

Development of Radiation-tolerant Components for the Quench Detection System at the CERN Large Hadron Collider

**Entwicklung von strahlungstoleranten Komponenten für das
Quench Detektorsystem am CERN Large Hadron Collider.**

Vom Fachbereich Physik
der Technische Universität Darmstadt
zur Erlangung des Grades
eines Doktors der Naturwissenschaften (Dr. rer. nat.)
genehmigte Dissertation
von Oliver Bitterling M. Sc.
aus Maintal

Referent: Prof. Dr. Joachim Enders
Koreferent: Prof. Dr. Rüdiger Schmidt

Tag der Einreichung: 07.02.2017
Tag der mündlichen Prüfung: 03.04.2017

Darmstadt 2017
D17



Development of Radiation-tolerant Components for the Quench Detection System at the CERN Large Hadron Collider

PhD. Thesis submitted by
Oliver Bitterling

Advisor:

Prof. Dr. Joachim Enders

Supervisor:

Dr. Reiner Denz

Day of submission: 7.2.2017

Abstract

This work describes the results of a three year project to improve the radiation tolerance of the Quench Protection System of the CERN Large Hadron Collider. Radiation-induced premature beam aborts have been a limiting factor for accelerator availability in the recent years. Furthermore, the future upgrade of the Large Hadron Collider to its High Luminosity phase will further increase the radiation load and has higher requirements for the overall machine availability. Therefore equipment groups like the Quench protection groups have used the last years to redesign many of their systems to fulfill those requirements. In support of the development of radiation-tolerant systems, several proton beam irradiation campaigns were conducted to determine the inherent radiation tolerance of a selection of varied electronic components. Using components from this selection a new Quench Protection System for the 600 A corrector magnets was developed. The radiation tolerance of this system was further improved by developing a filter and error correction system for all discovered failure modes. Furthermore, compliance of the new system with the specification was shown by simulating the behavior of the system using data taken from the irradiation campaigns. The resulting system is operational since the beginning of 2016 and has in the first 9 months of operation not shown a single radiation-induced failure. Using results from simulations and irradiation campaigns the predicted failure cross section for the full new 600 A Quench Protection System is $4.358 \pm 0.564 \cdot 10^{-10} \text{cm}^2$ which is one order of magnitude lower than the target set during the development of this system.

Abstract

Diese Arbeit beschreibt die Ergebnisse von drei Jahren Bemühungen die Strahlentoleranz des Quench Protection Systems am CERN Large Hadron Collider zu verbessern. Strahlungsinduzierte vorzeitige Strahlabbrüche waren ein limitierender Faktor für die Beschleunigerverfügbarkeit in den letzten Jahren. Darüber hinaus wird das zukünftige Upgrade des Large Hadron Collider auf seine High Luminosity-Phase die Strahlungsbelastung weiter erhöhen und höhere Anforderungen an die gesamte Maschinenverfügbarkeit haben. Aus diesem Grund haben Gerätegruppen wie die Quench-protection Gruppe die letzten Jahre genutzt, um viele ihrer Systeme neu zu gestalten, um diese Anforderungen zu erfüllen. Zur Unterstützung der Entwicklung strahlentoleranter Systeme wurden mehrere Protonenbestrahlungskampagnen durchgeführt um die inhärente Strahlungstoleranz einer Auswahl von verschiedenen elektronischen Komponenten zu bestimmen. Mit Komponenten aus dieser Auswahl wurde ein neues Quench Protection System für die 600 A Korrekturmagnete entwickelt. Die Strahlungstoleranz dieses Systems wurde weiter verbessert, indem ein Filter- und Fehlerkorrektursystem für alle erkannten Fehlermoden entwickelt wurde. Darüber hinaus wurde die Übereinstimmung des neuen Systems mit der Spezifikation gezeigt, indem das Verhalten des Systems unter Verwendung von Daten aus den Bestrahlungskampagnen simuliert wurde. Das daraus resultierende System ist seit Anfang 2016 in Betrieb und hat im ersten 9-Monatsbetrieb keinen einzigen strahlungsinduzierten Ausfall gezeigt. Mit den Ergebnissen von Simulationen und Bestrahlungskampagnen beträgt der vorhergesagte Ausfallquerschnitt für das neue 600 A Quench Protection System $4.358 \pm 0.564 \cdot 10^{-10} \text{cm}^2$ was um eine Größenordnung niedriger ist als das Ziel, das während der Entwicklung dieses Systems gesetzt wurde.

Contents

1	Introduction	10
2	Introduction to the LHC	12
2.1	Physics at the Large Hadron Collider	12
2.2	Availability and Physics Efficiency	14
3	Superconducting magnets of the LHC	19
3.1	Superconducting magnets	19
3.2	Magnet layout of the LHC	20
3.3	The 600 ampere magnet system	23
4	Magnet protection system	26
4.1	Quench protection methods	26
4.1.1	Quench Heater Strips	27
4.1.2	Cold Bypass Diode	28
4.1.3	Energy Extraction	28
4.2	Quench Detection System	28
4.2.1	Main Dipoles and Lattice Quadrupoles	29
4.2.2	Insertion Magnets	29
4.2.3	Corrector Magnets	30
5	Radiation effects	33
5.1	Effects of radiation on electronics	33
5.1.1	Displacement Damage	34
5.1.2	Total ionizing dose effects	34
5.1.3	Single-event effects	35
6	Radiation levels at CERN	39
7	Development of radiation-tolerant systems	42
7.1	Development and Protection Strategies	42
7.1.1	Radiation-Tolerant Technologies	43

7.1.2	Radiation-tolerