t$, respectively.

Protection System Sunglasses

As a CSCM test utilises fast current ramp-rates, large inductive voltages occur. The $\mathrm{d}U/\mathrm{d}t$ is therefore significantly higher during the current ramp than throughout the rest of the powering cycle. As such, setting $\mathrm{d}U/\mathrm{d}t$ thresholds high enough to prevent a trip during the current ramp would result in the thresholds being too high to safely detect a thermal runaway during the current plateau or decay. The solution to this problem was to implement threshold 'sunglasses', similar to a mask or filter. Threshold sunglasses allowed the $\mathrm{d}U/\mathrm{d}t$ threshold to initially be elevated by a factor at the onset of the current ramp, subsequently reducing to the nominal threshold after a defined time. This would, however, mean that any thermal runaway induced during the ramp-up may only be protected by the corresponding U threshold. Figure 4.9 illustrates the process as implemented. For the full qualification tests, the elevation factor was set to 4 and the trigger voltage, U_{trigger} was set to 40 mV.

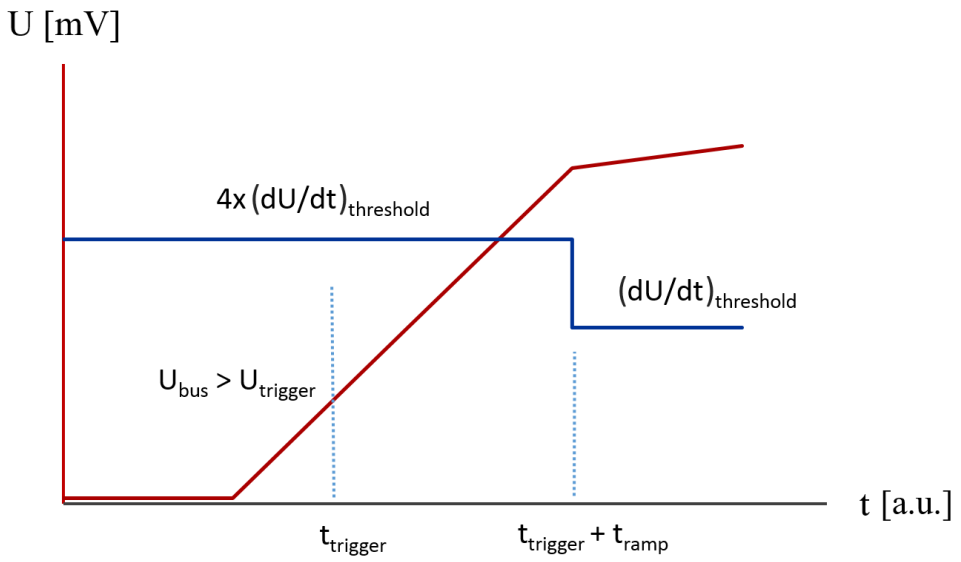

Figure 4.9: Illustration of the implemented PS sunglasses, allowing one to temporarily raise the $\mathrm{d}U/\mathrm{d}t$ *threshold during the current ramp to prevent false trips as a result of high inductive voltages.*

4.4 Sector 23 CSCM Type Test Results

To determine the safety of the proposed full qualification procedure, a type test was carried out to qualify that all involved systems can successfully operate under the atypical

operating conditions and that the calculated U and $\mathrm{d}U/\mathrm{d}t$ thresholds are sufficiently accurate to prevent excessive temperatures during a thermal runaway. The type test was carried out prior to interconnect consolidation and with discontinuities remaining in the system, thermal runaways were not only expected, but were deliberately induced. Such a high risk test was approved as any system damage that may occur could be repaired during LS 1 in the shadow of the consolidation campaign. This section details the results and conclusions drawn from the type test.

4.4.1 Powering Cycles

The CSCM type test was carried out on the Sector 23 main dipole circuit. Prior to carrying out the main powering cycles, all involved systems were commissioned to ensure correct operation. The Arc cryostat's temperature and pressure were regulated to specification one week prior to testing and all mBS boards were installed, each being individually tested for detection capabilities. Furthermore, the PCs current and voltage stability, ramp-rate and programmed exponential decay were tested at low current levels. However, during PC tests, multiple trips occurred at low current levels, <100 A, when attempting to ramp up. This was found to be an issue with opening all 154 diodes simultaneously as several were still sufficiently cold that the conducting voltage was 2.5-3 V. The solution was to initially start with a reduced ramp rate of 10 A/s and to implement a plateau at 200 A. This allowed all diodes to become conducting over a period of ∼60-80 s and, due to the varying characteristics of individual diodes, opening times of individual diodes were spread. The optimal length of the 200 A plateau was found to be ∼120 s and such a plateau was added to the beginning of each powering cycle. Subsequent to all diodes opening, the desired powering cycles could be carried out without issue.

Initial commissioning tests were carried out and the first few powering cycles were performed without problems. Upon reaching 6 kA there were, however, no signs of a runaway occurring and it was decided to increase the peak current to 7 kA. The results were positive and the first thermal runaway, see Figure 4.10, was induced in one of the cell 22 busbar interconnects, ∼24 s into the current plateau. The runaway was successfully detected and protected against, initiating a power abort as the voltage exceeded the defined threshold. Each successive cycle induced runaways at various current levels, with varying characteristics. All carried out tests are listed in Table 4.6.

The characteristics of the first thermal runaway were of particular interest as it was immediately apparent that the runaway was not exponential but, in fact, begins to slow-down. Such a slow-down phenomenon was not taken into account in the calculation of the thresholds and, as a result, it took significantly longer than expected for the protection

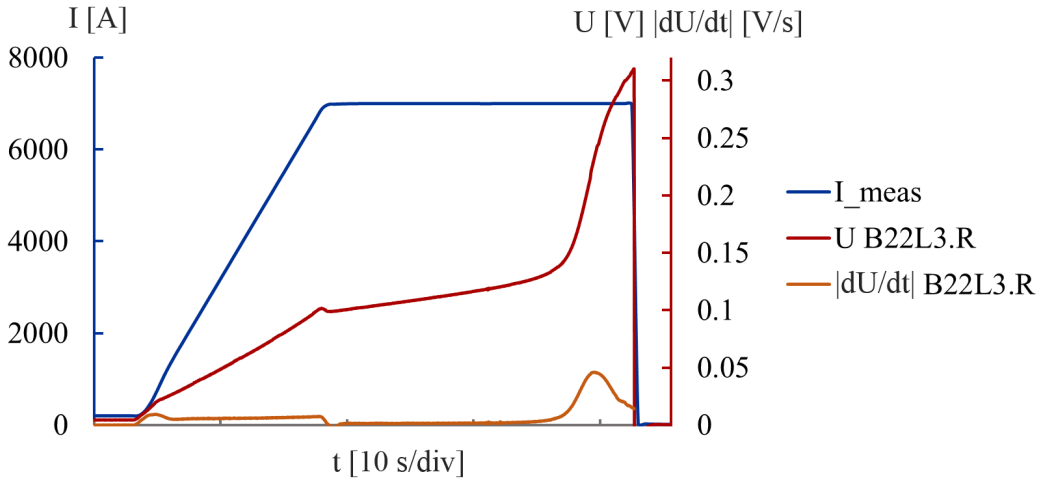

Figure 4.10: Plot of the 1st thermal runaway induced during the Sector 23 CSCM type test. Peak current 7 kA. Runaway detected ∼24 s into the plateau.

system thresholds to be exceeded. Furthermore, such a slow-down effect rendered the set dU/dt thresholds near-obsolete. Following analysis, it was decided to significantly reduce threshold margins such that the maximum additional margins were 200 mV, down from 500 mV, and 30 mV/s, down from 50 mV/s for U and dU/dt, respectively.

Figure 4.11 shows the 8 kA powering cycle, with an induced thermal runaway being detected and protected against ∼6.6 s into the current plateau. Reducing the maximum threshold margins did result in detection prior to the witnessed slow-down phenomenon.

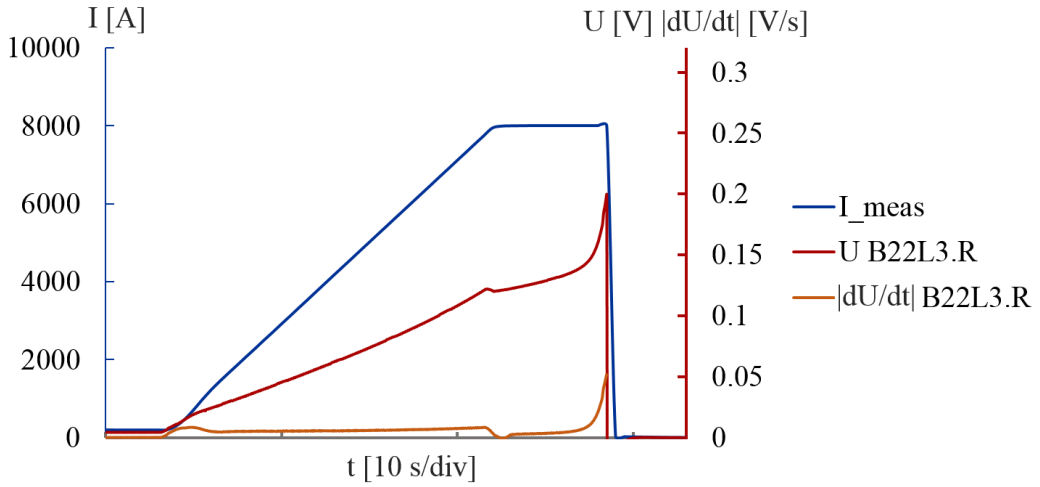

Figure 4.11: 2nd thermal runaway induced during the Sector 23 CSCM type test. Peak current 8 kA. Runaway detected ∼6.6 s into the plateau.

Subsequent to qualifying PC stability and thermal runaway detection, an *ad-hoc* test to determine the PC's maximum feasible ramp-rate and to fully qualify the protection system

Table 4.6: All tests carried out during the CSCM type test.

Peak Current [A]	Ramp Rate [A/s]	Nominal Plateau [s]	Runaway Detection [s]	Quench Integral [10^6A^{2}s]
1000	66.66	40	n/a	94
2000	150	40	n/a	360
4000	316.66	70	n/a	1900
6000	483.33	46	n/a	3400
7000	566.66	40	24	1400
8000	650	14	6.6	710
7000	566.66	40	31.5	1770
9000	1466.66	10	4.25	710
6000	483.33	90	76	2900

sunglasses was carried out, ramping at 1466 A/s to 9 kA. The test was successful following the first attempt, with the dU/dt threshold being successfully elevated throughout the current ramp and a power abort only being initiated by the detection of a thermal runaway during the current plateau. The fast ramp rate, however, resulted in significant current overshoot immediately tripping the PC's internal thresholds and initiating a power abort. Further tests concluded that a value no greater than ∼1400 A/s should be used.

4.4.2 Simulator Accuracy

To qualify the accuracy of the CSCM simulator, simulated and measured functions for U and dU/dt were compared for both BS_A and BS_B boards. Figures 4.12 shows comparative plots for a typical busbar segment during a 6 kA power cycle. Results show a good agreement with BS_A measurements but a large discrepancy for BS_B values (includes diode-leads contacts). This resulted in the thresholds for all BS_B boards being increased. The discrepancy of ∼22 $\mu\Omega$ is likely due to a previously observed phenomenon in which diode-lead contact resistances have a tendency to increase at current levels of 5-10 kA [95]. This phenomenon was discovered during cold-diode specific experiments carried out in CERN's SM18 test facility in 2011 where diode to heat-sink contacts resistances were shown to increase by up to ∼30 $\mu\Omega$, see Figure 4.13. Excess resistances were, however, shown to decrease as current levels exceeded 10 kA, *i.e.* an 11.08 kA full qualification test may stabilise or even reverse this effect. Neither of these effects are fully understood and further investigation is underway [96].

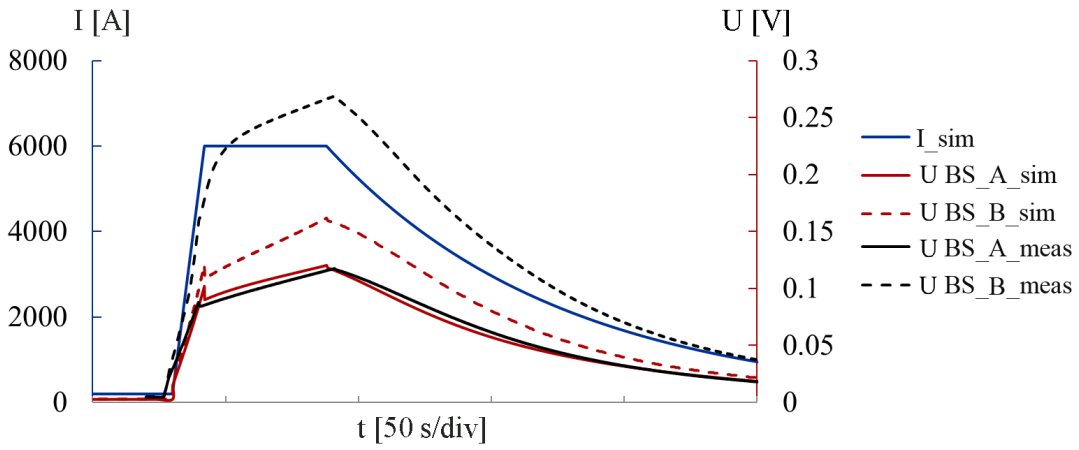

Figure 4.12: Comparative plots of the U simulated and measured values during a successful 6 kA CSCM cycle for both BS_A and BS_B protection system board types.

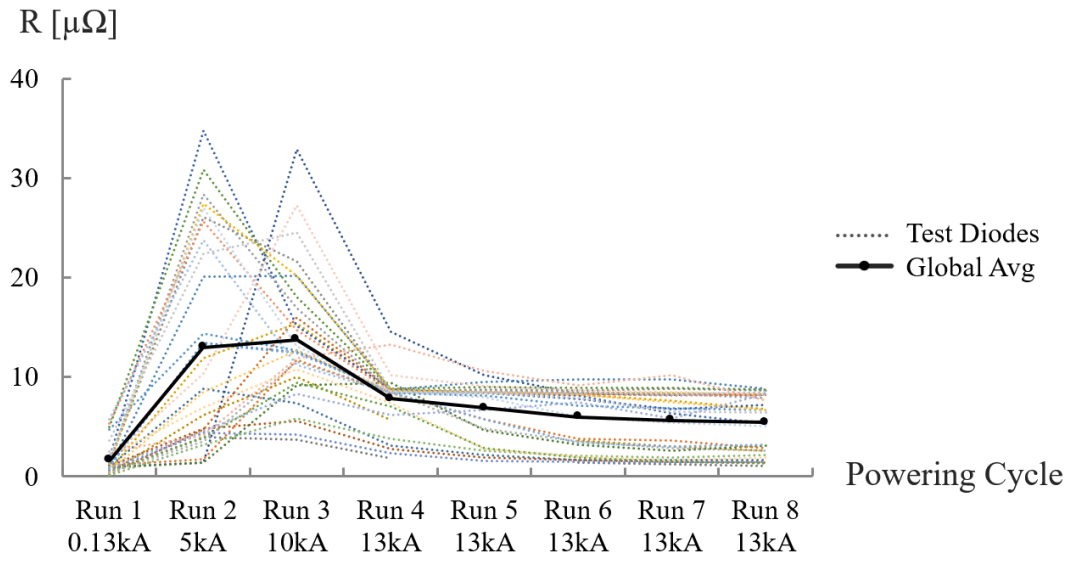

Figure 4.13: Plot of contact resistances as measured across 30 spare diodes during an SM18 test campaign. Resistances show a trend of increasing when submitted to current levels of 5-10 kA, but then decreasing to a plateau as current levels exceed 10 kA. Test data courtesy of G. Willering.

Figure 4.14 and 4.15 show a direct comparison of all simulated and measured values for U and the maximum U and $\mathrm{d}U/\mathrm{d}t$ simulation error, respectively. Overall, simulations are in good agreement with measured values and tuning for errors allow simulation accuracy to either remain consistent or improve with each current cycle. Overall, it was concluded that the CSCM simulator was of sufficient accuracy for at least the BS_A voltages and that it may be used for the calculation of thresholds for full qualification tests.

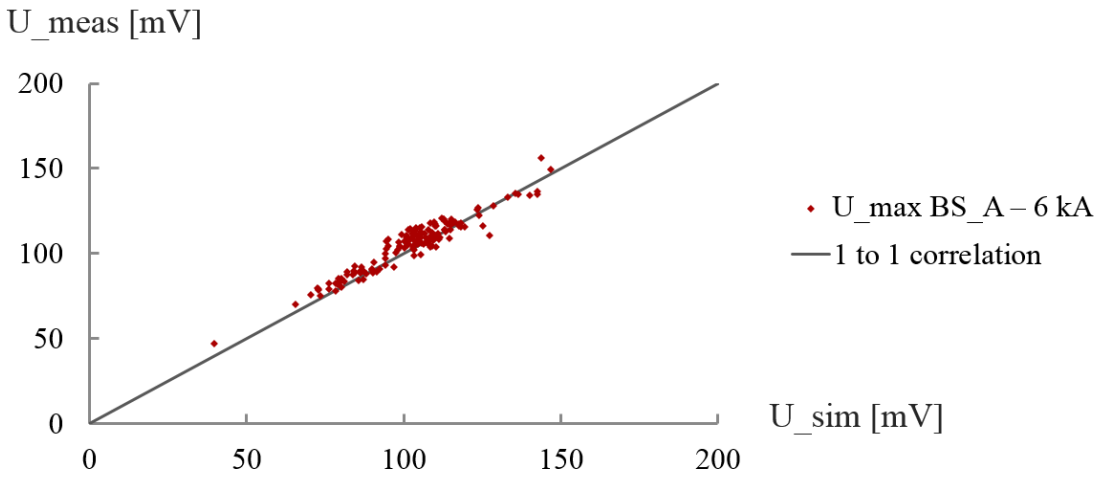

Figure 4.14: Comparative plot of measured and simulated values for $U_{\max}$ for all busbar segments. Results are in good agreement with each other.

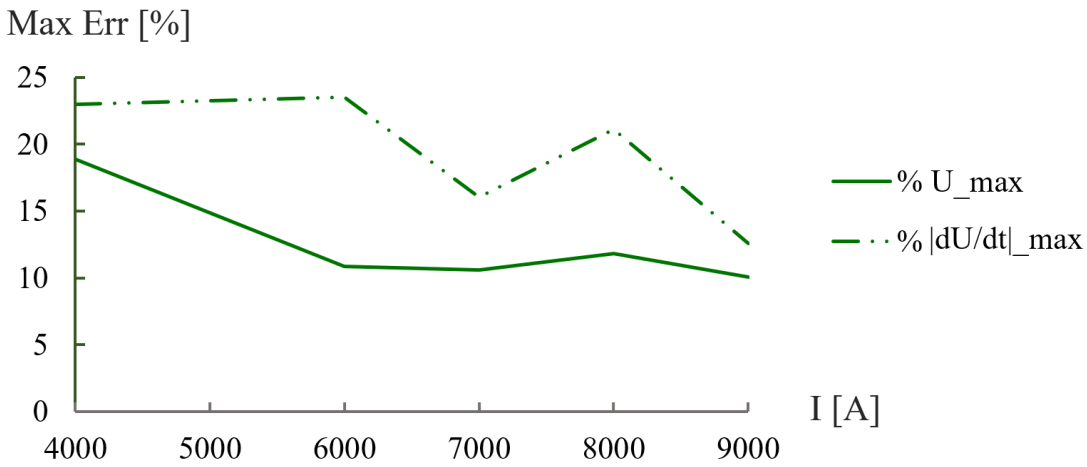

Figure 4.15: Maximum U and $\mathrm{d}U/\mathrm{d}t$ simulation errors for each current cycle. Errors remain consistent and or decrease from one cycle to the next as a result of simulation tuning.

4.4.3 Simulation of Thermal Runaways

To further understanding of the induced thermal runaways, in particular, the observed slow-down phenomenon (discussed in Section 4.4.1), detailed simulations were carried out

using CERN's own QP3 software [97]. QP3 is a 1–D numerical solver specifically designed to simulate longitudinal quench propagation (heat diffusion along the current flow axis). QP3 follows the same equations used in the CSCM solver, but has the additional capacity for longitudinal discretisation. This allows for the simulation of an interconnect discontinuity by means of specifying a localised highly resistive section.

QP3 was configured to simulate a powering cycle across a typical 41-m-long busbar segment, with an average RRR of 249 (as was calculated for the busbar segment in which the thermal runaways were occurring), corresponding material properties for Cu, Nb-Ti and Sn-Ag and a discretisation of 1 mm. Furthermore, a section of variable length and an arbitrarily reduced value for the RRR of 50, representing a interconnect discontinuity was included. Using QP3, by variation of the defect size, it was possible to not only simulate all induced thermal runaways with relative accuracy but also the observed slow-down phenomenon. Table 4.7 lists the defect parameters used to model each runaway as well as the corresponding temperature hot-spot. As shown, the defect resistance used to simulate each thermal runaway was approximately 33 $\mu\Omega$.

Table 4.7: Defect characteristics used in modelling the B22L3.1 thermal runaway.

Peak Current [A]	Defect RRR [-]	Defect Length [mm]	Defect Resistance [$\mu\Omega$@300 K]	Hot-Spot Temperature [K]
7000	50	25.05	32.6	560
8000	50	25.8	33.5	185
7000	50	25.2	32.8	340
9000	50	25.8	33.5	142
6000	50	21.75	28.5	440

Figure 4.16 shows a comparison between the measured data and simulation. In simulating the slow-down phenomenon, it was assumed that a plausible scenario could be that remnant solder was isolated in the interconnect gaps – depicted in Figure 4.17. As such, to match measured data, the slow-down effect was modelled by implementing a temperature dependant parallel resistance of 50 $\mu\Omega$ (at 20 K) across the discontinuity, which begins to share the current as the defect temperature and resistance increases.

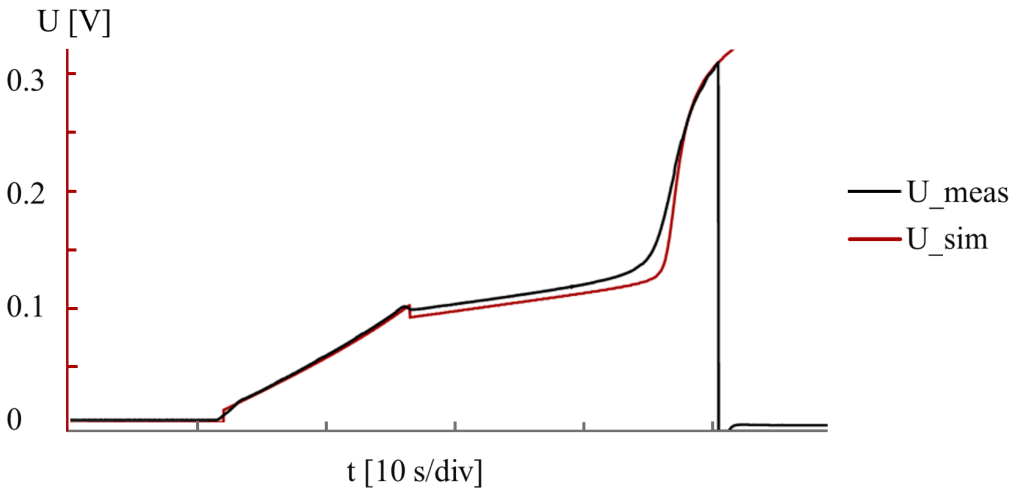

Figure 4.16: Comparative plot of the simulated and measured thermal runaway induced during the 1st 7 kA cycle. The slow-down phenomena was modelled by means of a parallel resistance across the discontinuity.

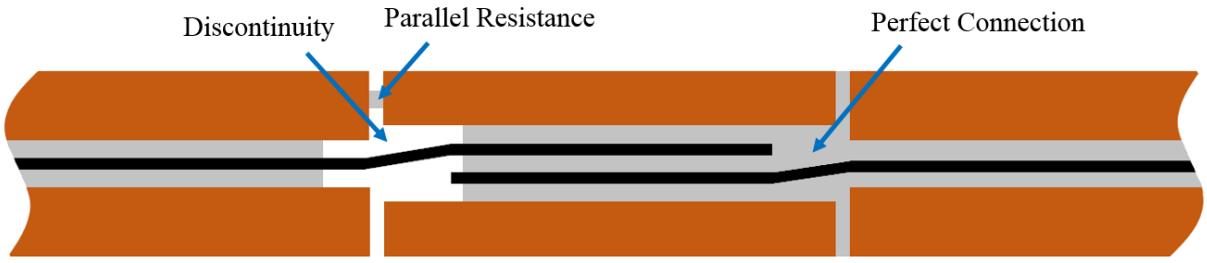

Figure 4.17: Illustration of a typical busbar interconnect, depicting a good and bad connection as well as a plausible location for remnant solder creating a parallel resistance.

Simulation results suggested a local hot-spot temperature of ∼560 K during the first thermal runaway – far exceeding the melting point of solder, ∼456 K. If the slow-down effect was stronger and the voltage threshold was not exceeded until much later, electrical arcing and subsequent structural damage would almost certainly have been the result. This is a clear example of the importance of accurate threshold calculation and conservative detection margins. All subsequent tests were carried out with significantly reduced thresholds margins, leading to much lower hot-spot temperatures.

4.4.4 Type Test Summary

The constructed 0–D CSCM simulator allowed for relatively accurate simulation of the busbar voltages across a range of peak current levels. It was, however, discovered that, due to erratic diode-lead contact resistances, accurate simulation of magnet bypass voltages was not possible and a constant value for all rather than segment specific margins were added to all BS_B thresholds. Furthermore, in carrying out the type test, an unpredicted phenomenon was observed, in which a thermal runaway would begin to slow down. As a result, both U and $\mathrm{d}U/\mathrm{d}t$ threshold margins were significantly reduced for subsequent tests to guarantee safe operation. The phenomenon was successfully modelled using QP3 software by simulating a parallel resistance across an interconnect discontinuity.

In summary, all powering and protection systems were qualified to specification. The PC was proven capable of supplying sufficient voltage, operating with ramp-rates exceeding $\sim$1 kA/s and successfully performing the programmed exponential current decay. The installed protection systems successfully detected and protected against induced thermal runaways without measurable damage. It was, therefore, concluded that a full qualification CSCM test, as proposed, would pose minimal system risk and that an LHC-wide campaign could and should be carried out post-consolidation.

4.5 LHC-Wide Full Qualification Test Results

Following the successes of the CSCM type test, full qualification tests were carried out post-consolidation on all main dipole circuits. This section details the results and conclusions of the test campaign.

4.5.1 Powering Cycles

Similar to the type test, each circuit's pressure and temperature were regulated and individual systems were commissioned prior to carrying out the CSCM tests. As discussed in Section 4.2.3, the protection system was now commissioned using normal operation BS boards in a CSCM mode rather than the CSCM specific mBS boards.

Following the full qualification procedure (Section 4.3.2), each circuit was subject to a series of powering cycles with a programmed exponential current decay, mimicking that of an equivalent quench event. Each successive step in the series is an incremental increase in peak current. In contrast to the type test, all circuits exceeded 7 kA without fault and were eventually subject to the full programmed current decay from 11.1 kA. Testament to the effectiveness of LS 1 consolidations, no thermal runaways were detected throughout the

entire campaign of ∼150 tests. As an example, Figure 4.18 shows all successful powering cycles carried out on the Sector 23 circuit during the full qualification CSCM – protection system commissioning tests and all 400 A resistance measurements are included.

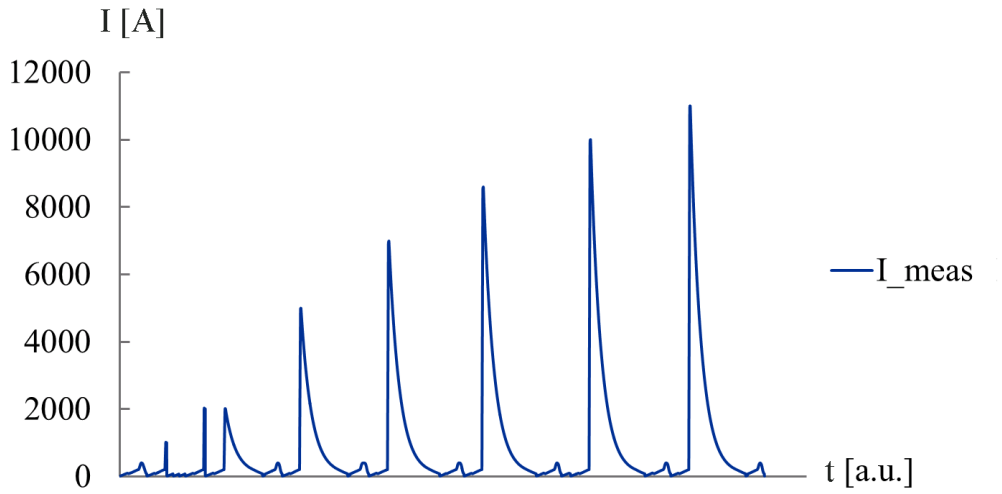

Figure 4.18: Plot of all carried out powering cycles used to fully qualify the Sector 23 circuit to 11.08 kA. Plot shows all main current cycles, protection system commissioning tests and intermittent 400 A resistance measurements.

RRR Measurements

For accurate calculation of protection system thresholds, individual busbar segment resistances were measured and RRR values were calculated in accordance with Equation 4.8. Figure 4.19 shows the statistical distribution of RRR values measured per sector. Values are in-line with manufacturing specifications and previous tests and predictions [93].

4.5.2 Analysis of 400 A Resistance Measurements

Due to the erratic diode characteristics observed during the type test, 400 A resistance measurements were carried out before and after each powering cycle to monitor busbar and magnet bypass resistance trends and outliers. This section details the results of these measurements.

Busbar Resistances

Figure 4.20 shows the global busbar resistance averages as measured by all BS_A boards (busbars only). In contrast to the type test, all resistances remained relatively low and consistent throughout, averaging 16 $\mu\Omega$. Results provide conclusive evidence that all interconnects were properly consolidated and all lyres are intact. The busbar's Cu stabiliser

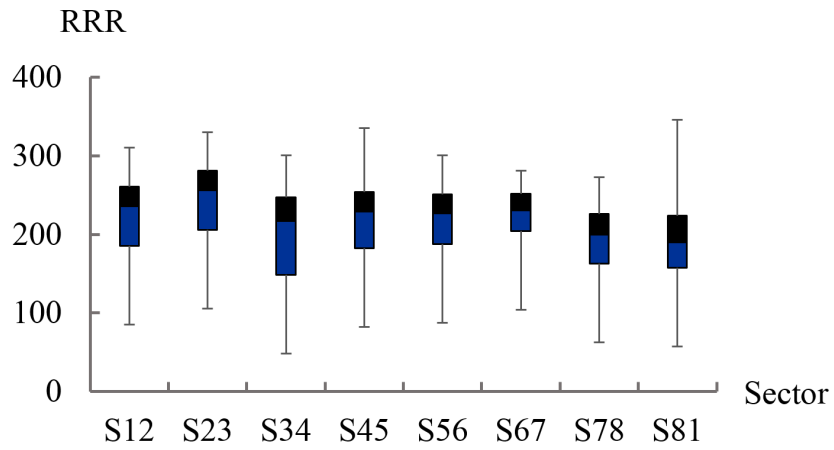

Figure 4.19: Boxplot of the per sector statistical distribution of busbar RRR values. Black and blue boxes represent the upper and lower interquartile ranges, respectively. Lines show the 5th and 95th percentiles. All sectors are in-line with the manufacturing specifications and previous tests and predictions.

can therefore be concluded capable of withstanding 6.5 TeV equivalent current levels of at least 11.1 kA.

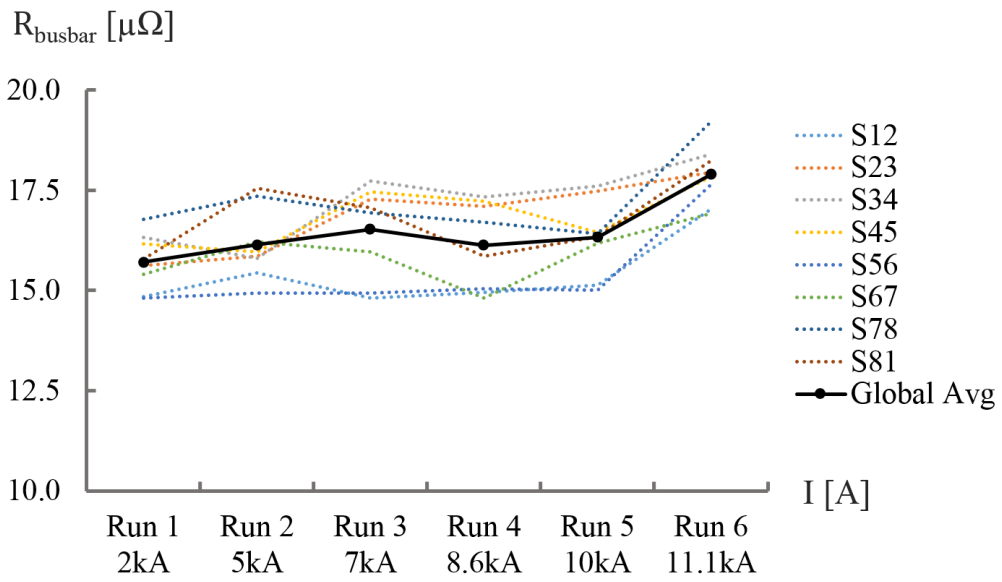

Figure 4.20: Global busbar resistance averages for all sectors as measured following CSCM powering cycles. All resistances remained relatively low and consistent throughout.

Magnet Bypass Resistances

Magnet bypass resistances were derived from the difference in bypass voltages, *i.e.* DS_B - DS_A. Results are shown in Figure 4.21. As predicted, they were notably less consistent than the busbar measurements. Similar to the type test, measured resistances were ob-

served to increase following powering. Furthermore, similar to the SM18 diode experiments (discussed in Section 4.4.2) resistances began to saturate and or decrease as current levels exceeded ∼8-10 kA.

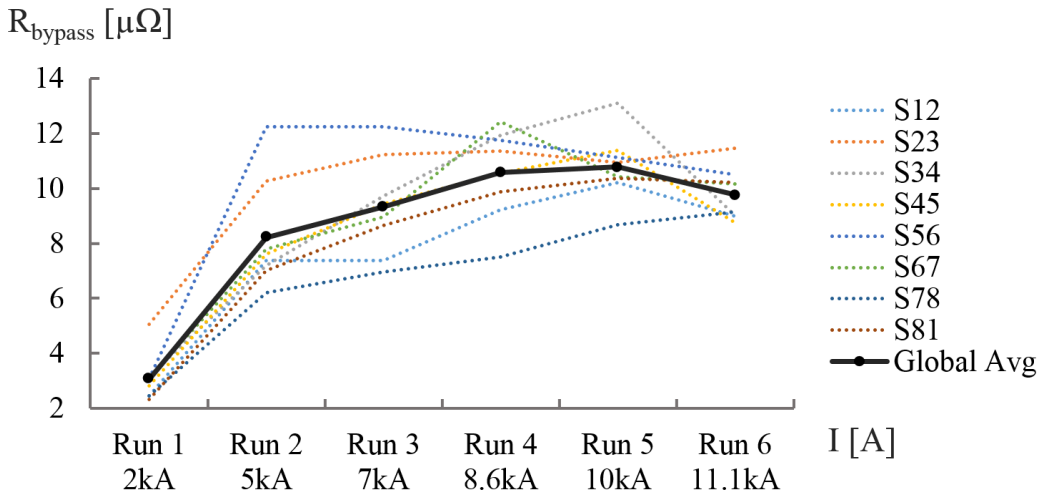

Figure 4.21: Global bypass resistance averages for all sectors as measured following CSCM powering cycles. Note the initial rise, saturation and recovery trend arising as current levels increase.

Table 4.8: Largest recorded resistance increases across a magnet bypass. Recall that the largest witnessed increase in SM18 diode experiments was only 30 μΩ.

Magnet Bypass [-]	Peak Current [A]	Initial Resistance [μΩ]	Peak Resistance [μΩ]	Resistance Increase [μΩ]	PS Trip [-]
B22R3	10000	29.0	384.1	355.1	Yes
B9L5	10000	35.4	155.3	119.9	No
C12L6	11080	14.6	127.5	112.9	No

Table 4.8 details the three largest resistance increases measured from one cycle to the next. As shown, the greatest increase in resistance from one cycle to next was measured at 355.1 μΩ – 10 times the largest increase observed in SM18 experiments. In this circumstance the protection electronics triggered ∼10 s into the programmed decay even without the notable occurrence of a thermal runway. Looking at the difference in measured diode-lead resistances from the both DS and BS measurements, *i.e.* (DS_B-DS_A) - (BS_B-BS_A), it was possible to deduce that the build-up of excess resistance occurred in at least two contact resistances and as a result pose much less risk. Furthermore, if

the excess resistance was to be generated in the diode stack contacts which have no adjacent heat sink, the temperature would have been at least an order of magnitude higher and would have undoubtedly resulted in a trip and or structural damage. Therefore, one can conclude that most, if not all, abnormally high excess resistances were present in the diode-to-heat sink contacts and pose minimal risk.

Due to the presence of such high resistances, it was decided not to proceed with the CSCM as normal, but to attempt to deliberately induce the recovery / resistance-reduction effect seen as current levels exceed 10 kA. However, rather than powering to 10+ kA, it was decided to first to perform a stepped powering cycle with 1 kA increments and relatively long ∼20 s plateau, up to a maximum of 5 kA. This was in attempt to deliberately heat up the highly resistive contacts to see if the recovery effect was in some way related to a critical energy density in the contacts, *i.e.* could be induced via excessive Joule heating. The results were successful, reducing the critical bypass resistance to just below 55 $\mu\Omega$, a 329 $\mu\Omega$ reduction. As a result, the stepped cycle was used twice more during the campaign for circuits in Sectors 45 and 56 to reduce the other contact resistances that exceeded 100 $\mu\Omega$. Note that these resistance increases occurred during otherwise successful powering cycles. Figure 4.22 shows the two largest resistance increases during otherwise successful powering cycles, *i.e.* no protection system trip. The linear trend of the increase throughout the decay does suggest a steady build of an excess resistance layer – hypothesised to be the production of a resistive nickel-oxide layer between the diode casing and its heat sink [96]. Figure 4.23 shows an example of the devised stepped cycle being implemented to deliberately induce the observed recovery phenomenon. As shown, the bypass resistance significantly reduces from ∼150 $\mu\Omega$ to below 50 $\mu\Omega$.

4.5.3 LHC-Wide Full Qualification Summary

In carrying out the full test procedure LHC-wide, all circuits were eventually subject to a full decay from 11.08 kA, fully qualifying all potentially weak components in the circuit and allowing magnet training to be carried out with the lowest achievable risk. The campaign also revealed invaluable information regarding the predictability of bypass contact resistances, including the true extent of potential bypass resistance increases from one powering cycle to the next. In view of the observations, a devised stepped cycle was successful in reducing excessive resistance build up in the diode-to-heat-sink contacts providing a direct and repeatable solution to the issue. Such a powering cycle could even be used in the event of critical resistance increase occurring during a training quench. Overall, the full qualification test campaign was deemed a success and magnet training was allowed to proceed.

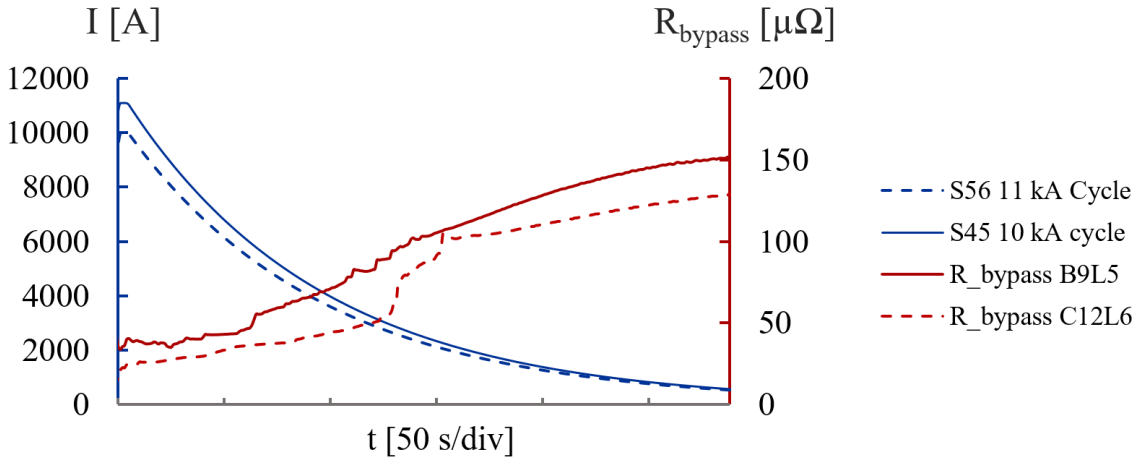

Figure 4.22: Plots of the two largest bypass resistance increases measured throughout otherwise successful powering cycles. Both shown increases exceed 100 μΩ.

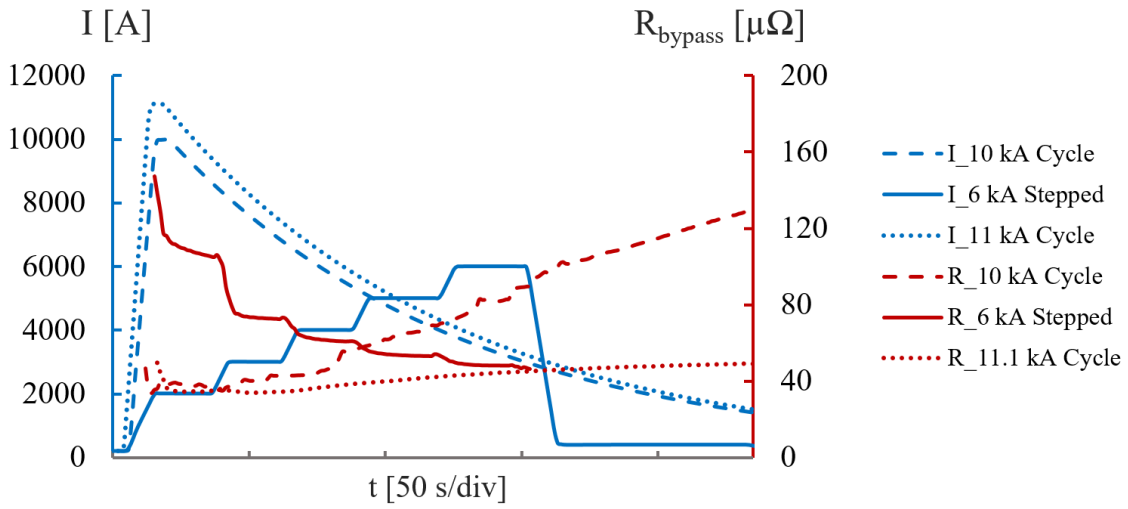

Figure 4.23: Plot of the devised stepped cycle (red) utilised to reduce the relevant bypass resistances by excessive Joule heating. Note that the resistance is decreased from ∼150 μΩ to below 50 μΩ and remains so throughout the following powering cycle. The preceding 10 kA (blue dashed) and subsequent 11 kA (blue dotted) cycles are also shown.

4.6 Magnet Training Outcome Statistics

Following the successes of the CSCM campaign, training of all main dipole magnets was carried out between December 2014 and April 2015. Figures 4.24 shows the training statistics by sector. Recall that predictions would see ∼84 quenches to achieve 6.5 TeV in total. Early in the training campaign it was clear that these predictions were unrealistic, with the first few sectors taking ∼20 quenches to reach 11.08 kA (I_{delta} was reduced to 100 A to minimise the required training). Sector 45 alone took over 50 quenches to fully train. In total, there were 179 training quenches, more than double that of predictions (Section 3.2.1), including a number of detraining quenches occurring intermittently. The discrepancy in predictions compared to empirical statistics is, however, thought to be related to how Firm 3 magnets have developed over the course of Run 1 operation. Taking into account existing quench data, it was expected that Firm 3 magnets would require more training, however, Figure 4.25 shows that ∼80% of all training quenches were for Firm 3 magnets – far greater than any prior estimates. There is also a correlation with the total number of training quenches per sector to the ratio of installed Firm 3 magnets. In total, 158 different magnets required retraining, with 1064 retaining their memory up to 6.5 TeV post-installation and Run 1 operation. Table 4.9 shows that the majority of quenches were on unique magnets with only 9 magnets having to be quenched twice and only one magnet having to be quenched 3 times. To further understand the statistics and phenomena involved as well as to make more accurate predictions for any future training campaigns, *e.g.* to 7 TeV, a specific working group, the 'Quench Behaviour Team', was formed as of July 2015 [98].

Table 4.9: Number of individual magnets that quenched more than once. 9 magnets required more than one training quench. Only one magnet required 3 training quenches.

Manufacturer (No. Magnets)	No. Magnets Quenched Once	No. Magnets Quenched Twice	No. Magnets Quenched 3 times
Firm 1 (400)	5	0	0
Firm 2 (420)	27	0	0
Firm 3 (412)	126	9	1
Total (1232)	158	9	1

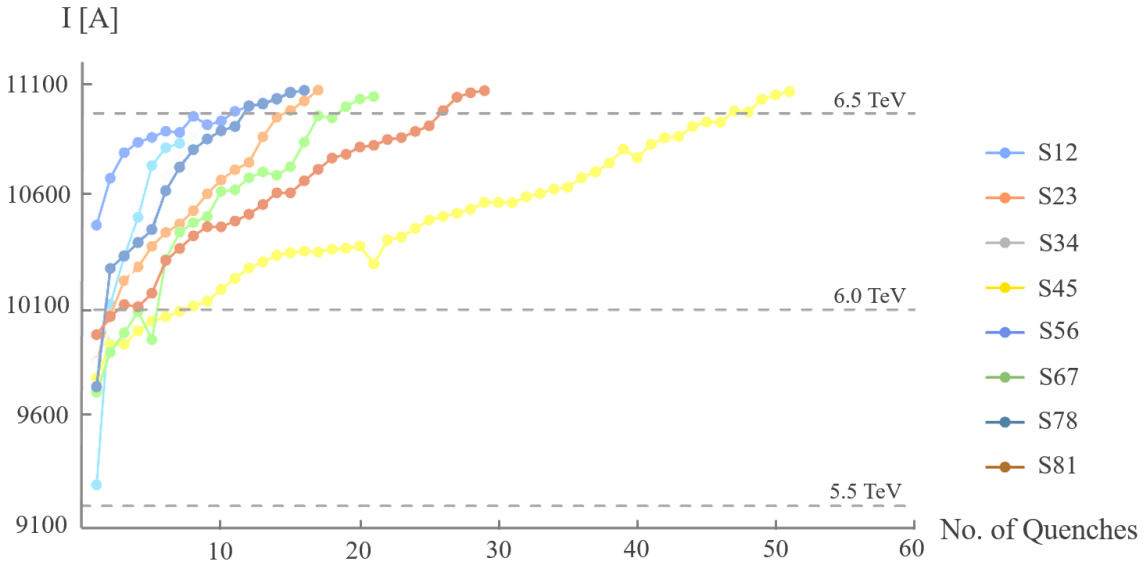

Figure 4.24: Training quench statistics across all sectors. Sector 45 is a notable outlier requiring >50 training quenches. A total of 179 training quenches were required to train all sectors to 11.08 kA.

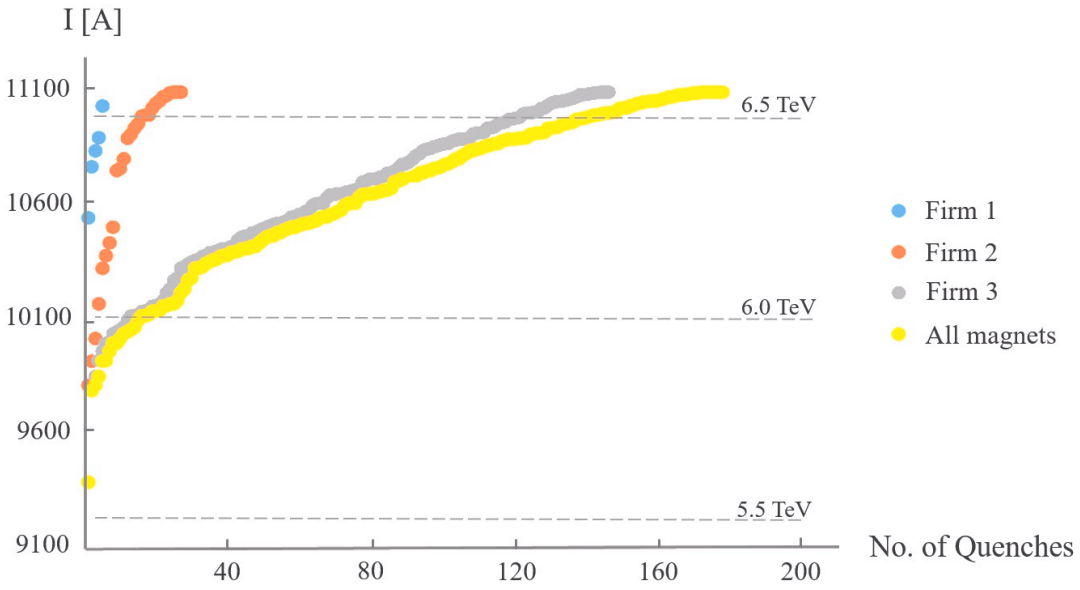

Figure 4.25: Training quench statistics, arranged by magnet manufacturer. It is clear that Firm 3 magnets required significantly more retraining than both Firm 1 and 2.

4.7 Chapter Discussion & Conclusions

Subsequent to the re-machining of all busbar interconnect discontinuities, it was proposed that a new qualification method be developed to ensure no issues remain, especially with magnet training predictions of >84 quenches. The result was the development of the copper stabilizer continuity measurement. The CSCM project saw the development and implementation of a new method of electrical qualification for the stabiliser and passive magnet bypass components of a superconducting magnet string. The aim was to fully qualify that a given circuit's stabiliser and magnet bypass circuitry can withstand a current decay equivalent to that of a quench event at the maximum expected operating current - in this case 11.08 kA. The test involved operating at 20 K, *i.e.* just above the critical temperature of Nb–Ti. This forced current to flow through the busbar's Cu stabiliser and each magnet bypass, qualifying the stability and continuity of all encompassed contacts and interconnects. A number of engineering challenges arose with solutions which involved the reconfiguring of the PCs and modification of the existing protection systems. Furthermore, with thermal runaways being a possibility, machine protection was of significant importance. A preliminary risk assessment concluded that the procedure should involve a series of powering cycles, with each cycle increasing in peak current until either a failure occurs or the maximum current profile has been carried out successfully, without measurable non-conformities or system damage. Furthermore, due to the non-standard operating conditions, all powering and protection systems were to be fully commissioned both before powering and after being switched back for normal operation.

To further investigate the risk involved in carrying out a CSCM test, a pre-consolidation type test was carried out to qualify all powering and protection systems. In addition, a CSCM simulation program was developed to allow for the calculation of accurate protection system thresholds for individual busbar segments and magnet bypasses. The type test concluded with all systems being successfully qualified under the atypical conditions and all induced thermal runaways being successfully detected and protected against. There were, however, unpredictable diode-lead characteristics similar to that witnessed in SM18 diode specific experiments as well as unexpected phenomena occurring wherein a thermal runaway induced in a busbar splice began to slow down, resulting in near-critical hot-spot temperatures. As a result, all threshold margin were significantly reduced. All subsequent tests were within specification. It was therefore deemed safe to proceed with an LHC-wide full qualification campaign. Overall, the type test was deemed a success and an LHC-wide full qualification campaign was scheduled for all main dipole circuits, post-consolidation.

The LHC-wide full qualification campaign also incorporated 400 A resistance measurements between each powering cycle to give information regarding busbar and bypass-diode

resistance trends. As expected following interconnect consolidation, all busbar resistances remained relatively low and consistent throughout. However, numerous diode-to-heat-sink contact resistances were noted to significantly increase from one powering cycle to the next, up to 10 times that of previous observations. To minimise the threat posed by such resistance increases, a stepped powering cycle was developed to deliberately induce the recovery trend observed in SM18 diode specific tests. The results were successful, reducing excess resistances by a factor of 7. The campaign continued and eventually concluded with all circuits being subject to a full programmed exponential decay from 11.1 kA, without measurable damage. As a consequence, magnet training could be carried out systematically with the minimum possible risk.

Magnet training to 6.5 TeV equivalent current levels was carried out on all circuits. Although exceeding twice that of predictions, further justifying having done CSCM, all sector's magnets were eventually qualified to 11.08 kA following 179 magnet quenches, 158 of which were on separate magnets. It was clear that, as predicted, Firm 3 magnets performed the worse, being the source of over 80% of all magnet quenches. All results must be taken into account in considering any further training beyond 6.5 TeV.

Overall, the CSCM format provides a systematic and conclusive means of qualifying the safe operation of a superconducting magnet string to a specific current level. The procedure is adaptable to any circuit of a similar type and, as such, could be implemented for other magnet circuits of the LHC for future qualification and or could be utilised in other or new accelerators as an integral part of hardware commissioning.

Chapter 5

UFO Event Modelling

Unidentified falling objects (UFOs), hypothesised to be falling dust particles interacting with the beam, were shown to be producing millisecond spikes in beam losses at various locations throughout the LHC, the energy deposition of which could potentially quench adjacent dipole magnets. To investigate, a numerical model was constructed which was capable of simulating a typical UFO event as well as calculating the resultant energy deposition in adjacent magnets. Following construction of the model, Monte-Carlo simulations were carried out to reproduce measured data from 2012 and to estimate quench probability predictions for 6.5 TeV operation. This chapter, based on works presented in the listed publications [p1] and [p4], details the development and subsequent use of the numerical model.

5.1 UFO Event Model

In order to determine the true extent of the threat posed by UFOs and to study plausible mitigation strategies, a UFO event model was constructed. The aim was to develop a model capable of reproducing measured 4 TeV UFO events, giving insights into the underlying phenomena, as well as allowing for extrapolation to 6.5 TeV equivalent scenarios, to determine whether or not the resultant particle shower would be capable of quenching adjacent dipole magnets.

The model determines the likelihood of a UFO event quenching an adjacent magnet by calculating the resultant beam losses and the corresponding BLM signals. In accordance with [99], the corresponding integrated loss rate picked up by a BLM at which an adjacent magnet would quench follows

$$\mathrm{BLM}_{\mathrm{signal@quench}} = \frac{\mathrm{BLM}_{\mathrm{response}}(T_{\mathrm{p}}, t_{\mathrm{int}}) \cdot T_{\mathrm{quench}}(t_{\mathrm{int}}) \cdot \mathrm{QLU}(t_{\mathrm{int}})}{T_{\mathrm{deposited}}(t_{\mathrm{int}})}, \tag{5.1}$$

where t_{int} is the integration time and T_{p}, $T_{\text{deposited}}$ and T_{quench} are the proton energy, the energy deposited in the magnet and the energy deposition required to quench, respectively. QLU is the quench level uncertainty factor – a remaining unknown (detailed in Section 5.1.3). The BLM signal itself is calculated following

$$\text{BLM}_{\text{signal}} = \dot{N}_{\text{p}}(t) \cdot \text{BLM}_{\text{response}}, \tag{5.2}$$

where $\dot{N}_{\text{p}}$ is the total proton loss for a given interaction, *i.e.* the integrated number of inelastic collisions, and $\text{BLM}_{\text{response}}$ is the BLM signal per lost proton as calculated numerically by FLUKA simulation. FLUKA is CERN's own general purpose software tool for calculating outcomes of particle interactions with matter [100].

5.1.1 BLM Configuration & UFO Event Recording

LHC BLMs have multiple running sums (RS), *i.e.* integration times, and the ability to set specific thresholds for each. It follows that, as the duration of a typical UFO event is of the order of 1 ms, RS4 ($t_{\text{int}} = 640\ \mu$s) is the most appropriate and therefore only the RS4 threshold is altered to protect against UFO events. Furthermore, for a UFO event to be recorded if the peak BLM signal does not exceed the RS4 threshold, it must exceed a lower post-processing threshold which is applied by the so called UFO Buster [83]. Algorithm 1 shows the filtering process which is applied by the UFO Buster to raw data in order for a UFO event to be recorded.

Algorithm 1 Run 1 UFO Buster Analysis

for All Arc BLMs **do**
  if 2 BLMs within 40 m detect $\max\left(\text{RS}(640\ \mu\text{s}, t)\right)/640\ \mu\text{s} > 0.1$ Gy/s **then**
    $N_{\text{ufo}} = N_{\text{ufo}} + 1$ (*Accumulative No. of UFOs)*
  end if
end for

The process defines the requirement for at least two BLMs within 40 m of each other to have picked up a BLM signal in RS4 which is at least five times greater than the background noise, *i.e.* >0.1 Gy/s. This prevents the recording of single BLM signal spikes and minimises the data to significant events. Note that this is all done in post-processing and that the UFO Buster algorithm can be altered such that a new and corresponding UFO database can be acquired at any time.

5.1.2 Run 2 Repositioning of BLMs

With dust and debris not being observed to have affected previous proton accelerators, it was not taken into consideration in designing the LHC. As a result, BLMs along the Arcs were only installed adjacent to the quadrupole magnets because in focusing and defocusing the beam they incur significant transverse effects as well as the fact the magnet aperture and transverse beam dimensions are at their largest at these locations [61]. However, with the discovery of the UFO phenomenon throughout Run 1, it was decided that redundant BLMs should be relocated in order to detect losses in the dipole magnet sections. Figure 5.1 depicts the planned BLM relocation within a typical Arc half-cell – carried out LHC-wide. This relocation significantly increases coverage, but is at the expense of redundancy.

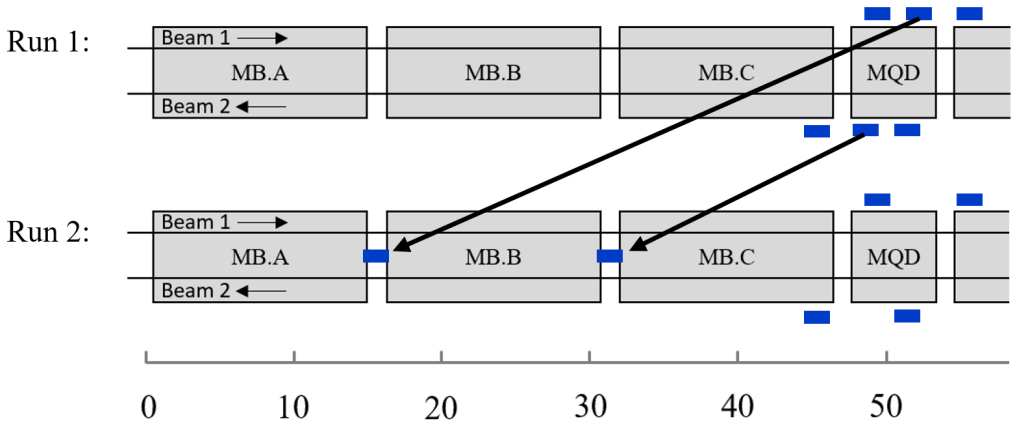

Figure 5.1: Illustration of the LS 1 relocation of BLMs in a typical Arc half-cell. Relocation reduces blind spots, allowing the monitoring of beam losses in the dipole regions.

Figure 5.2 shows the simulated BLM response for both Run 1 and Run 2 BLM locations. It is clear that the planned repositioning of the BLMs for Run 2 does significantly increase coverage. Note that proton losses adjacent to quadrupoles produce signals a factor of 3 larger than those adjacent to the dipoles and that this must be taken into account for Run 2 statistical analyses. Simulations were carried out by A. Lechner.

To take into account the fact that the factor of 3 difference in signal response for the quadrupoles would bias statistics, an adjustment was proposed for the BLM algorithm for which BLM thresholds in the quadrupoles would be reduced accordingly. However, as this was not necessary for machine protection and would only affect UFO event statistics, the algorithm was not altered and the factor of 3 difference in signal response was only taken into account for data analysis, *i.e.* post-processing.

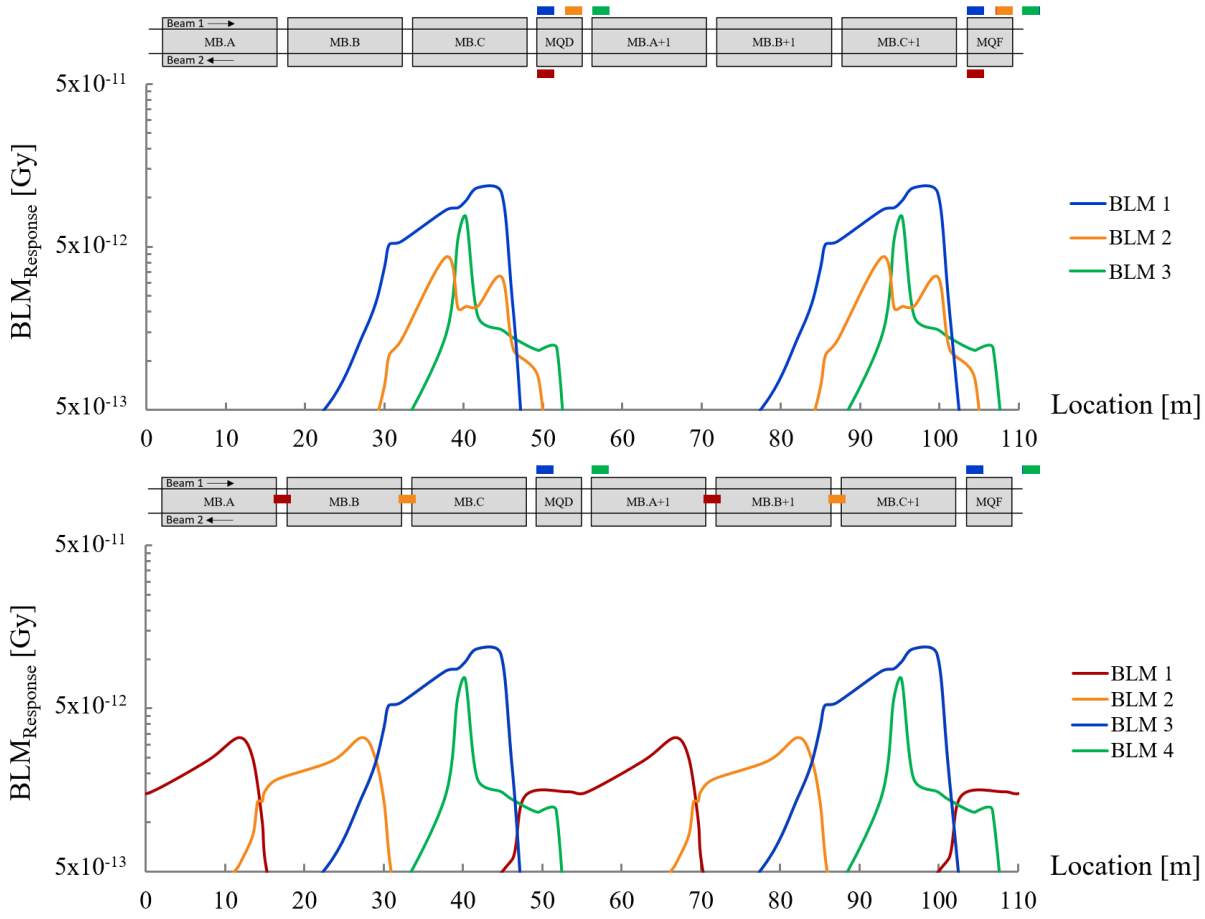

Figure 5.2: FLUKA-modelled 6.5 TeV per proton BLM signal response for a given location of an inelastic collision along a typical Arc cell. Signal response is shown for both pre-LS 1 (top) and post-LS 1 (bottom) BLM configurations. Results show the post-LS 1 configuration to have significantly increased coverage. Simulations data courtesy of A. Lechner.

5.1.3 Quench Level Uncertainty

As discussed in Section 3.3.4, there are numerous scenarios other than UFO events which may result in sufficient localised beam losses to quench an adjacent magnet. Simulations of such event have been carried out in the past [101], however, with little empirical data, significant discrepancies between the quench data and simulations results remain. Figure 5.3 shows the correlation between simulation and the only three available measured beam-induced quench data points. Simulations are in good agreement with measured data for loss events which last several seconds or are as short as a few microseconds. However, there is a significant discrepancy of up to a factor of 4 in the millisecond time range. Due to the lack of measured events, the validity of both simulation and measured data remained questionable and for the purposes of this study, a quench level uncertainty (QLU) factor was introduced.

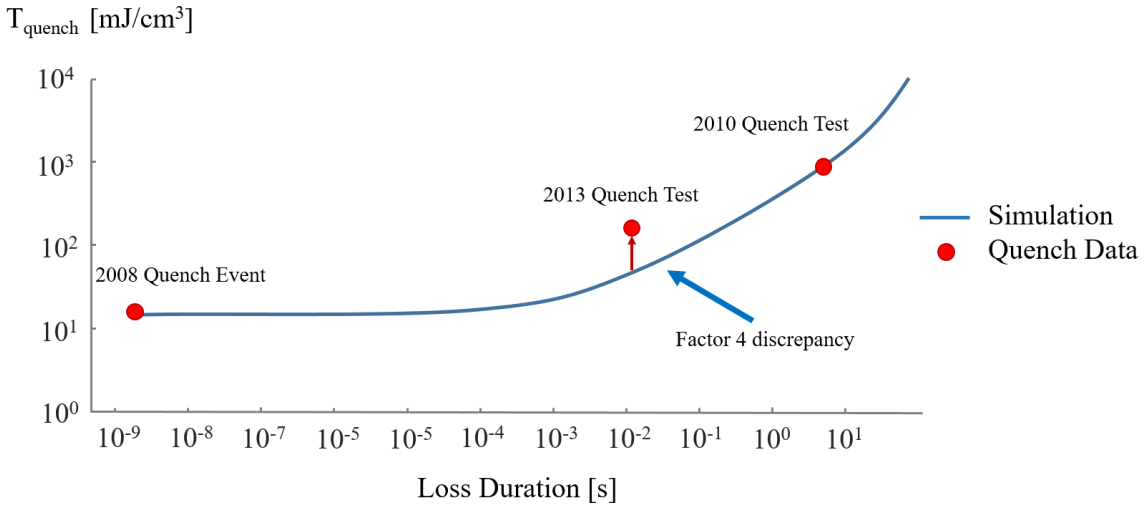

Figure 5.3: Comparative log-log plot of the minimum energy density required to quench as a function of loss duration. Note the discrepancy of a factor of 4 between measured data and simulation in the millisecond range. To account for this discrepancy in setting the BLM thresholds, a quench level uncertainty (QLU) factor is introduced. Data courtesy of B. Auchmann [99].

5.1.4 UFO Event Analytical Model

To calculate the proton loss rate, $-\mathrm{d}N_\mathrm{p}/\mathrm{d}t$, an analytical model was developed in order to simulate a typical Arc UFO event using plausible input parameters. This section describes the developed model.

Beam size

It was important that any variation in beam size and or shape along the Arc was taken into account. Horizontal and vertical beam sizes can be derived from the emittance and respective beta functions in accordance with [15], such that

$$\sigma_x(s) = \sqrt{\beta_x(s)\epsilon}, \tag{5.3}$$

$$\sigma_y(s) = \sqrt{\beta_y(s)\epsilon}, \tag{5.4}$$

where ϵ is the emittance and $\sigma_x, \sigma_y, \beta_x$ and β_y are the horizontal and vertical beam sizes and beta functions, respectively. Furthermore, transverse beam dispersion due to inhomogeneous bunch energies in the horizontal plane is also taken into account such that

$$\Delta x = D_x(s)\frac{\Delta p}{p_\mathrm{nom}} \tag{5.5}$$

where D_x is the dispersion function, p_{nom} and Δp are the nominal particle momentum and the deviation from nominal, respectively, and Δx is the difference in horizontal position with respect to nominal energy particles [37]. The dispersion function maxima and minima is given by

$$D_x^{\pm}(s) = \frac{L\theta}{\sin^2\frac{\psi}{2}}\left(1 \pm \frac{1}{2}\sin\frac{\psi}{2}\right), \tag{5.6}$$

where the bending angle

$$\theta = \frac{l_{\text{dipole}}}{r_{\text{dipole}}} \tag{5.7}$$

is acquired from the length of each dipole, l_{dipole} and the bending radius the dipole infers on the beam, r_{dipole}. For a magnetic lattice with dipoles of equal bending angles, the per cell phase advance

$$\psi \approx \frac{\pi}{2}. \tag{5.8}$$

The β (in meters) for a given location, accounts for transverse particle motion and is one of the main Twiss parameters used in the development of an accelerator's magnetic lattice (discussed in Section 1.6.4). Following [37], the maximum beta function is defined by

$$\beta_{\text{max}} = \frac{2L}{\sin\psi}\left(1 + \sin\frac{\psi}{2}\right), \tag{5.9}$$

and L is the length of a typical Arc half-cell, ∼55 m. Taking this into account, a parabolic fit for the beta functions along an Arc cell can be defined, such that

$$\beta_x(s) = \begin{cases} \frac{L}{2} + \frac{2(s+b)^2}{L}, & \text{if } s < L, \\ \frac{L}{2} + \frac{2(s-(b+2L))^2}{L}, & \text{if } s > L, \end{cases} \tag{5.10}$$

$$\beta_y(s) = \begin{cases} \frac{L}{2} + \frac{2(s-(b+L))^2}{L}, & \text{if } s < L, \\ \frac{L}{2} + \frac{2(s+(b-L))^2}{L}, & \text{if } s > L, \end{cases} \tag{5.11}$$

where s is the longitudinal distance along the cell and

$$b = \frac{1}{2}\sqrt{2L\beta_{\text{max}} - L^2} - L. \tag{5.12}$$

The emittance, ϵ, can be calculated from the normalised emittance, ϵ_{n} (4 μm for Run 1, 2.4 μm for Run 2) in accordance with [37], such that

$$\epsilon = \frac{\epsilon_{\text{n}}}{\gamma\beta} \sim \frac{\epsilon_{\text{n}}}{\gamma}, \tag{5.13}$$

where γ is the relativistic Lorentz factor, calculated from the proton mass, m_{p}, and kinetic energy, E_{p}, corresponding to

$$\gamma = 1 + \frac{E_{\mathrm{p}}}{m_{\mathrm{p}}c^2}, \tag{5.14}$$

and the relativistic beta

$$\beta = \frac{v_{\mathrm{p}}}{c}, \tag{5.15}$$

where c is the speed of light in a vacuum and v_{p} the proton velocity, *i.e.* 99.9999991% the speed of light.

Figure 5.4 shows the resultant 3–D beam profile for a typical Arc lattice cell as constructed from the above equations.

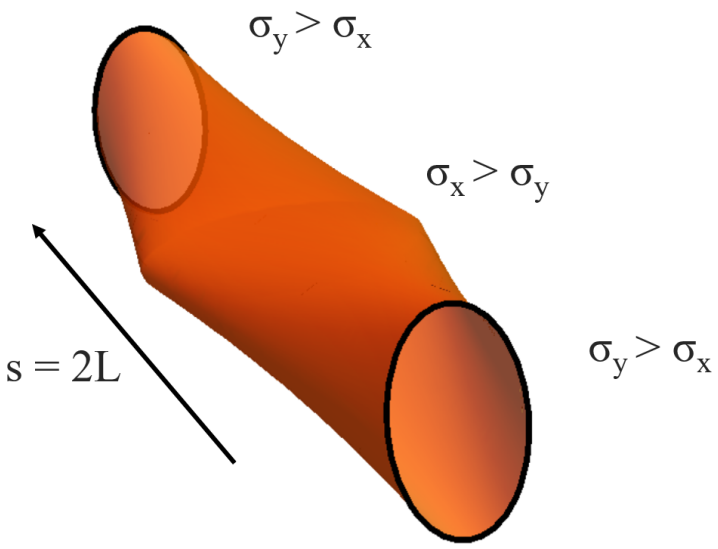

Figure 5.4: 3–D illustration of the transverse beam profile variation along a typical LHC Arc cell. Note that the profile varies from predominately vertical to predominately horizontal and back again, following the magnetic lattice. The image shown is longitudinally not to scale.

Beam-Charge Distribution

For simplicity, it is assumed that the transverse beam distribution follows a normal Gaussian distribution. Furthermore, as the time scale of interaction has been observed to be in the order of 1 ms and the circulation time for each bunch is approximately ∼89 μs, it is assumed that the beam can be simulated as continuous rather than being bunched or having significant gaps, for example, the abort gap [102]. Even though continuous,

following special relativity, the distance between charges, in the rest frame of the protons, still appears stretched by the Lorentz factor with respect to the laboratory frame, such that

$$\tau = \gamma\tau', \tag{5.16}$$

where τ and τ' are the laboratory and rest frame charge densities with respect to the beam.

It follows that τ' can be defined as

$$\tau'(x,y) = \frac{N_\mathrm{p} e}{\gamma C 2\pi\sigma_x\sigma_y} \mathrm{e}^{-\frac{x^2}{2\sigma_x^2}-\frac{y^2}{2\sigma_y^2}}. \tag{5.17}$$

where N_p is the number of protons in the beam, e is the electron charge, C the LHC circumference and the exponential component accounts for the Gaussian beam size and distribution. The function for current density in the laboratory frame can therefore be acquired in accordance with

$$J(x,y) = \frac{N_\mathrm{p} f e}{2\pi\sigma_x\sigma_y} \mathrm{e}^{-\frac{x^2}{2\sigma_x^2}-\frac{y^2}{2\sigma_y^2}}, \tag{5.18}$$

where the revolution frequency, f, follows

$$f = v_\mathrm{p}/C \approx c/C. \tag{5.19}$$

Electric Field

The transverse electrostatic and corresponding magnetic fields in reference to the laboratory frame, $\vec{E}_\mathrm{b}$ and $\vec{B}_\mathrm{b}$, are related to the those of the rest frame of the protons, in that

$$\vec{E}_\mathrm{b} = \gamma\vec{E}'_\mathrm{b} \tag{5.20}$$

$$\vec{B}_\mathrm{b} = \frac{\gamma}{c^2}\vec{v}_\mathrm{p} \times \vec{E}'_\mathrm{b}. \tag{5.21}$$

Note that the Lorentz factor of the charge density is balanced by the same factor in the fields. This allows one to calculate $\vec{E}_\mathrm{b}$ and thus $\vec{E}'_\mathrm{b}$ from Equation 5.18. Furthermore, note that the beam's magnetic field influence is a second order effect and for simplicity is not taken into account in this model.

In accordance with the Bassetti-Erskine formula for a 2-D Gaussian charge distribution

[103], one can find that for $\sigma_x > \sigma_y$,

$$\sigma_0 = \sqrt{2(\sigma_x^2 - \sigma_y^2)}, \tag{5.22}$$

$$z_1 = \frac{1}{\sigma_0}\left(x\frac{\sigma_y}{\sigma_x} + \mathrm{j}y\frac{\sigma_x}{\sigma_y}\right), \tag{5.23}$$

$$z_2 = \frac{1}{\sigma_0}(x + \mathrm{j}y), \tag{5.24}$$

such that

$$z_1^2 - z_2^2 = -\frac{x^2}{2\sigma_x^2} - \frac{y^2}{2\sigma_y^2}, \tag{5.25}$$

giving (similarly for $\sigma_y > \sigma_x$),

$$E_x = \frac{N_\mathrm{p} e}{2\varepsilon_0 C\sqrt{\pi}\sigma_0}\ \mathrm{Im}\left(w(z_2) - \mathrm{e}^{(z_1^2 - z_2^2)}w(z_1)\right), \tag{5.26}$$

$$E_y = \frac{N_\mathrm{p} e}{2\varepsilon_0 C\sqrt{\pi}\sigma_0}\ \mathrm{Re}\left(w(z_2) - \mathrm{e}^{(z_1^2 - z_2^2)}w(z_1)\right), \tag{5.27}$$

where $\omega(z)$ is the complex generalisation of the error function, sometimes referred to as the Fadeeva function [104],

$$w(z) = \mathrm{e}^{-z^2}\left(1 + \frac{2\mathrm{j}}{\sqrt{\pi}}\int_0^z \mathrm{e}^{\zeta^2}\,\mathrm{d}\zeta\right). \tag{5.28}$$

For a circular beam profile, *i.e.* $\sigma_y = \sigma_x = \sigma$, certain values tends towards zero, producing numerical instabilities. However, following [105], one can assume that the complex error function has an asymptotic expansion, such that

$$w(z) \approx \frac{\mathrm{j}}{\sqrt{\pi}z}, \tag{5.29}$$

resulting in,

$$E_x = -\frac{N_\mathrm{p} e}{2\pi\varepsilon_0 C}\mathrm{Im}\left(\frac{1}{\mathrm{j}x - y}\right)\left(1 - \mathrm{e}^{-\frac{x^2+y^2}{2\sigma^2}}\right), \tag{5.30}$$

$$E_y = -\frac{N_\mathrm{p} e}{2\pi\varepsilon_0 C}\mathrm{Re}\left(\frac{1}{\mathrm{j}x - y}\right)\left(1 - \mathrm{e}^{-\frac{x^2+y^2}{2\sigma^2}}\right). \tag{5.31}$$

Figure 5.5 depicts the resultant electric field as produced by the beam for various longitudinal locations and beam profiles.

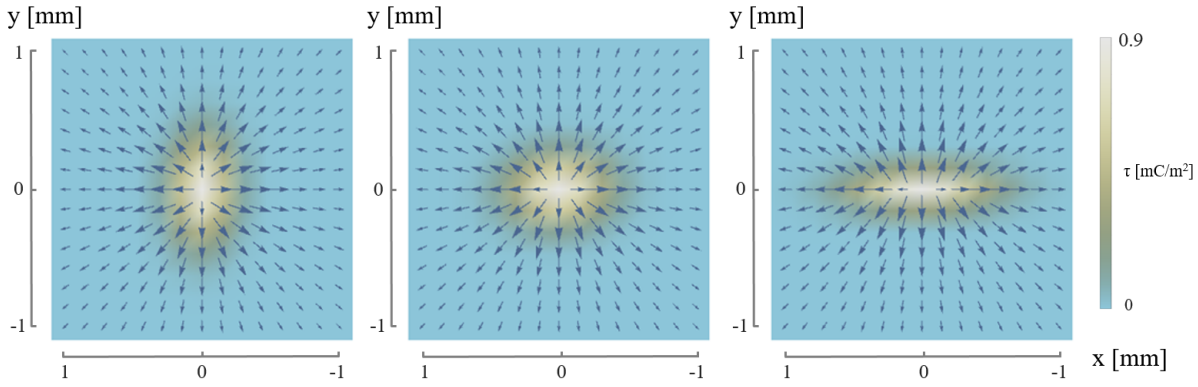

Figure 5.5: Resultant electric field map of the proton beam as it varies along the an Arc half-cell. Three locations shown, $s = 0, \frac{L}{2}, L$, where L is the half-cell length. Colour depicts the variation in charge density.

Mirror-Charge Effects

As the beam screen contains highly conductive copper lining (see Chapter 2 Section 1.6.2), mirror charges must also be considered. Following [79], the forces exerted on a macroparticle due to the mirror charge is given by

$$\vec{E}_{\mathrm{m}} = \frac{Qe}{16\pi\varepsilon_0 d_1^2}\vec{e}_{n_1} + \frac{Qe}{16\pi\varepsilon_0 d_2^2}\vec{e}_{n_2}, \tag{5.32}$$

where Q is the beam charge, d_1 and d_2 denote the vertical distance of the particle from the upper and lower horizontal surfaces of the beam screen, such that

$$d_1 = h - y, \tag{5.33}$$

$$d_2 = 2h - d_1 \tag{5.34}$$

with h being half the height of the beam screen and $\vec{e}_{n_1}$ and $\vec{e}_{n_2}$ the respective normal vectors. The result is that mirror charges are therefore most prevalent when the macroparticle is close to the beam screen surface.

5.1.5 UFO Dynamics

Progressing on previous theoretical works [79], [106] and [107], the following details the implemented macroparticle dynamics.

UFO Equation of Motion

The macroparticle acceleration at any given time is defined as a function of its transverse location, such that

$$\ddot{\vec{r}}(x, y, t) = \frac{Q(t)e}{m_{\mathrm{mp}}}\vec{E}(x, y) + \vec{g}, \tag{5.35}$$

taking into account the macroparticle charge, Q, gravity, g, and the forces exerted by the beam's electric field, $\vec{E}$, where e is the electron charge and m_{mp} is the macroparticle mass.

UFO Charge Rate & Electron Escape Energy

Elastic collisions during a UFO event may result in the positive ionisation of the macroparticle as electrons are ejected. As a result, the charge rate, $\dot{Q}$, plays an integral role in the electric field influence. The charge rate can be derived from the distribution of knock-on electrons in accordance with [109] (related to the Bethe-Bloch formula for stopping power [108]), taking into account relativistic influences and proton specific charge and spin-dependences following [110], such that

$$\frac{\partial^2 N_{\mathrm{e}}}{\partial T \partial s_{\mathrm{mp}}} \approx 2\pi r_{\mathrm{e}}^2 m_{\mathrm{e}} c^2 n \frac{1}{T^2}, \tag{5.36}$$

where N_e is the number of knock-on electrons, T is the kinetic energy transferred to the electron and s_{mp} is the incident proton penetration depth. The classic electron radius (all constants) follows

$$r_{\mathrm{e}} = e^2/(4\pi\varepsilon_0 m_{\mathrm{e}} c^2), \tag{5.37}$$

where m_{e} is the electron mass and ε_0 the vacuum permeability. Finally, the electron density

$$n = (N_{\mathrm{A}} Z \rho)/(A M_{\mathrm{u}}), \tag{5.38}$$

where N_{A} is Avogadro's constant, Z and A are the atomic number and relative atomic mass of the macroparticle material, respectively, ρ the macroparticle density, and M_{u} the molar mass constant. The resultant charge rate formula is therefore

$$\dot{Q}(x, y, t) = \int_a \int_S \int_{T_{\min}}^{\infty} J(x, y) \frac{\partial^2 N_{\mathrm{e}}}{\partial T \partial s_{\mathrm{mp}}} \,\mathrm{d}T \,\mathrm{d}s_{\mathrm{mp}} \,\mathrm{d}a, \tag{5.39}$$

where a is the macroparticle's cross-sectional area, S the average path-length of the incident proton through the particle, and J is the proton beam current density, see Equation 5.18, with a total current of $N_{\mathrm{p}} e f$, where f is the LHC revolution frequency. Evalu-

ation of Equation 5.39 gives

$$\dot{Q}(x,y,t) = -\frac{2N_{\mathrm{p}} f R^3 \pi N_{\mathrm{A}} r_{\mathrm{e}}^2 m_{\mathrm{e}} c^2 \rho}{3\sigma_x \sigma_y T_{\mathrm{min}}(Q(t),R) M_{\mathrm{u}}} \mathrm{e}^{-\frac{-x^2}{2\sigma_x^2} - \frac{-y^2}{2\sigma_y^2}}, \tag{5.40}$$

where T_{min} is the minimum transferred energy required to eject a free electron. For simplicity, it is assumed that the maximum energy transfer is infinite as $T\mathrm{p}¿¿T\mathrm{min}$.

In previous models, the Coulomb potential, φ, was taken as the minimum energy for an electron to be ejected. However, for the electron to leave the macroparticle completely, it must traverse several μm through the macroparticle and in doing so has a chance of being reabsorbed.

Following [111], the practical range of knock-on electrons in a given material is

$$PR(T) = \frac{aT}{\rho}\left(1 - \frac{b}{1+cT}\right), \tag{5.41}$$

where T is the energy of the knock-on electron and a, b and c are empirical constants.

To determine the minimum energy required to escape from the macroparticle, the average path length, l, that an electron ejected from any location has to traverse in order to reach the macroparticle edge, must be determined.

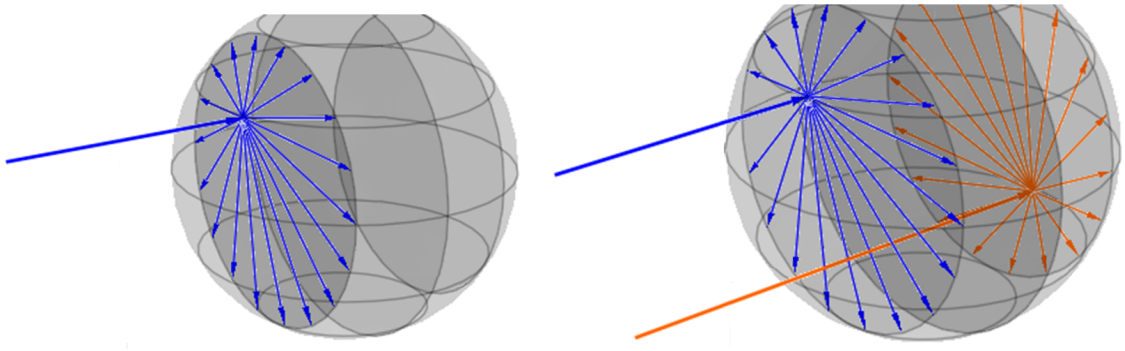

Figure 5.6: Graphical illustration of plausible knock-on electron scattering trajectories for a given incident proton of specific transverse location and penetration depth prior to elastic collision. Arbitrary colouring.

For simplicity, the macroparticle is assumed spherical and the transverse cross-section is therefore a circle of radius

$$\bar{R}(R, s_{\mathrm{mp}}) = \sqrt{R^2 - s_{\mathrm{mp}}^2}, \tag{5.42}$$

where R is the macroparticle radius and s_{mp} is the longitudinal incident proton penetration depth. Also stated by [111] is that knock-on electrons following elastic collisions with high-energy incident particles are predominately ejected normal to the incident particle path. Assuming this to be the case, for a given transverse collision location, r_{mp}, it follows that the average integrated electron path for all azimuthal angles θ, see Figure 5.6 (left), is

$$l''(R, r_{\mathrm{mp}}, s_{\mathrm{mp}}) = \frac{1}{2\pi} \int_0^{2\pi} r_{\mathrm{mp}} \cos\theta + \sqrt{r_{\mathrm{mp}}^2 \cos^2\theta + \bar{R}^2(R, s_{\mathrm{mp}}) - r_{\mathrm{mp}}^2}\, \mathrm{d}\theta. \tag{5.43}$$

Furthermore, integrating over the possible penetration depth, $2\bar{R}$, gives

$$l'(R, r_{\mathrm{mp}}) = \frac{1}{2\bar{R}} \int_0^{2\bar{R}} l''(R, s_{\mathrm{mp}}, r_{\mathrm{mp}}) \mathrm{d}s_{\mathrm{mp}} \tag{5.44}$$

and surface integration over all possible transverse starting locations, see Figure 5.6 (right), provides a formula for the average knock-on electron escape distance, such that

$$l(R) = \frac{1}{\pi R^2} \int_0^{R} \int_0^{2\pi} l'(R, r_{\mathrm{mp}}) r_{\mathrm{mp}}\, \mathrm{d}r_{\mathrm{mp}}. \tag{5.45}$$

Numerical evaluation of this equation gives $l(R)$=0.736 R. Finally, inverting equation (5.41), the minimum escape energy required for the average path length, $T_{\mathrm{esc}}(R)$, can therefore be obtained from the following

$$T_{\mathrm{esc}}(R) = \frac{a(b-1) + cl(R)\rho + \sqrt{4acl(R)\rho + (a(b-1) + cl(R)\rho)^2}}{2ac}. \tag{5.46}$$

and consequently,

$$T_{\mathrm{min}}(Q, R) = \varphi(Q) + T_{\mathrm{esc}}(R). \tag{5.47}$$

Figure 5.7 shows the resultant average energy a given knock-on electron requires to traverse a specific distance.

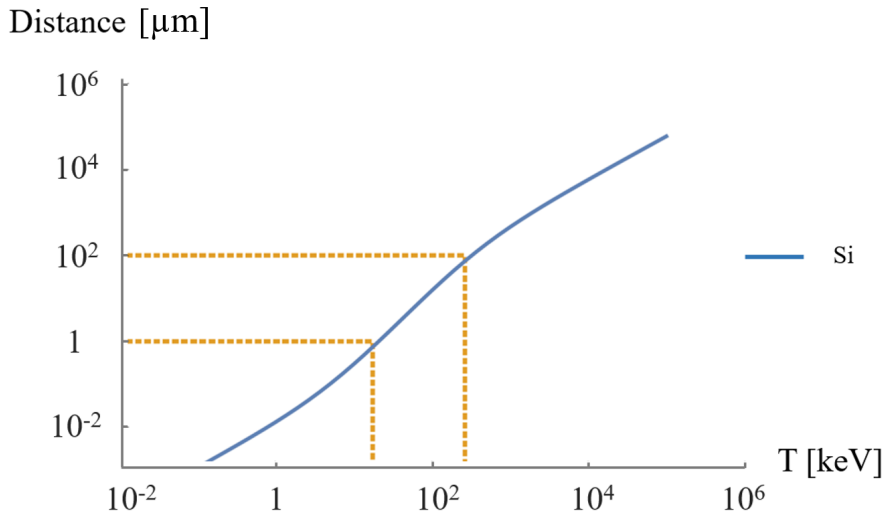

Figure 5.7: Plot of the energy transfer required for a knock-on electron to travel a specific distance within a material – example shown is for silicon material properties. The range of expected macroparticle radii, 1<R<100 μm, is marked.

Comparison to Previous Models

Figure 5.8 shows a comparison of previous and newly developed analytical formula for the minimum energy transfer and the charge rate. As expected, the new model requires significantly higher energy to be transferred before an electron is ejected. The resultant difference to the charge-rate, which ultimately determines the macroparticle–beam interaction time, would see ionisation to the point of repulsion take significantly longer.

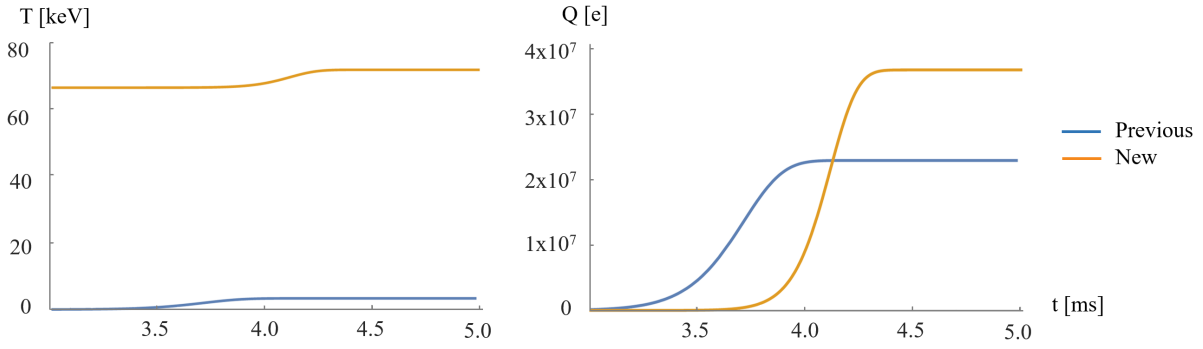

Figure 5.8: Minimum energy transfer (left) and integral of charge (right) comparison between newly developed and previous analytical formula. The difference is, as expected, significant.

To cross-check the proposed modifications, numerical simulations of the charge per proton ionisation of a spherical silicon (Si) particle by incident high-energy protons were carried out independently using CERN's own otherwise validated particle tracking software Garfield++ [112], by H. Schindler. Figure 5.9 shows the comparison between Garfield++

and the new analytical results. The average charge per proton was derived from the analytical model in accordance with (for $Q_0 = 0$)

$$Q_{\mathrm{pp}} = \frac{\dot{Q}}{\int_a J \, \mathrm{d}a}, \tag{5.48}$$

such that

$$Q_{\mathrm{pp}}(R) = \frac{4\pi R N_{\mathrm{A}} r_{\mathrm{e}}^2 m_{\mathrm{e}} c^2 \rho}{3\, T_{\mathrm{min}}(Q, R)\, M_{\mathrm{u}}}. \tag{5.49}$$

Analytical and numerical simulations agree to within 30-40% of each other, which, considering the orders of magnitude difference in comparison to previous analytical models, was deemed acceptable. The discrepancy is, however, still under investigation and is possibly due to the assumption of solely orthogonal electron scattering trajectories.

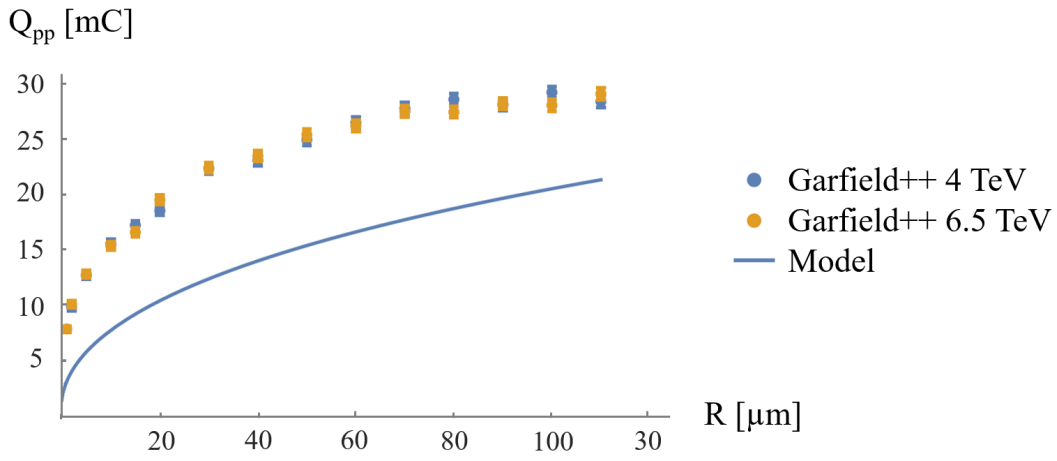

Figure 5.9: Per incident proton charge comparison between new analytical formulae and numerical simulations carried out using Garfield++ software.

Proton Loss Rate

Finally, the proton loss rate, $\dot{N}_{\mathrm{p}}$, corresponding to the rate of inelastic collisions produced by incident protons as they pass through the macroparticle, is calculated in accordance with [79], such that

$$\dot{N}_{\mathrm{p}}(x, y) = \int_a \int_S J(x, y) \Sigma_{\mathrm{iel}} \, \mathrm{d}s_{\mathrm{mp}} \, \mathrm{d}a, \tag{5.50}$$

where $J(x, y)$ is the laboratory frame beam current density, see Equation 5.18, and the macroscopic cross-section of the possible inelastic interaction region

$$\Sigma_{\text{iel}} = \sigma_{\text{iel}} \rho_{\text{A}} \tag{5.51}$$

is the product of the corresponding microscopic cross-section, σ_{iel}, and the atomic density

$$\rho_{\text{A}} = (N_{\text{A}} \rho)/(A M_{\text{u}}). \tag{5.52}$$

The resultant formula can be written as shown,

$$\dot{N}_{\text{p}}(x, y) = \frac{2 N_{\text{p}} f \sigma_{\text{iel}} R^3 N_{\text{A}} \rho}{3 \sigma_x \sigma_y A M_{\text{u}}} \mathrm{e}^{-\frac{-x^2}{2\sigma_x^2} - \frac{-y^2}{2\sigma_y^2}}. \tag{5.53}$$

5.2 Simulation Results

The following section details all simulation results, including direct comparison to measured 4 TeV UFO events, parametric studies and the extrapolation to 6.5 TeV operation.

5.2.1 Macroparticle Flight Path

In constructing the model, it was important to establish that a given falling macroparticle acts as would be expected, *i.e.* initially falls straight down and then depending on starting location and macroparticle size, has a corresponding path of repulsion.

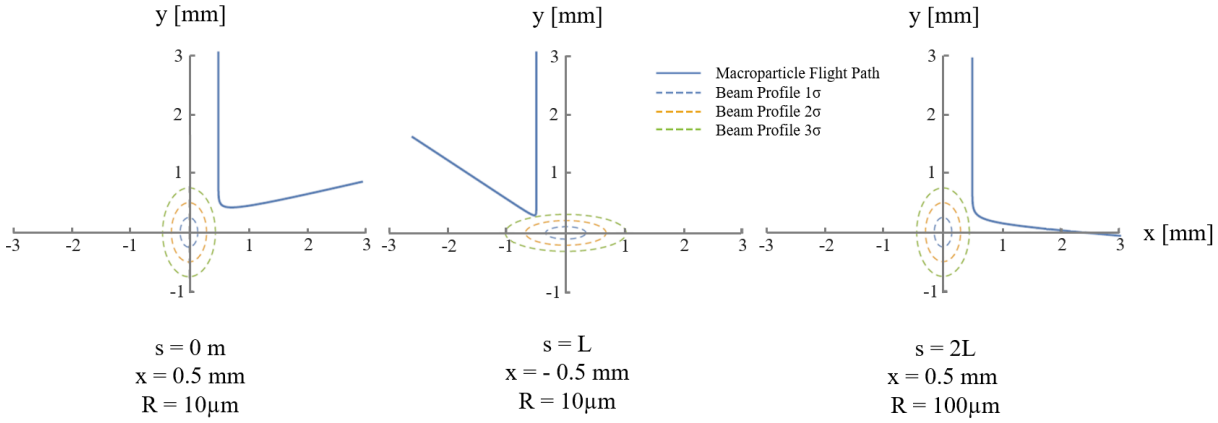

Figure 5.10: Plots of macroparticle flight paths for various longitudinal and transverse starting locations, s=0, L, 2L and x=±0.5 mm, and macroparticle radii, 10 and 100 μm.

Figure 5.10 plots the flight path for longitudinal and transverse starting location for

three events at various longitudinal and transverse starting locations and macroparticle radii. Note that larger macroparticles take longer to ionise to the point of repulsion and therefore fall further into the beam, increasing the overall interaction time.

5.2.2 Single Event Comparison

Figure 5.11 demonstrates that BLM signals of measured UFO events can be reproduced by the analytical model with realistic input parameters. Results are in very good agreement, with the model matching both the asymmetric radiation signal profile and signal magnitude of all involved BLMs. The example shown was chosen as it was the only Run 1 event which resulted in particle showers large enough to exceed the BLM thresholds, both dumping the beam and having the event data recorded. Simulation involved a Cu macroparticle of R=45 μm with a longitudinal starting location s=35 m. The initial vertical position of the macroparticle was at $y = h - R$, *i.e.* attached to the beam screen ceiling, and the initial macroparticle charge was zero.

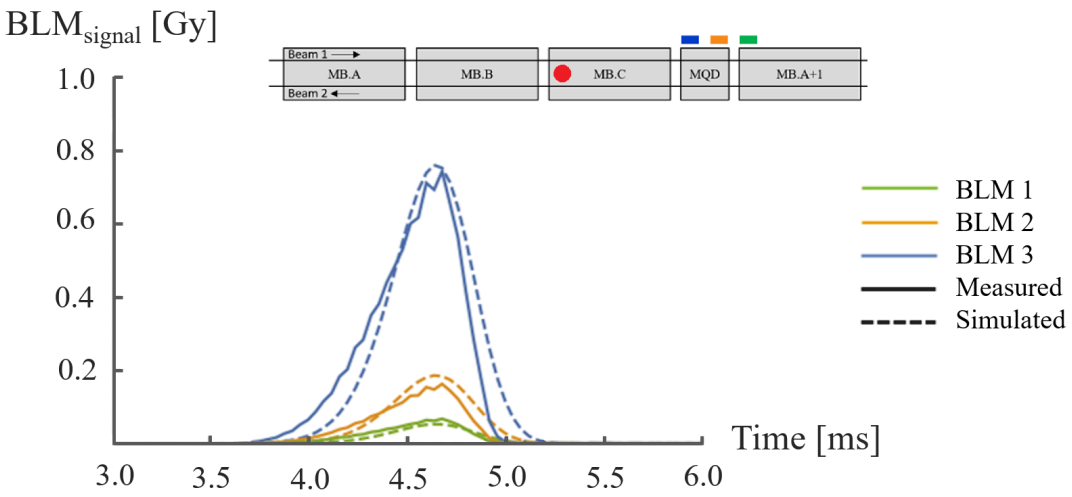

Figure 5.11: Comparison between the developed analytical model and a 4 TeV Arc UFO event measured in 2012. The event was reproduced using Run 1 BLM locations and simulating a Cu macroparticle of radius 45 µm falling at s=35 m, with an initial charge Q_0=0. BLM locations and UFO event location (red circle) are marked in lattice insert.

5.2.3 Matching Run 1 UFO Event Data

To allow for extrapolation to 6.5 TeV operation, it was important that a reasonable distribution for macroparticle radii could be defined. It was therefore determined that an iterative process of running Monte-Carlo simulations, varying the radii distribution, be

carried out until a reasonable fit for 4 TeV data was obtained. As a result, a UFO generator was added to the model allowing for input distributions for longitudinal location, transverse offset and macroparticle radii. It was assumed that the distribution for both longitudinal and transverse starting location be uniform along a typical Arc cell. For computational reasons, the transverse offset from the beam centre was reduced to ±2 mm to increase the probability of UFO scattering exceeding the UFO algorithm threshold.

Simulations proved successful, with the UFO generator being capable of recreating the 2012 UFO event data set, approximately 1800 recorded events, in a single Monte-Carlo simulation of 10,000 UFO events. A post-processing algorithm similar to that of the UFO Buster (discussed in Section 5.1.1) was implemented, *i.e.* not all UFO events result in high enough peak losses such that they are recorded. Figure 5.12 shows a cumulative histogram of 4 TeV peak BLM signals, as recorded in three BLM positions, as well as the well-matched simulated equivalents.

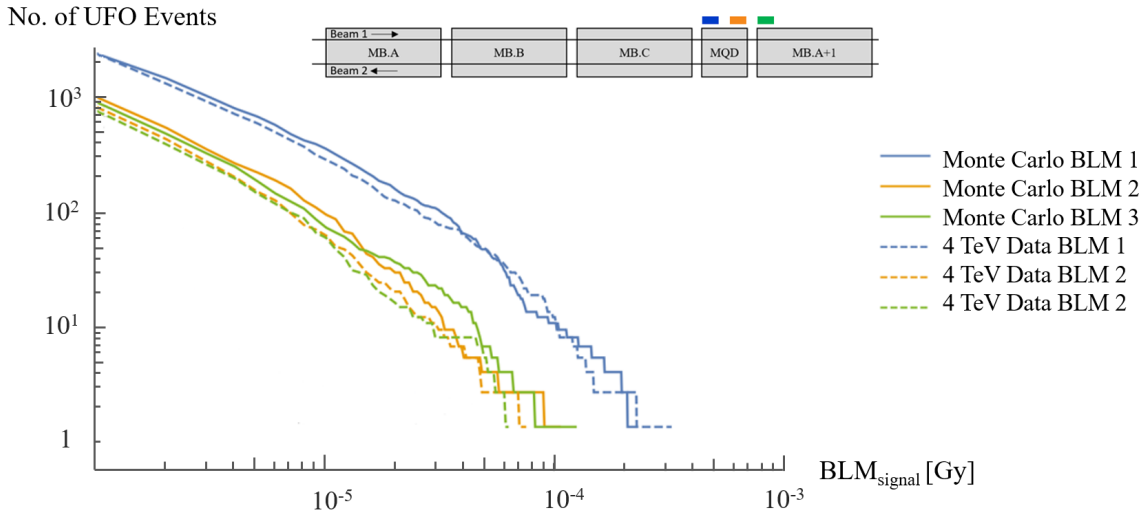

Figure 5.12: Cumulative histogram of peak BLM signals during UFO events in 2012 (dashed) and Monte-Carlo simulations (solid). Insert shows the corresponding BLM locations – Run 1.

To achieve the shown level of accuracy, the fit distribution for macroparticle radii, shown in Figure 5.13, was found to be a truncated distribution ranging from $5<R<100$ μm. This was acquired via simulations carried out with a Cu macroparticle with an initial charge of $Q_0 = 0e$. Considering the initial ballpark figure of 1-100 μm observed in the aforementioned dust study (Section 3.3.1) as well as assuming that largest of these macroparticles would be less likely to remain following cleaning procedures, the fit was deemed reasonable. Furthermore, even if the chosen material was different, the resultant distribution would simply be offset by a certain factor and, as such, uncertainties of this nature should have minimal effect on extrapolation results.

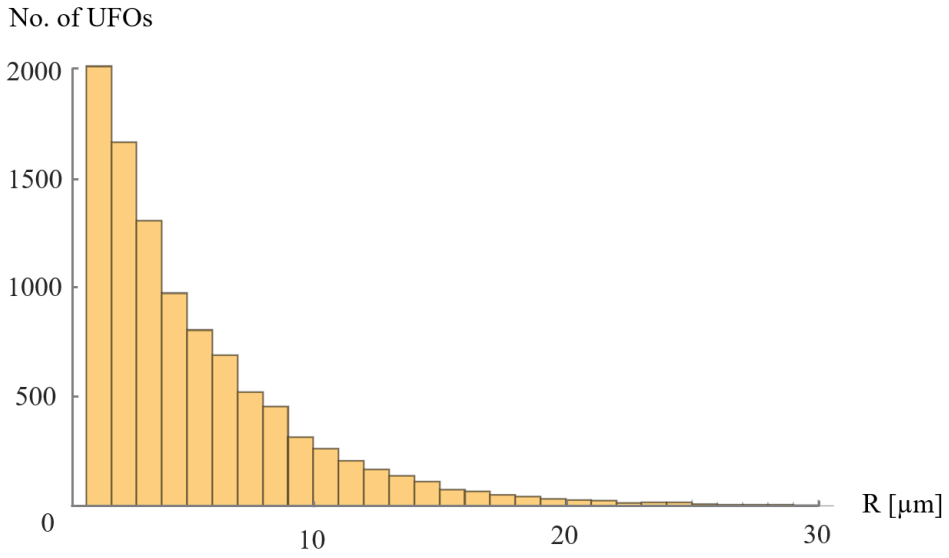

Figure 5.13: Resultant distribution for macroparticle radii used to match the full 2012 UFO data set of ∼1800 events. Note that there are significantly more macroparticles than events as smaller macroparticles, depending on location, may not produce enough losses to exceed the UFO algorithm threshold.

5.2.4 Extrapolation to 6.5 TeV & Quench Probability

Subsequent to acquiring a sufficiently accurate distribution for macroparticle radii, it was possible to extrapolate to 6.5 TeV. Beam parameters and BLM locations were configured as would be for Run 2.

Figure 5.14 shows the simulation results as well as the Run 2 $BLM_{Signal@Quench}$ thresholds for both simulated and measured values of the QLU factor. Table 5.1 shows the resultant predictions in terms of percentage of UFOs that would result in a quench.

Table 5.1: Model predictions of UFO-induced dipole quenches for a given QLU factor.

QLU factor -	Quench Probability [%]
1	1.24
4	0.11

It is clear that the QLU factor has a significant impact on the quench probability,

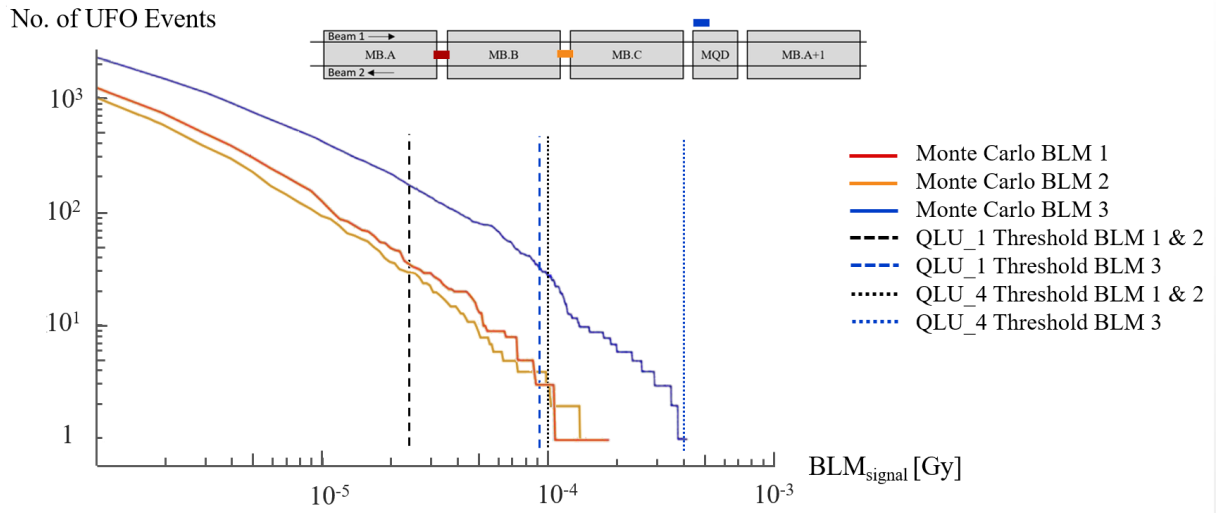

Figure 5.14: 6.5 TeV Monte-Carlo simulation results including quench limits for conservative and progressive QLU factors. Insert shows the corresponding BLM locations – Run 2.

varying predictions by up to a factor of 10. To obtain absolute numbers, an assumption must be made on the UFO rate. For example, consider the possibility that the UFO rate following LS 1 resumes the levels as was initially recorded during Run 1, *i.e.* ∼10 UFOs/hour. Furthermore, recent LHC availability predictions [114] can be taken into account, *i.e.* stable beam ∼7 hours/day, 7 days/week. The resultant worst case scenario (QLU=1) would, without appropriate BLM thresholds, see ∼6 quenches/week and the availability predictions would not hold. With an average quench recovery time being 8-12+ hours, mostly cryo-recovery time, this would have a near catastrophic impact on LHC availability. Also recall that such a UFO rate may even be optimistic as conditioning may have already occurred in 2010 prior to UFO events initially being monitored. Results give a clear statement that UFO-induced dipole quenches are not only a possibility at 6.5 TeV, but even with moderate estimates, are likely to be at the forefront of availability issues for Run 2.

5.2.5 Parametric Studies

To further understand how specific beam parameters influence the outcome of a UFO event, *i.e.* affects the integrated loss rate, a series of parametric studies were carried out. Table 5.2 lists the default beam and macroparticle parameters used for each study.

Table 5.2: Default beam and UFO parameters for parametric studies

Parameters	-
Beam Energy (T_b)	6.5 TeV
No. Bunches (N_b)	2808
Bunch Intensity (I_b)	1.3x10^{11}
Normalised Emittance (ϵ_n)	2.4 μm
Macroparticle Material	Cu
Macroparticle Radius (R)	10 μm
Initial Charge (Q_0)	0 e
Longitudinal Location (s)	53.46 m
Transverse Offset (x)	0.25 mm

Energy & Number of Bunches

Figure 5.15 shows how the loss rates scale with beam energy, T_b, and the number of bunches, N_b, *i.e.* beam intensity. Losses generally scale linearly with T_b, however, for N_b, there is a notable trend of an initial peak at ~200 bunches and a subsequent gradual reduction as N_b increases. The gradual reduction in loss rates for higher N_b can be explained by the increased probability of elastic collisions, resulting in earlier ionisation and macroparticle repulsion, out-weighing the increase in probability of inelastic collisions, *i.e.* losses. Overall, the variation in losses for T_b and N_b, between injection and near-nominal conditions is approximately one order of magnitude and is therefore considered to have significant influence on the event outcome.

Macroparticle Material & Radius

Figure 5.16 shows how the loss rates scale with macroparticle material and radius. Of the most likely macroparticle materials [83], Cu, produces the highest losses, with losses linearly scaling with nuclei cross-section and atomic density. Nonetheless, with loss rates scaling one order of magnitude depending on the macroparticle material, the material is also considered to have significant influence. However, in varying the macroparticle radius, even within the most common range of 1-30 μm, results differ by up to 5 orders of magnitude. This is due to it taking longer to ionise to the point of repulsion, inherently increasing the interaction time. It is clear that size and therefore interaction time are two of the most influential parameters.

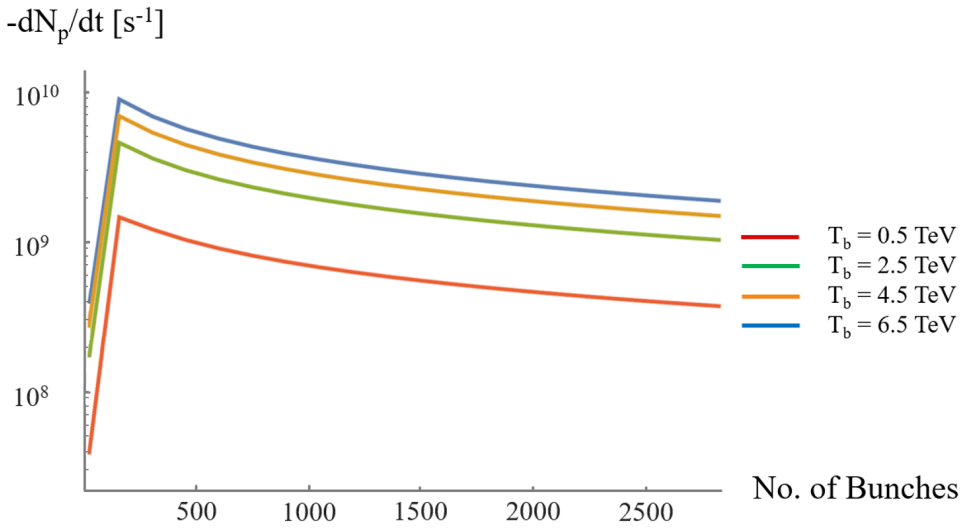

Figure 5.15: Log plot of the proton loss rate as a function of T_b and N_b. It shows that the losses saturate as N_b exceeds ~200 bunches. Overall variations in loss rates are not only significant, but would suggest that the UFO rate should decrease are the number of bunches is increased.

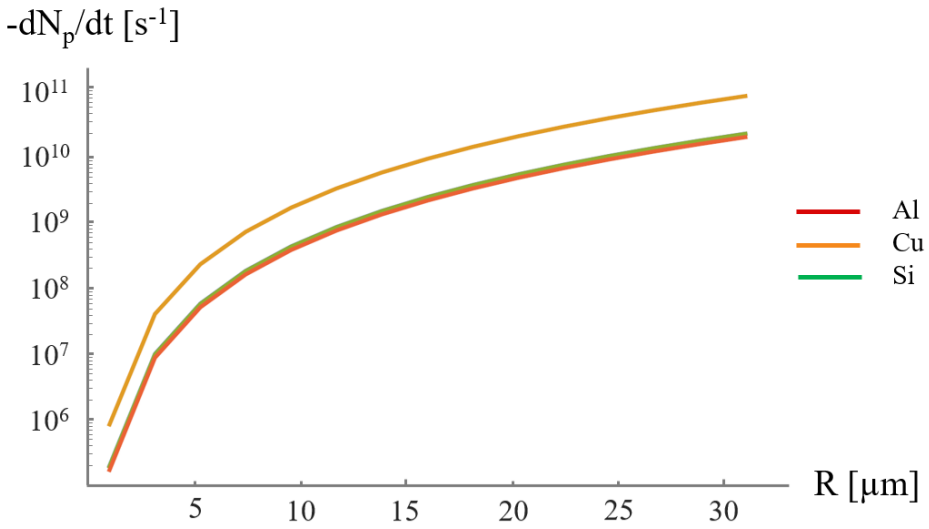

Figure 5.16: Log plot of the proton loss rate as a function of macroparticle radius and material. A relatively significantly increase is noted with change in material, however, loss rates are shown to scale considerably with macroparticle radius, with losses varying 5 orders of magnitude within the established range.

Longitudinal Location & Transverse Offset

Figure 5.17 shows how variations in the beam profile affects the loss rate. Similar to the influence of macroparticle radius, with a predominantly vertical beam, at s=0 m, peak loss rates vary by up to 6 orders of magnitude from being directly above the beam, x=0 mm, to being transversely offset by as little as 1 mm. This is due to the fact that, as the transverse offset increases, the falling macroparticle interacts less with the beam. For the same reason, when the beam is predominantly horizontal, s$\sim$55 m, transverse offsets of $\sim$1 mm are still close enough to the beam to result in a significant number of interactions and therefore produce peak losses similar to being directly above the beam centre. Note that, though not shown, for a predominately horizontal beam profile, offsets of up to $\sim$2 mm still produce significant losses, $> 10^8$ s^{-1}. From the results one can conclude that, assuming a uniform transverse distribution of dust and or debris, regions with a horizontal beam would produce the highest losses for a given macroparticle radius.

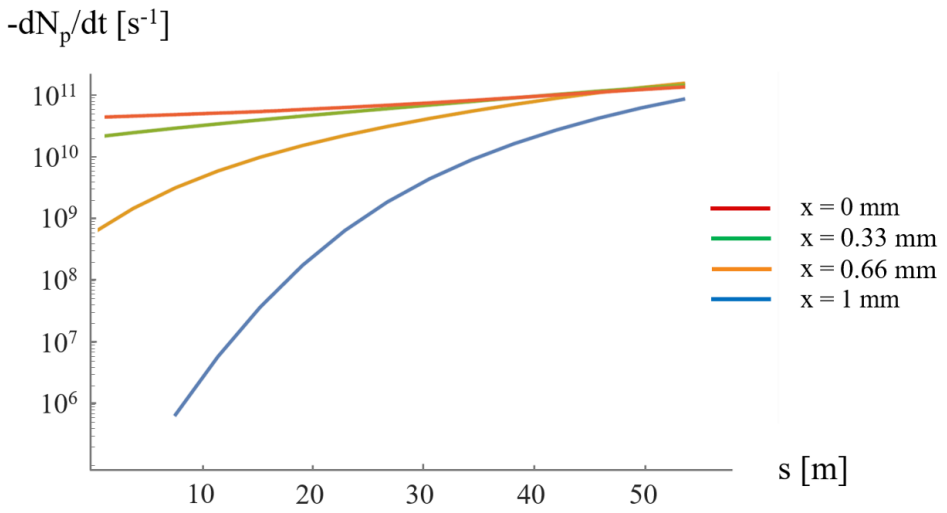

Figure 5.17: Log plot of the proton loss rate as function of longitudinal and transverse starting locations. Loss rates are notably dependant on both parameters. The most likely region to produce quenches is where the beam is predominantly horizontal, i.e. $s \sim 55$ m.

Normalised Emittance

Figure 5.18 shows the impact that the emittance has on the loss rate. It is shown that larger emittances can slightly reduce the peak loss rates. This is due to the process of ionisation beginning earlier and the reduced proton density reducing the probability of inelastic collisions. The variation in loss rates is approximately 50% and, as such, the emittance is considered to have a moderate influence of the outcome. It is also of note that the emittance has a notable influence on the rise and fall times and, even with the shown loss rate reduction, integrated event losses remain fairly consistent.

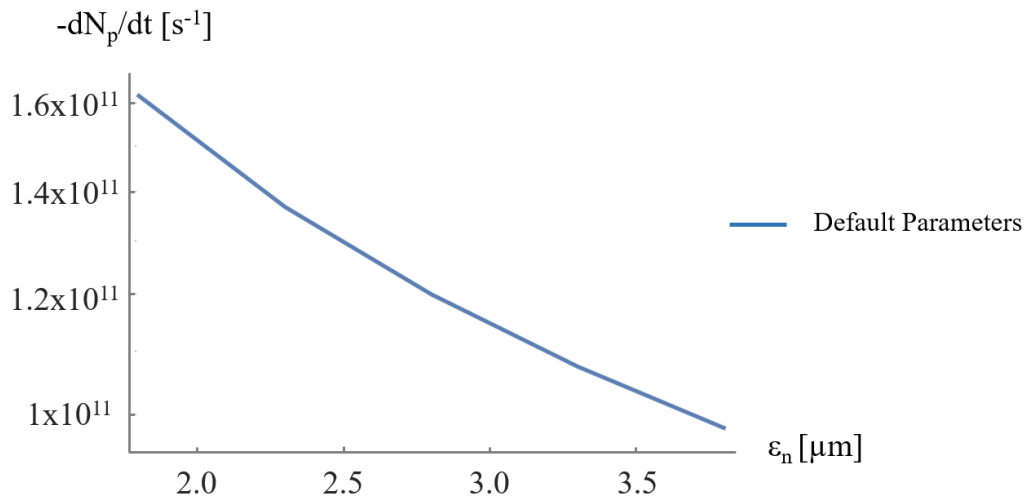

Figure 5.18: Plot showing the proton loss rate as function of beam emittance. The emittance is shown to have a moderate influence on the peak loss rates.

Initial Charge & Electron Cloud

As a significant increase in UFO rate was observed when transitioning to 25 ns bunch spacing (Figure 3.9), it is thought that the electron cloud (discussed in Section 1.6.4) may impact the UFO outcome [113] . Investigation resulted in the proposal that the presence of the electron cloud may negatively offset the initial charge of a macroparticle, either at rest or whilst falling. As such, it was important to investigate how this would affect the outcome by offsetting the initial charge, Q_0.

Figure 5.19 shows that, as the initial charge drops below $-10^5 e$, the integrated losses significantly increase. This is due to an initial increase in momentum as a result of the attractive forces as well as an increase in time required for the macroparticle to ionise to the point of repulsion. The result is that the macroparticle falls closer to the beam centre, greatly increasing interaction time. If correct, whilst operating with 25 ns bunch spacing, smaller macroparticles would be capable of producing the same integrated losses as larger neutral equivalents and would increase the probability of magnet quenches without a physical change in the release of macroparticles and or radii distribution. Furthermore, the attractive forces due to being negatively charged may contribute to the UFO release mechanism. Note, however, that this increase in loss rate peaks as the initial charge drops below -10^8, sharply reducing thereafter. This is due to the attractive forces sufficiently accelerating the particle such that it passes through the beam in the order of microseconds, *i.e.* not enough time for ionisation to occur. Follow-up studies throughout Run 2 are to be carried out in attempt to confirm these hypotheses [115].

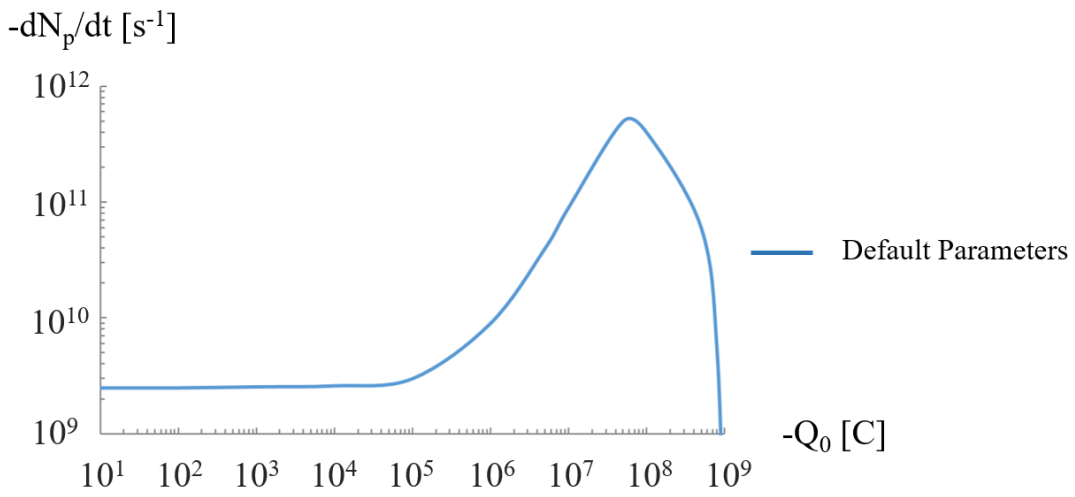

Figure 5.19: Log-Log plot showing that if the initial charge, Q_0, of the macroparticle is less than $-10^5 e$ the resultant peak loss rates rise significantly, peaking as the charge drops below $-10^8 e$ and sharply reducing thereafter.

5.3 Mitigation and Conditioning Strategies

Simulations predict that UFO-induced dipole quenches are likely to cause significant disruption to initial Run 2 operation. As such, plausible mitigation and conditioning strategies were investigated. This section details several options.

5.3.1 Iterative BLM Threshold Optimisation

The current strategy to maximise availability is to set BLM thresholds to dump the beam and initiate a slow power abort (2-3 hours recovery time) prior to losses resulting in a magnet quench (8-12+ hours recovery time). However, with the QLU factor having such a large influence on any calculations for BLM thresholds, it is proposed that the thresholds be set initially just below the quench level (QLU=4), such that the majority of UFOs would not trigger a beam dump, but rare events producing relatively large integrated losses could still, according to the model, quench adjacent magnets. This would give important empirical data allowing for accurate values for T_{quench} to be acquired, reducing uncertainty and allowing for optimised threshold calculations, *i.e.* continual revision of BLM thresholds based on T_{quench} input.

5.3.2 Defender Bunches

The interaction duration of a UFO event was found to be strongly dependant on the ionisation process of the macroparticles. As such, the concept of 'defender bunches' was proposed, in which multiple defender bunches, N_{b}^*, would have a larger emittance, ϵ_{n}^*,

relative to the rest of the beam and would essentially 'take the hit', so to speak, such that the macroparticle would be ionised to the point of repulsion prior to falling far enough as to interact with the rest of the beam. Figure 5.20 plots the simulation results carried out to determine the feasibility of defender bunches.

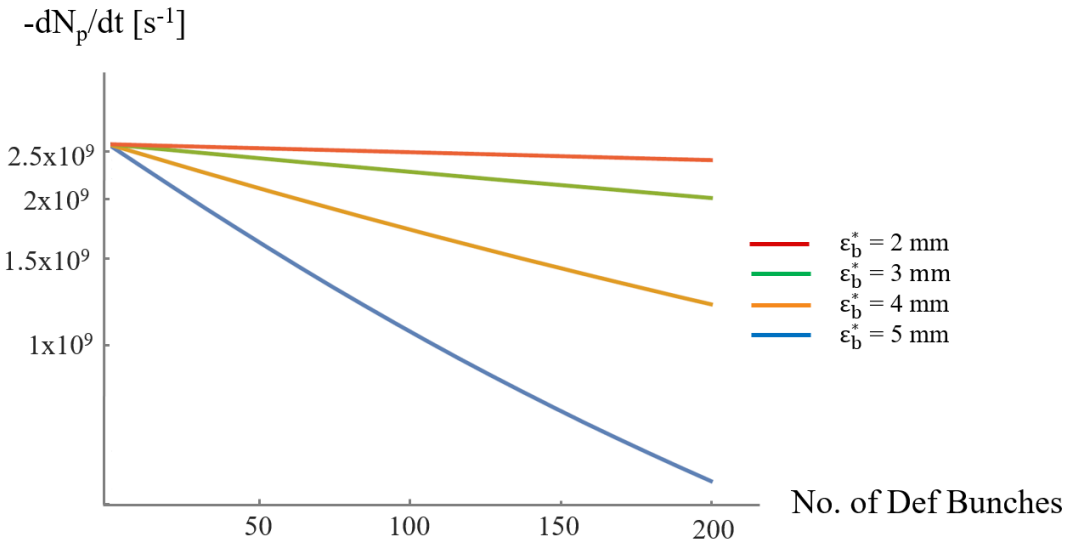

Figure 5.20: Feasibility study of the proposed defender bunches mitigation strategy. Unfortunately, similar to normalised emittance, the variation in losses is minimal.

In line with observations that the emittance has only a moderate influence on peak losses, the defender bunches were only able to reduce peak losses by a factor of 2-3 with reasonable parameters. Furthermore, this process would also reduce the quality and production rate of collision data as well as requiring significant changes to beam cleaning systems. As a result, it was concluded that such a method of UFO mitigation is not likely to be a viable option.

5.3.3 Conditioning Strategies

Run 1 statistics found that the UFO rate generally decreases with time, discussed in Chapter 3 Section 3.3.3 Figure 3.9. However, with the effect that LS 1 will have on UFO rate being unknown, it is uncertain if the dust consistency will remain 'conditioned' or will remain close to or exceed that of initially observed Run 1 rates. As such, any method of accelerating the condition process may be invaluable.

Beam Pipe Scrubbing & Secondary Electron Yield

Following the hypothesis that the electron cloud has an influence on both the initial charge of a macroparticle, leading to higher loss rates, as well as on the UFO release mechanism, leading to higher UFO rates, LHC scrubbing may assist in UFO conditioning.

LHC scrubbing is a process of low intensity beam bombardment of the LHC beam pipe to deliberately induce secondary electron emissions and over time reduce the secondary electron yield (SEY). Such a process reduces the rise time of the electron cloud density to reduce heating effects and inherent unwanted pressure gradients – irrespective of UFOs. However, in theory, if either the electron cloud and or the beam pipe heat-load plays a role in macroparticle dynamics, any reduction in SEY should have a quantifiable effect on the UFO rate. As such, it was proposed that during initial Run 2 beam commissioning, several LHC fills be carried out before and after a scrubbing period with the same parameters to allow for direct comparison.

5.4 Chapter Discussion & Conclusion

Following the observation of UFO events during Run 1, it was proposed that, due to relatively low quench margins at higher energies, UFO-induced dipole quenches may be a possibility at 6.5 TeV operation. Investigation led to the construction of a numerical model, advancing on previous theoretical works, incorporating more accurate formulae for macroparticle ionisation as well as the implementation of 3–D beam field influences and BLM signal responses for both Run 1 and Run 2 beam parameters and BLM locations. Furthermore, the model allowed for parametric studies and Monte-Carlo simulations.

Advancements to previous analytical formulae were found to be within 30% of independent numerical simulations carried out using previously validated software, Garfield++. With previous models being several orders of magnitude different, this relative accuracy was deemed sufficient. Furthermore, with reasonable input parameters, the model was analytically capable of reproducing all involved BLM signals of individual Run 1 UFO events. Parametric studies found that the longitudinal and transverse starting locations as well as the initial charge of the macroparticle were the most influential parameters in terms of integrated loss rates. Furthermore, taking into account that the UFO rate has been observed to increase significantly whilst operating with 25 ns bunch spacing, it is hypothesised that the electron cloud, produced by the beam, may offset a macroparticle's initial charge, either at rest or whilst falling. This may result in an increase in time required for positive ionisation to the point of repulsion, ultimately reducing integrated losses and the probability of a magnet quench. The attractive forces as a result of this phenomenon may also explain and or contribute to the UFO release mechanism - a remaining unknown.

In an attempt to estimate the UFO-induced dipole quench probability for Run 2, Monte-Carlo simulations were first carried out to determine feasible distributions for macroparticle radii and starting locations by matching 4 TeV measured data. Following

successful reproduction of 4 TeV data, input parameters were adjusted to that of 6.5 TeV operation and, taking into account the QLU factor, the quench probability was calculated. Unfortunately, predictions vary from a manageable 0.11% to a possibly catastrophic 1.24% within plausible values the QLU factor. In light of this, as well as the unknown effect LS 1 may have had on the UFO rate, the true extent of the threat cannot yet be quantified. As a result, it was decided that Run 2 BLM thresholds for the millisecond integration time be initially set such that a number of UFO-induced dipole quenches may occur to acquire more accurate quench level data, ultimately reducing uncertainty. Subsequent simulations would then provide more accurate predictions and allow for the calculation of optimal BLM thresholds.

Investigating mitigation and conditioning strategies, it was concluded that the best course of action would be to monitor and continually optimise BLM thresholds to maximise availability, *i.e.* dump the beam (2-3 hours recovery time) rather than risk a magnet quench (8-12+ hours recovery time). The numerical model also allowed investigation into proposals such as the 'defender' bunches, in which a number of proton bunches would have a larger emittance than the rest, ionising and repelling incoming UFOs prior to interaction with the remainder of the beam. Simulations suggest that this strategy would have small, but still potential useful impact on beam losses and the feasibility of implementing such a strategy is currently being discussed. Regarding conditioning effects, with observations of an increase in UFO rate at 25 ns bunch spacing, it is proposed that the electron cloud may play an integral role in the UFO event outcome. If true, LHC scrubbing to reduce the beam pipe SEY, would have a quantifiable effect on the UFO rate and may be a viable mitigation solution. Note that, even with this proposal, it was concluded that to develop a full-proof means of accelerating the conditioning effect, more information is required regarding the UFO release mechanism, that is, what makes the dust or debris fall in the first place. As a result, a follow-up study, investigating the electron cloud influence and comparing initial Run 2 UFO statistics to model predictions was subsequently carried out. Furthermore, an experimental physics doctoral thesis, investigating the UFO release mechanism was proposed [115].

To conclude, simulations suggest that UFO-induced dipole quenches are not only a possibility, but are likely to be at the forefront of availability issues for initial 6.5 TeV operation. As a result, detailed statistical tracking and continual analysis as well as continued optimisation of BLM thresholds is advised. Furthermore, the UFO event phenomenon as a whole should be taken into account when considering any further increases to sustainable beam energy or intensity and hardware upgrades and or the design of future accelerators, particularly with regard to quench margins and a clean manufacturing processes.

Chapter 6

6.5 TeV Operation: Initial Observations

Following the successes of the interconnect consolidation and magnet training of the main dipole circuits, the LHC was, for the first time, ramped-up to sustain 6.5 TeV beam energies. This chapter will present the most recent observations relating to the operation of the LHC at 6.5 TeV, in particular, the following will discuss initial Run 2 UFO statistics as is presented in the listed publication [p1].

6.1 Energy Ramp-Up

As a direct consequence of the successes of the interconnect consolidation, system qualification and the main dipole magnet training campaigns, on April 10^{th} 2015, following hardware and beam commissioning, the LHC successfully ramped-up all circuits to support and sustain a stable, low-intensity beam at 6.5 TeV. A capture of the live LHC operations window recording this event is shown in Figure 6.1 [116].

As is the case for case for any electrical system, the LHC is only as strong as its weakest component, and in achieving this feat, it is therefore self-evident that not only are the main dipole circuits now capable of operating with 6.5 TeV beam energies, but such is the case for all involved systems.

Following the repeated success of energy ramp-up tests, the main dipole orbit correctors prior to the interaction points were activated and first of many 13 TeV collision events were recorded [117]. A 13 TeV collision example as measured by the ATLAS experiment is shown in Figure 6.2. Overall, it was clear that the system changes and upgrades carried out throughout LS 1 were a success, and that operating with 6.5 TeV beam energies was likely to be a viable option for Run 2. Note that, at this point, the long-term sustainability of operating at 6.5 TeV, *i.e.* the impact of UFO events, was still an unknown.

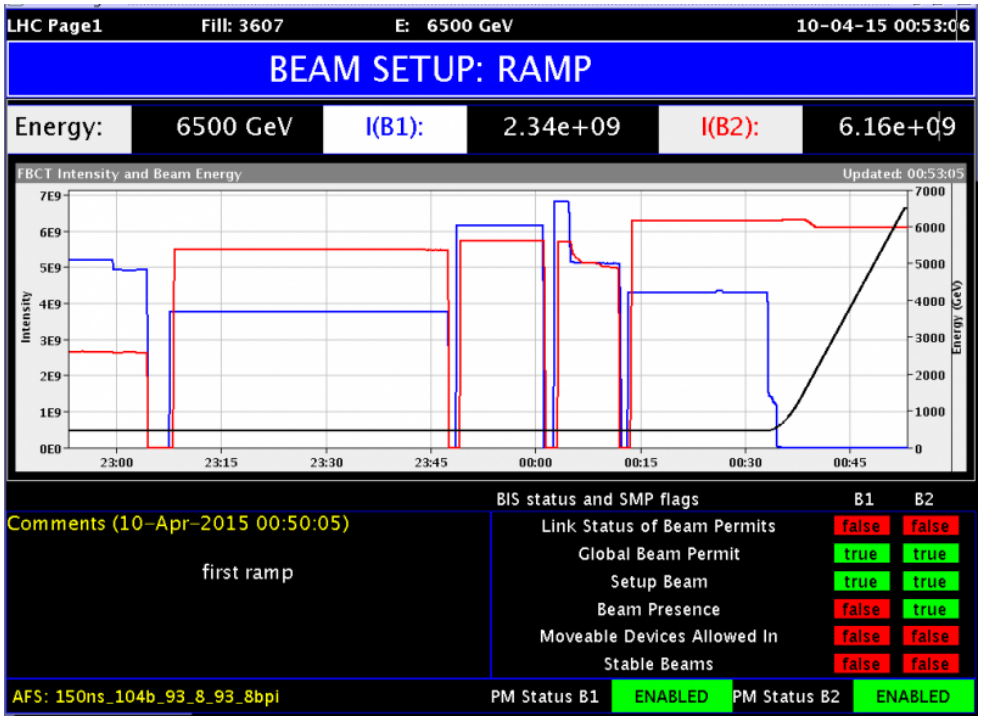

Figure 6.1: Live LHC operations plot, showing the intensity of both beams (red and blue) and the corresponding beam energy (black). Note the successful ramp up in energy to 6.5 TeV. Image courtesy of CERN [116].

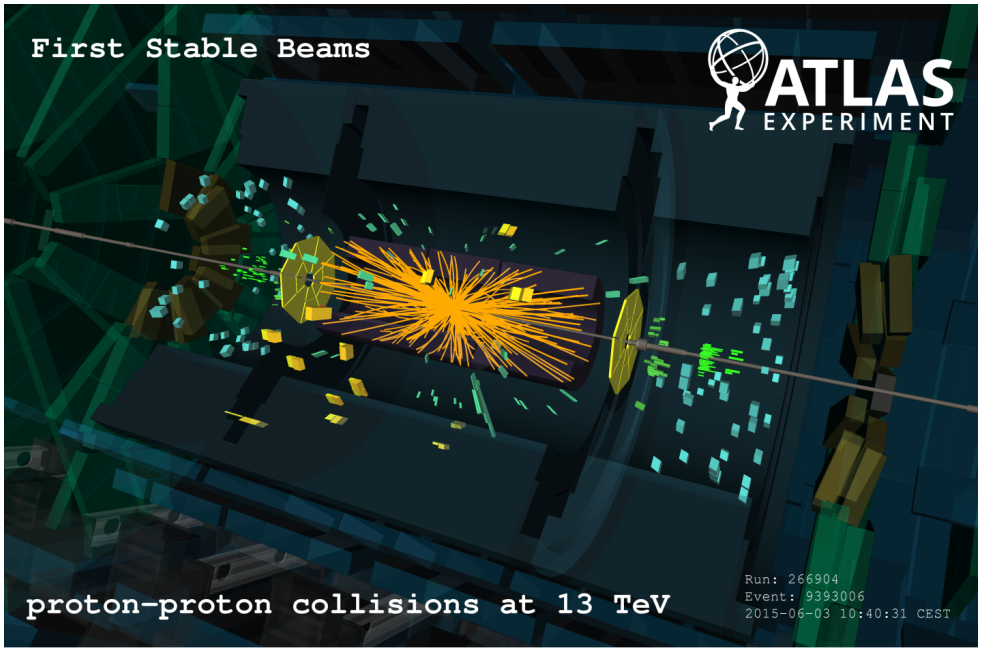

Figure 6.2: 3-D graphical reproduction of a 13 TeV collision event showing the ejected particle trajectories and relevant detector structures as measured by the ATLAS experiment, circa 2015. Image courtesy of CERN [117].

6.2 Initial Run 2 UFO Statistics

Author's note: The thesis author was no longer at CERN during the time in which the follow statistics were collected. Statistics were originally presented at Evian 2015 [118] and Chamonix 2016 [119] by B. Auchmann, Chamonix 2016 [120] by G. Iadarola and IPAC 2016 [121] by G. Papotti. The following written works are, however, the author's interpretation of the results, in particular, with regard to model predictions.

6.3 UFO Rate & Conditioning

Figure 6.3 shows the recorded UFO events for the first 6 months of 6.5 TeV operation. Data presented includes all recorded Run 2 UFO statistics from June up to and including November 2015. To discuss all relevant observations methodically, the analysis will be presented individually for 50 ns bunch spacing, 25 ns bunch spacing and the intensity ramp-up. Note that all following references to Run 1 observations is in relation to that which was presented in Chapter 3 Figure 3.9.

50 ns Bunch Spacing

Prior to the main intensity ramp-up which followed TS 2, two week periods of low-intensity operation with 50 ns bunch spacing were carried out, each followed by a 3-week-long scrubbing period to minimise the SEY and inherent electron cloud heat-load.

In looking at the absolute UFO rate for this period, it is clear that the post-LS 1 UFO rate was, as predicted, subject to 'de-conditioning', with UFO rates being notably worse than previous recordings. Recall, however, that the UFO event logging and monitoring systems were not implemented until ∼6 months after the start of Run 1 and that initial recorded UFO rates may have already underwent significantly conditioning. Nonetheless, the threat posed by UFO events is much higher than was originally considered, with UFO rates being consistently between 40 to 50 UFOs/hour. Absolute UFO rates do, however, appear to saturation even with a relatively low beam intensity. At first appearance, this would coincide with the model predictions that losses begin to saturate and subsequently decrease after ∼100-200 bunches. However, looking at the relative UFO rates (normalised to neglect the influence of beam intensity) for the same period, it is clear that relative conditioning had already begun to occur and that losses do not scale with intensity as predicted. Note that this does not necessarily contradict the model as losses do not automatically translate to recorded UFO events. Rather, the lack of correlation gives insight into what may be missing, *i.e.* external influences. The relative conditioning is likely due to the unaccounted for and remaining unknown release mechanism, *i.e.* it is therefore

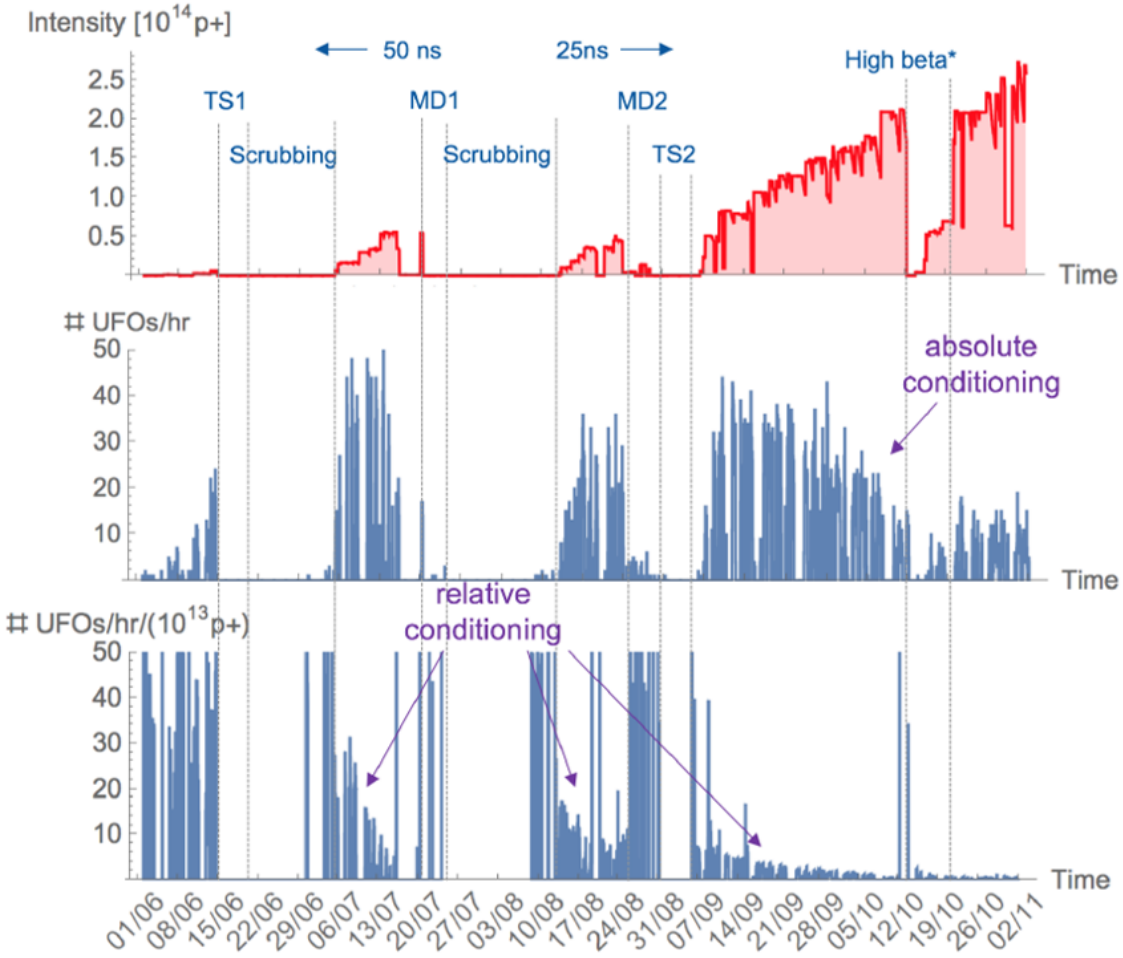

Figure 6.3: Plot showing the beam intensity ramp-up (corresponds to the number of bunches) (top) as well as both the absolute (middle) and relative (normalised with intensity) (bottom) UFO rates. Beam bunch spacing as well as Technical Shutdown (TS), Scrubbing and Machine Development (MD) periods are marked were appropriate. Plot courtesy of B. Auchmann [118].

likely that the increase in beam intensity does not only influence the macroparticles initial charge, ionisation and or interaction time, but also the probability that a macroparticle would fall in the first place. It is also of note that there was no clear evidence of relative conditioning prior to the first scrubbing period which may suggest scrubbing to be beneficial as a surface treatment if not directly responsible for the onset of relative conditioning. At such low intensities, however, no concrete conclusions can be drawn.

25 ns Bunch Spacing

Subsequent to the second scrubbing period it was concluded that the electron cloud heat load would now be low enough to operate with 25 ns bunch spacing. For reference, recall that towards the end of Run 1, a significant spike in the UFO rate whilst operating with 25 ns bunch spacing was recorded and as a result it was predicted that UFO events, while

operating with 25 ns bunch spacing, would pose a significant threat to availability.

In an attempt to quantify the specific influence that either scrubbing or a reduction in bunch spacing would have had independently, a special fill was carried out following the second scrubbing period with 50 ns and at the same intensity as a previous run. The recorded UFO rate for this special fill was ∼16 UFOs/hour compared to ∼35 UFOs/hour before scrubbing. Furthermore, in then resuming 25 ns the UFO rate then was recorded to be ∼41 UFOs/hour. This was, however, at a higher intensity. The results would suggest that scrubbing is indeed beneficial as a surface treatment, however, due to the schedule for intensity ramp-up, the request for more fills with 50 ns bunch spacing and or with 25 ns bunch spacing and a comparable intensity was rejected. As a result, there is not currently enough statistics to draw any verifiable conclusions. Nonetheless, with relative UFO rates being similar, *i.e.* within previously observed fluctuation limits, for both 50 ns and 25 ns bunch spacing, it is likely that scrubbing has had at least some impact on the outcome.

Intensity Ramp-Up

Following TS 2, the main beam intensity ramp-up began. This period is of particular interest because both relative and absolute conditioning was observed. It is also clear that the relative conditioning effect had begun to saturate. However, it is not clear to what extent the absolute conditioning effect is as a result of an increase in intensity or if the SEY and electron cloud bore any influence. Note that absolute condition would also appear to be asymptotically tending towards a UFO rate of ∼10 UFOs/hour. However, as was the case during Run 1, long-term conditioning may occur over the course of the following year.

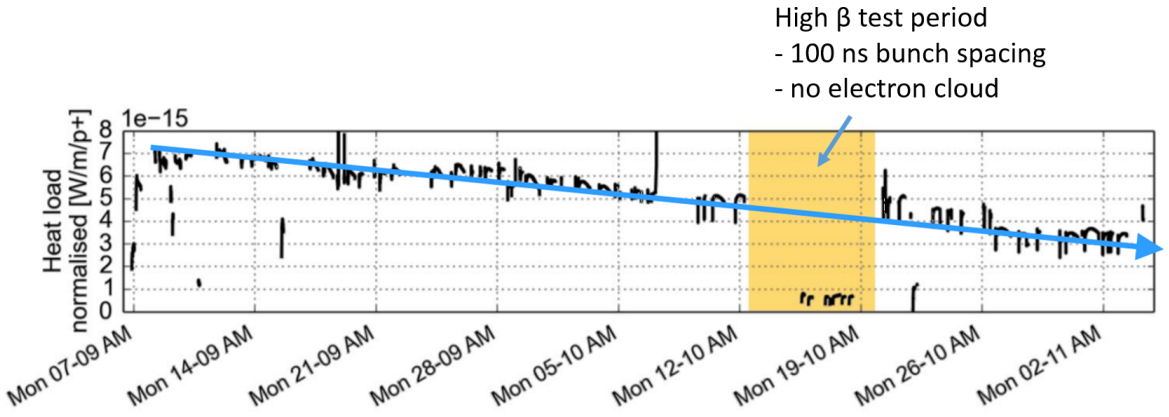

Figure 6.4: Plot of the LHC beam-pipe average heat-load throughout the intensity ramp-up period, normalised to beam intensity. It is clear that there is a continued linear reduction irrespective of intensity. A high-β test with 100 ns bunch spacing is highlighted in orange (there is little to no electron cloud build-up for this period). Plot courtesy of G. Iadarola [120].

Figure 6.4 shows the beam pipe heat-load normalised with respect to beam intensity over the main intensity ramp-up period. There is a clear gradual linear reduction in heat-load with no other changes to system operation. This gradual reduction in heat-load through periods of operation is known to be due to gradual surface treatment similar to that of a dedicated scrubbing period, reducing the SEY over time. It would normally therefore also be assumed that the electron cloud density has gradually reduced over this period, however, the electron cloud density was known to be consistently at maximum operating levels throughout this period and, as a result, any recorded reduction and or corresponding conditioning is therefore irrespective of the electron cloud density. Moreover, a specific high-β test period with 100 ns bunch spacing and little to no electron cloud was carried out – highlighted in orange. During this period, several fills were carried out with comparable intensity levels and the recorded UFO rates did not change. It can therefore be concluded that the density of the electron cloud has little to no effect on the measurable UFO rate or conditioning. As a result, the correlation between normalised heat-load and conditioning is more likely to be due to other proposed theories such as that which relates cryo-condensation and cryo-sorption. Further research is required.

6.3.1 Beam Dumps & Magnet Quenches

On September 20th 2015, near the start of the intensity ramp up, the first ever UFO-induced magnet quench was recorded, see Figure 6.5.

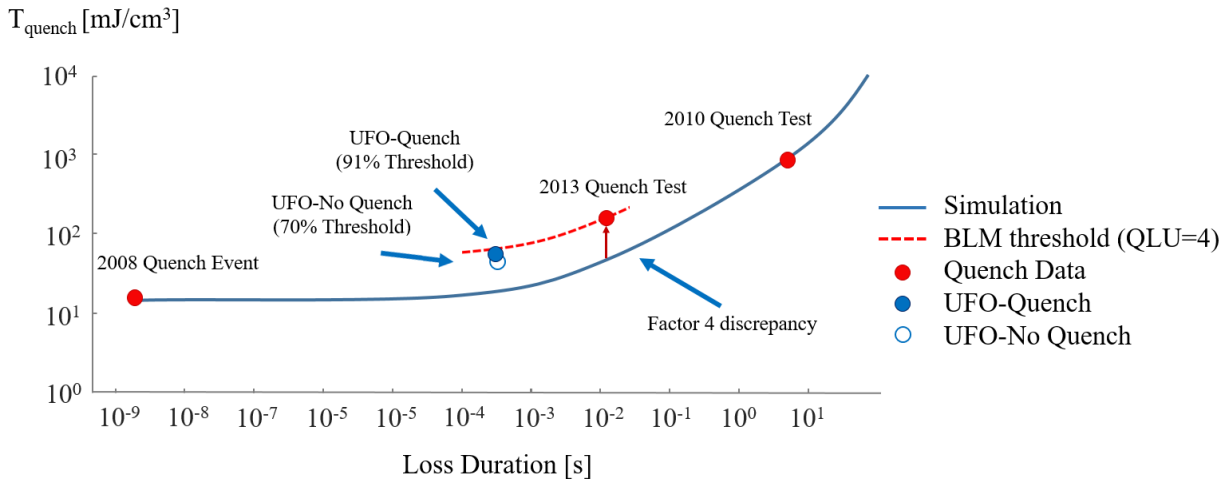

Figure 6.5: Plot of the energy deposition required to quench as a function of loss duration, comparing simulation to measured data. The recent UFO-induced quench data is included. Data courtesy of B. Auchmann [119].

The UFO event resulted in a BLM signal corresponding to 91% of the BLM threshold (QLU=4). Furthermore, subsequent confirmed UFO events were recorded to have a

corresponding signal of 70% of the implemented BLM threshold (QLU=4) and did *not* result in a magnet quench. It was therefore evident that initial quench level estimates from simulations were too conservative and that the implementation of a QLU factor for the BLM thresholds was not only justified, but necessary. This would also suggest that, in light of fast conditioning, the long-term impact of UFO events is likely to be more in-line with progressive estimates, *i.e.* manageable.

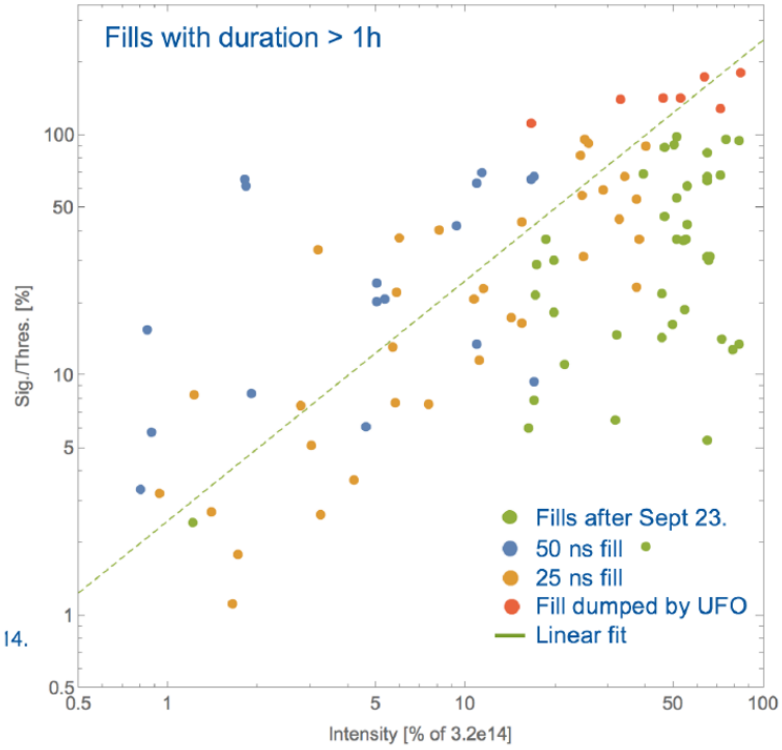

Figure 6.6: Log-log plot of BLM signal as a percentage of the dump threshold as a function of beam intensity as a percentage of nominal, showing the peak BLM signal during all recorded UFO events with a fill duration >1 hour. Events recorded after September 23rd, the onset of absolute conditioning, are marked in green. Events which resulted in a beam dump or magnet quench are marked in red. Linear fit shown is related solely to recorded event prior to September 23rd. Plot courtesy of B. Auchmann [118].

Figure 6.6 plots the peak BLM signal across all recorded UFOs from each fill that exceeded 1 hour in duration. It is clear that, prior to absolute conditioning (end of September), linear extrapolation of recorded UFOs would see the system being limited in operation to ∼50% of the nominal number of bunches if it were to remain at 6.5 TeV. Fortunately, just as numerous UFO events were resulting in beam dumps and magnet quenches, 8 dumps between September 20th and October 5th alone, absolute conditioning began and the average peak BLM signals also reduced. As a result, it can now be concluded that 6.5 TeV operation remains a viable option. However, assuming conditioning is indeed asymptotic UFOs are still likely to be of constant threat to availability for the

duration of 6.5 TeV operation and may, in fact, be a physical limitation if the LHC is ever commissioned to the design nominal of 7 TeV.

6.3.2 UFO Distribution

Another unknown was the longitudinal UFO distribution along a typical Arc cell – assumed homogeneous throughout in the model. Figure 6.7 compares the distribution of UFO events as would be predicted by the model and that which have been recorded. Recorded data has been processed such that factor of 3 difference in the BLM response of the BLMs surrounding the quadrupoles has been taken into account.

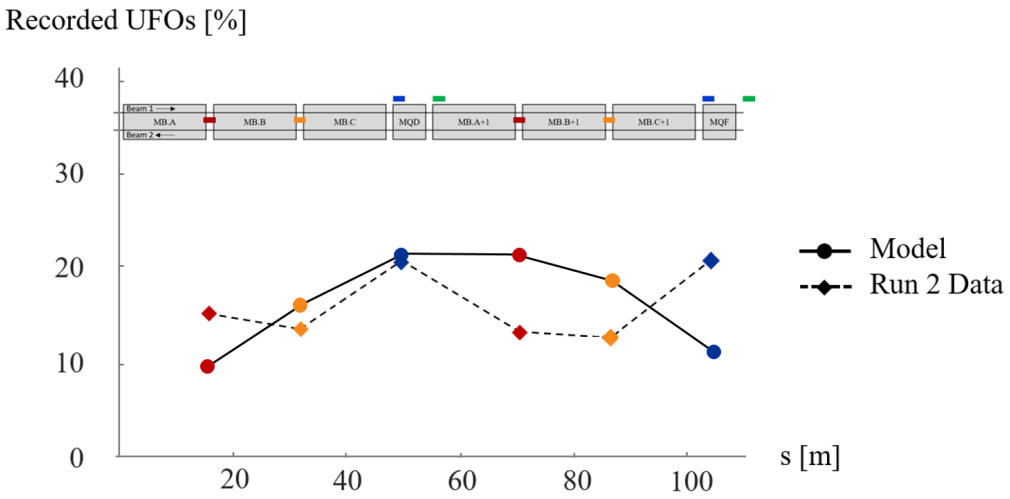

Figure 6.7: Comparison between model predictions and recorded events relating to the UFO distribution as would recorded by the BLMs along a typical Arc cell. Recorded data shows a peak in UFO events surrounding defocusing quadrupoles which is not present in model predictions. Data courtesy of B. Auchmann [119].

Although similarly distributed around the defocusing quadrupole BLMs, recorded UFO events do not follow the same beam-shape-dominated relation as the model predicts, with the biggest difference being a notably higher percentage of UFO events being recorded following the focusing quadrupoles, *i.e.* where the beam is predominately vertical. It is also unlikely that the physical distribution of dust in the LHC is more likely to accumulate around such magnets as the environment within the beam pipe does not change. It is therefore assumed that the model has not taken an external influence into account. With the transverse beam profile at these locations being predominately vertical, *i.e.* the beam is closest to the beam screen ceiling as it can get, the unaccounted for external influence is most likely to be the release mechanism. If this is the case, a predominately vertical beam is therefore more likely to interact with a macroparticle such that it falls, either by means of offsetting the initial charge or heat transfer (further backing the theory that

cryo-condensation and or cryo-sorption could be influencing factors). More research is required.

6.4 Chapter Summary & Outlook

As a direct consequence of the successes of the interconnect consolidation, qualification and magnet training campaigns, the LHC was found to be repeatedly capable of successfully ramping up in beam energy to support and sustain 6.5 TeV proton beams, exceeding previous records. Not long after, the first 13 TeV collisions were being carried out with sufficient stability and Run 2 of LHC operation commenced with an official beam energy of 6.5 TeV.

In monitoring initial Run 2 UFO statistics a number of notable trends arose. Firstly, the initial absolute UFO rate was much higher ($\sim$40-50 UFOs/hour) than was previously recorded during Run 1 ($\sim$10-12 UFOs/hour), even when considering the increase in BLM coverage (would only account for the rate increasing by a factor of 2). This likely confirms the hypothesis that, as UFO event logging was not implemented until the start of 2011, after $\sim$7 months initial operation, significant conditioning may already have occurred in 2010. Nonetheless, looking at relative UFO rates, *i.e.* irrespective of beam intensity, for the first 3 weeks, it was soon apparent that a relative conditioning effect had already begun.

Throughout initial operation with 50 ns bunch spacing, two 3-week-long scrubbing periods were carried out. Relative UFO rates subsequently showed that this is likely to have been beneficial as a surface treatment as relative conditioning was shown to occur immediately following these periods, with no evidence to suggest this would have happened otherwise (no relative condition in the period prior to scrubbing).

Operation with 25 ns bunch spacing commenced following the second scrubbing period. Most notably, the absolute UFO rate was not subject to a significant increase as was observed in Run 1 and, in fact, relative conditioning proceeded to continue. Again, this is likely to be as a result of the long scrubbing periods. At this point it was deemed safe to proceed with the intensity ramp-up.

Throughout the intensity ramp-up, absolute UFO rates also began to reduce even with relative conditioning beginning to saturate. The underlying cause of absolute conditioning remains an unknown, however, in looking at the relative heat-load (irrespective of intensity), which correlates to the SEY as well as the electron cloud density, there is also a notable reduction across the same time period. This is almost certainly as a result of a

reduction in SEY, however, a specific high-β test proved that UFO rates remain consistent even without the presence of an electron cloud, all but ruling out the electron cloud density a potential contributor to the recorded UFO rate.

At this point it is noted that, prior to the onset of absolute conditioning, peak BLM signals were high enough and several UFO-induced magnet quenches were recorded, confirming predictions. Furthermore, the quench level was at 91% of the set threshold (QLU=4) which would suggest that previous simulations were too conservative and that both the real quench level and therefore impact predictions are more in-line with that of the existing quench test data.

Looking specifically at the peak BLM signals, it is clear that if it were not for absolute conditioning, 6.5 TeV operation would likely have been limited to only 50% nominal intensity. Fortunately, absolute conditioning did commence and was also relatively fast, leading to the majority of UFOs being far below the threshold throughout the following month. As a result, it was concluded that 6.5 TeV operation would continue to be a viable option and that operation of the LHC throughout 2016 would continue at this energy level.

Chapter 7

Thesis Summary & Recommendations for Further Work

This thesis investigated remaining issues which may have prevented operation of the main dipole circuits at near-nominal energy levels. Research led to the hypothesis that, even with observed issues, operation with at least 6.5 TeV beam energies was still a viable option. This final chapter will recall the hypothesis and established prerequisites, discussing the outcome of related projects. The chapter, and thesis, will then conclude with the author's recommendations for further work.

7.1 Thesis Summary

Recall the following hypothesis and established prerequisites:

Thesis Hypothesis

Following successful circuit qualification and dipole magnet training, the LHC's main dipole magnet circuits can safely support and sustain near-nominal proton beam energies of at least 6.5 TeV.

Prerequisites

1 To ensure magnet training can be carried out both systematically and without risk, the development and subsequent implementation of a new method of superconducting magnet bypass qualification is required.
2 Magnet quenches following dust-to-beam interaction scattering are thought to be a possibility at 6.5 TeV operation. Further understanding of the phenomena and research into potential mitigation strategies is therefore required.

In order to address each of the above statements, a series of goals were set out. The following discusses the work involved for each of the set goals individually.

Addressing 1st Prerequisite

- Develop and investigate both the risk and feasibility of a global superconducting magnet circuit continuity qualification method.

Advancing on the work of a previously proposed test method, research led to the development of the Copper Stabiliser Continuity Measurement. The procedure would be capable of providing conclusive evidence that all interconnections, contacts and potential weak sections of the superconducting busbars or magnet bypasses could withstand a full current decay similar to that which would occur following a near-nominal quench event.

To establish the potential risk and feasibility of the developed method, a CSCM simulation program was constructed and a type test was carried out in Sector 23, pre-LS 1. Due to the known presence of interconnect discontinuities at the time, this allowed for thermal runaways to be induced in a controlled manner for the purposes of qualifying detection and protection systems. Results were a success, with multiple thermal runaways being induced as well as successfully detected and protected against.

- Carry out the developed qualification method on all main dipole circuits prior to magnet training.

Following the successes of the type test, a full qualification campaign was scheduled for all main dipole circuits. All circuits were eventually qualified for 6.5 TeV operation, with each circuit eventually being subject to a full current decay from 11.08 kA, without measurable damage.

Alongside the devised full qualification test program, 400 A resistance measurements were carried out to monitor resistance trends in the interconnects and magnet bypasses. As a testament to the successes of the interconnect consolidation campaign, the interconnect resistances remained relatively low and consistent throughout. However, magnet bypass resistances were notably more unpredictable, with diode-to-heat-sink contacts increasing in resistance by up to a factor of 10 on several occasions. Fortunately, a stepped powering cycle was devised which was capable of reversing this effect in a controlled manner.

Overall, the full qualification campaign was considered a success and magnet training was subsequently carried out to 6.5 TeV equivalent current levels.

- Upon successful qualification of all circuits, carry out magnet training as planned (up to 6.5 TeV equivalent current levels).

Subsequent to all main circuits being qualified to ensure safe operation in the event of a quench at full current, magnet training was carried out for the first time since the 2008 incident. Training was successful, eventually resulting in all circuits' magnets being qualified to support 6.5 TeV beam energies, *i.e.* trained to 11.08 kA. Most notably, the number of required training quenches was almost double that of earlier predictions, in particular, this was as a result of Firm 3 magnets requiring significantly more training than expected.

Conclusion

Overall, the CSCM was proven to be both capable and repeatable in providing conclusive evidence for the qualification and safe operation of each main dipole circuit to 11.08 kA. Subsequently, each circuit underwent magnet training to the same current level without issues and as a direct consequence the LHC was now capable of supporting 6.5 TeV beam energies.

Addressing 2nd Prerequisite

- Develop a numerical model capable of simulating dust-to-beam interactions as well as quantifying the resultant particle shower energy deposition in adjacent magnets.

Progressing on previous theoretical works, a numerical model was constructed capable of simulating dust-to-beam interactions as well as calculating the corresponding signals as would be read by adjacent BLMs. The model was proven capable of accurately reproducing both single 4 TeV UFO events as well as, via Monte-Carlo simulation, the entire 4 TeV data set. All simulations were run with feasible input parameters.

- Use the model to determine the probability of a 6.5 TeV UFO event resulting in a magnet quench, quantifying probable losses to system availability.

Following Monte-Carlo simulation of 4 TeV UFO events, it was possible, using the same input distributions, to extrapolate to 6.5 TeV. Simulations at 6.5 TeV resulted in the discovery that UFO-induced magnet quenches were, in fact, a possibility. However, determining the true extent of the threat to availability was restricted due to the lack of available data relating to magnet quench limits. This quench level uncertainty varied predictions on the likelihood of magnet quenches by a factor of 11, varying from a manageable 0.11% of recorded UFOs resulting a magnet quench to a near-catastrophic 1.24% – at this stage an exact number could not be given as there was little to no information regarding

the physical quench level or Run 2 UFO rates.

- Investigate and propose mitigation strategies.

With simulations predicting UFO events to have a significant impact on system availability, it was important to investigate plausible mitigation strategies. An interesting idea regarding potential defender bunches was proposed, however, simulations suggest that such bunches would have minimal impact on the outcome. As such, for initial Run 2 operation it was concluded that the best course of action would be to bias the BLM thresholds, allowing for UFO-induced quenches to occur, after which, having obtained more accurate values for magnet quench limits and the UFO rate, thresholds can be optimised to balance the trade-off in availability between magnet quenches and beam-dumps. Note that a more permanent solution is still desired and an experimental thesis was proposed to further understanding of the UFO release mechanism.

Initial Run 2 Observations

Towards the end of commissioning, the LHC was shown to be repeatedly capable of successfully ramping up in beam energy to support and sustain 6.5 TeV proton beams, including being stable with 13 TeV collisions. However, not long after the onset of 6.5 TeV operation, the first UFO-induced magnet quench was recorded. Furthermore, initial absolute UFO rates were recorded to be 4 times greater than in 2011, increasing the threat to availability. In fact, as of the start of October it was looking like the system may even be limited to half the nominal beam intensity for 6.5 TeV operation. Fortunately, however, just as the intensity ramp-up was resulting in UFO losses frequently resulting in beam dumps, absolute conditioning began to occur. The result was that the majority of UFO events in the following month, even with the intensity continuing to increase, produced signals that were far below the BLM threshold and 6.5 TeV operation was allowed to continue. There are, however, to date not enough statistics to draw conclusions on the long-term availability of 6.5 TeV operation.

Conclusion

To conclude, the devised model was capable of simulating 4 TeV UFO events and allowed for outcome predictions of 6.5 TeV equivalent scenarios. Research led to numerous interesting insights into the involved phenomena, in particular, with regard to UFO dynamics and the probability of UFO-induced magnet quenches. Initial Run 2 observations not only confirmed the possibility of UFO-induced magnet quenches but also confirmed thoughts that UFO events will be at the forefront of availability issues for 6.5 TeV operation and beyond. Nonetheless, as was also predicted, absolute conditioning was eventually observed

and the frequency of beam dumps and magnet quenches reduced to within manageable numbers. Overall, operating the LHC with 6.5 TeV beam energies therefore remains a viable option and the thesis hypothesis is considered confirmed (albeit with an element of luck involved).

[This space has been intentionally left blank]

7.2 Recommendations for Further Work

The research which has been undertaken for this thesis has highlighted a number of topics for which further research could be beneficial. The following lists the author's recommendations for further work.

6.5 TeV UFO Impact on Availability

- UFO events were both predicted and shown to have a significant impact on LHC availability. It is suggested that research in the immediate future not only involves, but focuses on the following.
  - Continued collection and subsequent analysis of 6.5 TeV UFO event statistics, focusing on the optimisation of BLM thresholds and the trade-off in availability been magnet quenches and beam dumps.
  - Investigate the most influential factors relating to the UFO release mechanism, either by experiment or modelling.
  - Research and develop mitigation strategies and methods for accelerating conditioning (such research is also likely be a deciding factor in whether or not the LHC will further increase operation to the design-nominal beam energy of 7 TeV).
  - Update the existing modelling to incorporate external influences such as the release mechanism, heat-load, *etc.* Use the model to make long-term availability predictions.

Standardisation of CSCM

- The CSCM was proven to be capable in providing conclusive evidence for the qualification and safe operation of a superconducting magnet string to a specific current level. The method should be standardised and all possible applications should be investigated.
  - Standardise the method to be repeatable by non-experts.
  - Consider the development a full 'CSCM mode' for the magnet protection system, which is easily switched too for testing.
  - Consider CSCM for the qualification of other major circuits, both existing and in development (for example, High-Luminosity LHC upgrades).

7 TeV (Design-Nominal) Operation

- With the overall success of supporting and sustaining 6.5 TeV beam energies, there is no reason not to investigate the feasibility of further increasing the maximum sustainable beam energy to the design-nominal value of 7 TeV.
  - Using new training data, provide more accurate predictions for magnet training quench numbers and estimate the required time.
  - Consider circuit qualification to 7 TeV via CSCM.
  - Model and investigate the 7 TeV UFO impact, taking into account yet higher beam energies and yet further reduced magnet quench margins. Also take into account accelerated conditioning, if methods for doing so are developed.

References

[1] D. W. Kerst *et al.*, "Attainment of Very High Energy by Means of Intersecting Beams of Particles," *Phys. Rev.*, vol. 102, pp. 590–591, Apr. 1956. DOI: 10.1103/PhysRev.102.590.

[2] R. Oerter, *The Theory of Almost Everything: The Standard Model, the Unsung Triumph of Modern Physics.* Plume, 2006, ISBN: 978-0-452-28786-0. DOI: 10.1063/1.2337829.

[3] S. L. Glashow, "Partial-symmetries of weak interactions," *Nuclear Physics*, vol. 22, pp. 579–588, 1961. DOI: 10.1016/0029-5582(61)90469-2.

[4] F. Englert and R. Brout, "Broken Symmetry and the Mass of Gauge Vector Mesons," *Phys. Rev. Lett.*, vol. 13, pp. 321–323, 1964. DOI: 10.1103/PhysRevLett.13.321.

[5] P. W. Higgs, "Broken Symmetries and the Masses of Gauge Bosons," *Phys. Rev. Lett.*, vol. 13, pp. 508–509, 1964. DOI: 10.1103/PhysRevLett.13.508.

[6] G. S. Guralnik *et al.*, "Global Conservation Laws and Massless Particles," *Phys. Rev. Lett.*, vol. 13, pp. 585–587, 1964. DOI: 10.1103/PhysRevLett.13.585.

[7] S. Weinberg, "A Model of Leptons," *Phys. Rev. Lett.*, vol. 19, pp. 1264–1266, 1967. DOI: 10.1103/PhysRevLett.19.1264.

[8] ATLAS Collaboration, "Observation of a new particle in the search for the Standard Model Higgs boson with the ATLAS detector at the LHC," *Physics Letters B*, vol. 716, pp. 1–29, 2012. DOI: 10.1016/j.physletb.2012.08.020.

[9] CMS Collaboration, "Observation of a new boson at a mass of 125 GeV with the CMS experiment at the LHC," *Physics Letters B*, vol. 716, pp. 30–61, 2012. DOI: 10.1016/j.physletb.2012.08.021.

[10] P. Langacker, "The Standard Model and Beyond," *AIP Conference Proceedings*, vol. 150, pp. 142–164, 1986. DOI: 10.1063/1.36182.

[11] V. Trimble, "Existence and Nature of Dark Matter in the Universe," *Annual Review of Astronomy and Astrophysics*, vol. 25, pp. 425–472, 1987. DOI: 10.1146/annurev.aa.25.090187.002233.

[12] P. J. E. Peebles and B. Ratra, "The cosmological constant and dark energy," *Rev. Mod. Phys.*, vol. 75, pp. 559–606, 2003. DOI: 10.1103/RevModPhys.75.559.

[13] H. Baer and X. Tata, *Weak Scale Supersymmetry.* Cambridge, 2006, ISBN: 978-0-521-85786-4. DOI: 10.2277/0521857864.

[14] E. Eriksen and K. Vøyenli, *The Classical and Relativistic Concepts of Mass*, 1. 1976, vol. 6, pp. 115–124. DOI: 10.1007/BF00708670.

[15] E. Wilson, *An Introduction to Particle Accelerators*. 2001, ISBN: 978-0-198-50829-8. DOI: 10.1093/acprof:oso/9780198508298.001.0001.

[16] O. S. Brüning *et al.*, *LHC Design Report*. CERN, 2004, vol. 1, ISBN: 978-9-290-83239-3. DOI: 10.5170/CERN-2004-003-V-1.

[17] J. Bardeen *et al.*, "Theory of Superconductivity," *Phys. Rev.*, vol. 108, pp. 1175–1204, 5 Dec. 1957. DOI: 10.1103/PhysRev.108.1175.

[18] R. Schmidt *et al.*, "Protection of the CERN Large Hadron Collider," *New Journal of Physics*, vol. 8, p. 290, 2006. DOI: 10.1088/1367-2630/8/11/290.

[19] A. P. Verweij, "Electrodynamics of Superconducting Cables in Accelerator Magnets," *PhD Thesis*, 1995. [Online]. Available: http://cern-verweij.web.cern.ch/cern-verweij/title-contents.pdf.

[20] K.-H. Mess *et al.*, *Superconducting Accelerator Magnets*. World Scientific, 1996, ISBN: 978-9-810-22790-6. DOI: 10.1142/9789814261029.

[21] M. Spakinsi *et al.*, "Beam Induced Quenches of LHC Magnets," *Proceedings of IPAC13*, pp. 52–55, 2013. [Online]. Available: http://inspirehep.net/record/1314083.

[22] O. S. Brüning *et al.*, *LHC Design Report*. CERN, 2004, vol. 2, ISBN: 978-9-290-83224-9. DOI: 10.5170/CERN-2004-003-V-2.

[23] B. Tuchming, "Tevatron Higgs results," *EPJ Web Conf.*, vol. 60, pp. 1–6, 2013. DOI: 10.1051/epjconf/20136002003.

[24] P. I. Kalmus, "High-energy physics: W and Z particles from CERN," *Nature*, vol. 304, 1983. DOI: 10.1038/304686b0.

[25] L. Lyons, "Discovering the Significance of 5 sigma," *ArXiv.org*, 2013. [Online]. Available: https://arxiv.org/abs/1310.1284.

[26] ATLAS Collaboration, "The ATLAS experiment at the CERN Large Hadron Collider," *IOP Journal of Instrumentation*, vol. 3, 2008. DOI: 10.1088/1748-0221/3/08/S08003.

[27] ALICE Collaboration, "The ALICE experiment at the CERN LHC," *IOP Journal of Instrumentation*, vol. 3, 2008. DOI: 10.1088/1748-0221/3/08/S08002.

[28] CMS Collaboration, "The CMS experiment at the CERN LHC," *IOP Journal of Instrumentation*, vol. 3, 2008. DOI: 10.1088/1748-0221/3/08/S08004.

[29] LHCb Collaboration, "The LHCb Detector at the LHC," *IOP Journal of Instrumentation*, vol. 3, 2008. DOI: 10.1088/1748-0221/3/08/S08005.

[30] ATLAS Collaboration, "Collision event at 7 TeV," *ATLAS Experiment*, 2010. [Online]. Available: http://atlas.web.cern.ch/Atlas/public/EVTDISPLAY/atlas2010-vp1-152166-316199.png.

[31] H. Satz, "The Quark-Gluon Plasma A Short Introduction," *ICPAQGP-2010*, vol. 862, 2011. DOI: 10.1016/j.nuclphysa.2011.05.014.

[32] WMAP Science Team, "Cosmology: The Study of the Universe," *NASA's Wilkinson Microwave Anisotropy Probe*, 2011. [Online]. Available: http://map.gsfc.nasa.gov/universe/WMAP_Universe.pdf.

[33] I. Bigi, "CP violation - A probe of Nature's Grand Design," *Surveys in High-Energy Physics*, vol. 12, 1998. DOI: 10.1080/01422419808228861.

[34] P. Kaptiza, "Viscosity of Liquid Helium below the Lambda-Point," *Nature*, vol. 141, p. 74, 1938. DOI: 10.5170/CERN-2004-003-V-1.

[35] P. Cruikshank *et al.*, "Mechanical design aspects of the LHC beam screen," *Proceeding of PAC1997*, vol. 3, pp. 3586–3588, 1997. DOI: 10.1109/PAC.1997.753282.

[36] P. Loïez *et al.*, "Beam screens for the LHC beam pipes. Les conduits pour les faisceaux LHC," May 1997, [Online]. Available: http://cds.cern.ch/record/39110.

[37] H. Wiedemann, *Particle Accelerators Physics*. 1993, ISBN: 978-3-540-49045-6. DOI: 10.1007/978-3-540-49045-6.

[38] G. Rumolo, "LHC Experience with Different Bunch Spacings in 2011 (25, 50 and 75ns)," *Proceeding of Chamonix 2012 Workshop on LHC Performance*, pp. 88–89, 2012. DOI: 10.5170/CERN-2012-006.89.

[39] O. Grobner, "Bunch Induced Multipactoring," *Proceedings of HEACC 1977*, pp. 277–282, 1977. [Online]. Available: http://inspirehep.net/record/127252.

[40] M. Lamont, "Status of the LHC," *IOP J. Phys.: Conf. Ser. 455*, p. 012 001, 2013. DOI: 10.1088/1742-6596/455/1/012001.

[41] A. Verweij, "Performance of the Main Dipole Magnet Circuits of the LHC during Commissioning," in *Proceedings of EPAC2008*, 2008, pp. 2473–2475. [Online]. Available: https://cds.cern.ch/record/1124069.

[42] D. Nisbet and H. Thiesen, "Review of the Initial Phases of the LHC Power Converter Commissioning," in *Proceedings of EPAC2008*, 2008, pp. 3670–3672. [Online]. Available: https://cds.cern.ch/record/1123729.

[43] G. Gumer, "Description of the CERN Electrical Network," *CERN EDMS*, vol. 559013, 2005. [Online]. Available: https://edms.cern.ch/document/559013/3.

[44] J. Ahlback *et al.*, "Construction of a 56 mm aperture high-field twin-aperture superconducting dipole model magnet," *IEEE Transactions on Magnetics*, vol. 32, pp. 2097–2100, 4 1996. DOI: 10.1109/20.508576.

[45] L. Rossi, "LHC dipole production begins to take off," *CERN Courier*, 2004. [Online]. Available: http://cerncourier.com/cws/article/cern/29011.

[46] O. Napoly, "Beam Delivery System and Beam-Beam Effects," 2009, [Online]. Available: https://agenda.linearcollider.org/event/3475/session/3/contribution/21/material/paper/2.pdf.

[47] L. Rossi and L. Bottura, "Superconducting Magnets for Particle Accelerators," *Reviews of Accelerator Science and Technology*, vol. 5, pp. 25–50, 2012. DOI: 10.1142/S1793626812300034.

[48] B. Auchmann and S. Russenchuck, "Coil End Design for Superconducting Magnets Applying Differential Geometry Methods," *IEEE Transactions on Magnetics*, vol. 40, pp. 1208–1211, 2 2004. DOI: 10.1109/TMAG.2004.824710.

[49] J.-L. Caron, "Beam screens for the LHC beam pipes. Les conduits pour les faisceaux LHC," Mar. 1998, [Online]. Available: https://cds.cern.ch/record/841511.

[50] A. P. Verweij and L. Buchsbaum, "Experimental results of current distribution in rutherford-type LHC cables," *Cryogenics*, vol. 40, pp. 663–670, 2 2001. DOI: 10.1016/S0011-2275(01)00023-6.

[51] CERN, "Superconducting cable," 2005, [Online]. Available: https://cern.ch/lhc-machine-outreach/components/cable.htm.

[52] L. Belova *et al.*, "Design and manufacture of the superconducting bus-bars for the LHC main magnets," *IEEE Transactions on Applied Superconductivity*, vol. 12, pp. 1305–1309, 1 2002. DOI: 10.1109/TASC.2002.1018642.

[53] A. Ballarino *et al.*, "13000-A HTS current leads for the LHC accelerator: From conceptual design to prototype validation," *CERN-LHC-PROJECT-REPORT*, vol. 696, 2004. [Online]. Available: http://cds.cern.ch/record/729433.

[54] CERN, "A 13 kA current lead, measuring 1.5 m in length. une amene de courant 13 kA. la pice mesure 1,5 m de longueur.," May 2004, [Online]. Available: https://cds.cern.ch/record/841511.

[55] K. Dahlerup-Petersen *et al.*, "Energy extraction for the LHC superconducting circuits," *Proceeding of PAC2001*, pp. 3448–3450, 2001. DOI: 10.1109/PAC.2001.988140.

[56] K. Dahlerup-Petersen, "World of protection for 13kA circuits at LHC," Feb. 2009, [Online]. Available: http://te-mpe-cp.web.cern.ch/te-mpe-cp/index.php?sub=goals.

[57] Y. Iwasa, *Case Studies in Superconducting Magnets: Design and Operational Issues*. Springer, 2009, ISBN: 978-0-387-09799-2. DOI: 10.1007/b112047.

[58] L. Coull *et al.*, "LHC Magnet Quench Protection System," *IEEE Transactions on Magnetics*, vol. 30, pp. 1742–1745, 4 1994. DOI: 10.1109/20.305593.

[59] A. V. Fernndez and F. Rodriguez-Mateos, "Reliability of the quench protection system for the LHC superconducting elements," *Nucl.Instrum.Meth.A*, vol. 525, pp. 439–446, 3 2004. DOI: 10.1016/j.nima.2004.01.081.

[60] F. Rodriguez-Mateos *et al.*, "Quench heater studies for the LHC magnets," *Proceedings of PAC 2001*, vol. 5, pp. 3451–3453, 2001. DOI: 10.1109/PAC.2001.9881416.

[61] B. Dehning *et al.*, "Beam loss monitoring system for the LHC," *Proceedings of IEEE NSSMIC 2005*, pp. 1052–1056, 2005. DOI: 10.1109/NSSMIC.2005.1596433.

[62] E. B. Holzer *et al.*, "Beam Loss Monitoring for LHC Machine Protection," *Physics Procedia*, vol. 37, pp. 2055–2062, 2011. DOI: 10.1016/j.phpro.2012.04.110.

[63] M. Stockner *et al.*, "Measurements and Simulations of Ionization Chamber Signals in Mixed Radiation Fields for the LHC BLM System," *Proceedings of IEEE NSS/MIC 2006*, vol. 3, pp. 1342–1345, 2006. DOI: 10.1109/NSSMIC.2006.354151.

[64] M. Bajko *et al.*, "Report of the Task Force on the Incident of 19th September 2008 at the LHC," *CERN-LHC-PROJECT-REPORT*, vol. 1168, 2008. [Online]. Available: http://cds.cern.ch/record/1168025.

[65] A. P. Verweij, "Thermal Runaway of the 13 kA Busbar Joints in the LHC," *IEEE Transactions on Applied Superconductivity*, vol. 20, pp. 2155–2159, 2010. DOI: 10.1109/TASC.2010.2043350.

[66] R. Denz *et al.*, "Upgrade of the Protection System for Superconducting Circuits in the LHC," *Proceedings of PAC09*, vol. 1168, pp. 244–246, 2009. DOI: 10.1109/TASC.2013.2283101.

[67] M. Brice, "Damage of the LHC magnets in sector 3-4 of the LHC. dommages sur les aimants dans le secteur 3-4 du LHC.," Nov. 2009, [Online]. Available: https://cds.cern.ch/record/1185822.

[68] A. P. Verweij *et al.*, "Consolidation of the 13 kA Interconnects in the LHC for Operation at 7 TeV," *IEEE Transactions on Applied Superconductivity*, vol. 21, pp. 2376–2379, 2011. DOI: 10.1109/TASC.2010.2083622.

[69] L. Cifarelli and L. Maritato, *Superconducting Materials for High Energy Colliders*. World Scientific, 1999, ISBN: 978-9-810-24319-7. DOI: 10.1142/4435.

[70] N. Bhattacharya, "Degradation and training in superconducting magnet: Revisited," *Physica A: Statistical Mechanics and its Applications*, vol. 270, pp. 226–236, 1999. DOI: 10.1016/S0378-4371(99)00150-8.

[71] P. P. Granieri *et al.*, "SlipStick Mechanism in Training the Superconducting Magnets in the Large Hadron Collider," *IEEE Transactions on Applied Superconductivity*, vol. 21, pp. 3555–3560, 2011. DOI: 10.1109/TASC.2011.2162727.

[72] P. Pugnat and A. Siemko, "Review of Quench Performance of LHC Main Superconducting Magnets," *IEEE Transactions on Applied Superconductivity*, vol. 17, pp. 1091–1096, 2 2007. DOI: 10.1109/TASC.2007.898501.

[73] C. Lorin *et al.*, "Predicting the Quench Behavior of the LHC Dipoles during Commissioning," *IEEE Transactions on Applied Superconductivity*, vol. 20, pp. 135–139, 3 2010. DOI: 10.1109/TASC.2010.2043076.

[74] A. P. Verweij, "Training the Dipoles," *Proceedings of Chamonix 2009 workshop on LHC Performance*, pp. 25–28, 2009. [Online]. Available: https://cds.cern.ch/record/1172379.

[75] M. Bednarek, "ELQA during LS1 - the unexpected activities," Jun. 2015, [Online]. Available: https://indico.cern.ch/event/396125/.

[76] A. Chao, *Physics of Collective Beam Instabilities in High Energy Accelerators*. Wiley, 1993, ISBN: 978-0-471-55184-3. [Online]. Available: http://www.slac.stanford.edu/~achao/wileybook.html.

[77] T. Baer *et al.*, "UFOs in the LHC," *Proceedings of IPAC2011*, pp. 1347–1349, 2011. [Online]. Available: http://www-linac.kek.jp/mirror/IPAC2011/papers/tupc137.pdf.

[78] T. Baer *et al.*, "UFOs in the LHC:Observations, Studies and Extrapolations," *Proceedings of IPAC2012*, pp. 3936–3938, 2012. [Online]. Available: http://inspirehep.net/record/1126583.

[79] F. Zimmermann, "Trapped dust in HERA and DORIS," *DESY HERA 93-08*, 1993. [Online]. Available: http://inspirehep.net/record/359539.

[80] F. Zimmermann *et al.*, "Trapped macroparticles in electron storage rings," *Proceedings of PAC1995*, pp. 517–519, 1995. DOI: 10.1109/PAC.1995.504705.

[81] D. Sagan, "Mass and charge measurement of trapped dust in the CESR storage ring," *Nucl.Instrum.Meth.A*, vol. 330, pp. 371–379, 1993. DOI: 10.1016/0168-9002(93)90566-Z.

[82] E. Jones *et al.*, "Transverse instabilities due to beam-trapped ions and charged matter in the cern antiproton accumulator," *IEEE Transactions on Nuclear Science*, vol. 32, pp. 2218–2220, 1985. DOI: 10.1109/TNS.1985.4333865.

[83] T. Baer *et al.*, "Very Fast Losses of the Circulating LHC Beam, their Mitigation and Machine Protection," *CERN-THESIS*, vol. 233, 2013. [Online]. Available: https://cds.cern.ch/record/1637966.

[84] W. Umrath *et al.*, *Fundamentals of Vacuum Technology*. Oerlikon, 2007. [Online]. Available: https://www3.nd.edu/~nsl/Lectures/urls/LEYBOLD_FUNDAMENTALS.pdf.

[85] B. Salvachua *et al.*, "Machine Protection Strategy for LHC Commissioning and Outlook for Future Challenges," *Proceedings of Chamonix 2014 Workshop on LHC Performance*, pp. 134–137, 2014. [Online]. Available: https://cds.cern.ch/record/2031197.

[86] E. B. Holzer *et al.*, "Beam Loss Monitoring for LHC Machine Protection," *Proceedings of TIPP2011*, 2055–2062, 2012. DOI: 10.1016/j.phpro.2012.04.110.

[87] A. Valishev *et al.*, "Tevatron Accelerator Physics and Operation Highlights," *FERMILAB-CONF-11-129-APC*, 2012. [Online]. Available: http://arxiv.org/ftp/arxiv/papers/1202/1202.5525.pdf.

[88] A. P. Verweij, "CSCM tests description," Oct. 2011, [Online]. Available: https://indico.cern.ch/event/154441/.

[89] H. Thiesen *et al.*, "CSCM-7TeV Status," Aug. 2012, [Online]. Available: https://indico.cern.ch/event/204695/.

[90] P. Hopkin, *Fundamentals of Risk Management. 2nd Edition.* Kogan Page, 2012, ISBN: 978-074-946-539-1. [Online]. Available: http://www.koganpage.com/product/fundamentals-of-risk-management-9780749472443.

[91] E. Todesco, "Quench Limits in the Next Generation of Magnets," *Proceedings of WAMSDO2013*, pp. 10–16, 2013. [Online]. Available: http://arxiv.org/ftp/arxiv/papers/1401/1401.3931.pdf.

[92] E. Todesco, "Thermal Conductivity and Electrical Resistivity of Very High-Purity Copper at Low Temperatures," *Proceedings of ICEC12*, pp. 414–417, 1988. DOI: 10.1016/B978-0-408-01259-1.50082-6.

[93] A. Apollonio *et al.*, "Measurement of the Residual Resistivity Ratio of the Bus Bars Copper Stabilizer of the 13 kA Circuits of the LHC," *Proceedings of IPAC2012*, pp. 3572–3574, 2012. [Online]. Available: http://inspirehep.net/record/1126699.

[94] L. Walckiers *et al.*, "Field Quality of the LHC Dipole Magnets in Operating Conditions," *Proceedings of EPAC 2002*, pp. 260–262, 2002. [Online]. Available: http://epaper.kek.jp/e02/PAPERS/WEALA002.pdf.

[95] V. Roger *et al.*, "Contact resistances in the cold bypass diode leads of the main LHC magnets," *IEEE Transactions on Applied Superconductivity*, vol. 24, 2014. DOI: 10.1109/TASC.2013.2294407.

[96] G. Willering, "Private Communications," 2015.

[97] A. Verweij, "QP3: Users Manual," *CERN EDMS*, vol. 1150045, 2008. [Online]. Available: https://edms.cern.ch/document/1150045/.

[98] B. Auchmann *et al.*, "Quench Behaviour Team," [Online]. Available: http://qbt.web.cern.ch/QBT/.

[99] B. Auchmann *et al.*, "Blm Threshold Strategy (UFOs and Quenches)," *Proceedings of Chamonix 2014 Workshop on LHC Performance*, pp. 143–148, 2014. [Online]. Available: https://e-publishing.cern.ch/index.php/CYR/article/view/133/77.

[100] A. Ferrari *et al.*, "Fluka: A multi-particle transport code," *CERN-2005-10*, 2005. [Online]. Available: http://www.fluka.org.

[101] B. Auchmann *et al.*, "Testing beam-induced quench levels of LHC superconducting magnets," *Phys. Rev. ST Accel. Beams*, vol. 18, 2015. DOI: 10.1103/PhysRevSTAB.18.061002.

[102] J. M. Zazula and S. Peraire, "Design studies of the LHC beam dump," *Proceedings of SARE3*, pp. 150–164, 1997. [Online]. Available: http://inspirehep.net/record/444977.

[103] M. Bassetti and G. A. Erskine, "Closed expression for the electrical field of a two-dimensional gaussian charge," *CERN-ISR-TH/80-06*, 1980. [Online]. Available: http://inspirehep.net/record/152962.

[104] V. N. Faddeeva and N. M. Terent'ev, *Tables of Values of the Function w(z)*. Pergamon Press, 1961. [Online]. Available: http://trove.nla.gov.au/version/26370929.

[105] V. Ziemann, "Beyond Bassetti and Erskine: Beam-beam deflections for non-Gaussian beams," *SLAC-PUB- 5582*, 1991. [Online]. Available: http://inspirehep.net/record/316705.

[106] F. Zimmermann *et al.*, "Interaction of Macro-Particles with LHC Proton Beam," *Proceedings of IPAC2010*, pp. 492–494, 2010. DOI: https://cds.cern.ch/record/1269302.

[107] N. F. Martinez *et al.*, "Simulation studies of macroparticles falling into the lhc proton beam," *Proceedings of IPAC2011*, pp. 634–636, 2011. [Online]. Available: http://inspirehep.net/record/1183564.

[108] H. Bethe and J. Ashkin, *Experimental Nuclear Physics Vol 1*. Wiley, 1953. [Online]. Available: https://archive.org/details/ExperimentalNuclearPhysicsVol1.

[109] Particle Data Group, "Review of Particle Properties," *Phys. Rev. D*, vol. 45, 1992. DOI: 10.1103/PhysRevD.45.S1.

[110] B. Rossi, "High Energy Particles," *Science*, vol. 117, 1952. DOI: 10.1126/science.117.3042.417-a.

[111] W. Blum *et al.*, *Particle Detection with Drift Chambers*. Springer, 2008, ISBN: 978-3-540-76684-1. [Online]. Available: http://www.springer.com/cn/book/9783540766834.

[112] H. Schindler, "Garfield++ User Guide (Version 2014.1)," 2014. [Online]. Available: http://cern.ch/garfieldpp.

[113] G. Iadarola *et al.*, "Analysis of the Electron Cloud Observations with 25 ns Bunch Spacing at the LHC," *Proceedings of IPAC2014*, 2014. [Online]. Available: http://inspirehep.net/record/1313694.

[114] A. Apollonio *et al.*, "Machine Protection: Availability for Particle Accelerators," *CERN-THESIS*, vol. 23, 2015. [Online]. Available: https://cds.cern.ch/record/2002820.

[115] L. Grob, "UFOs in the LHC," *Presented at 9th Gentner Day Meeting*, 2016. [Online]. Available: https://indico.cern.ch/event/472073.

[116] CERN, "First successful beam at record energy of 6.5 TeV," *CERN home*, 2015. [Online]. Available: http://home.cern/about/updates/2015/04/first-successful-beam-record-energy-65-tev.

[117] CERN, "LHC season 2: New frontiers in physics," *CERN*, 2015. [Online]. Available: https://run2-13tev.web.cern.ch/.

[118] B. Auchmann *et al.*, "How to Survive a UFO Attack," *Proceedings of 6th Evian Workshop*, 2015. [Online]. Available: https://indico.cern.ch/event/434129/.

[119] B. Auchmann *et al.*, "UFOs, ULO, BLMs," *Presented at Chamonix 2016 Workshop on LHC Performance*, 2016. [Online]. Available: https://indico.cern.ch/event/448109.

[120] G. Papotti, "Electron Cloud Effects," *Proceedings of Chamonix 2016 Workshop on LHC Performance*, 2016. [Online]. Available: https://indico.cern.ch/event/448109.

[121] G. Papotti *et al.*, "Macroparticle-induced Losses During 6.5 TeV LHC Operation," *Proceedings of IPAC2016*, 2016. [Online]. Available: http://ipac16.org/proceedings/papers/tupmw023.pdf.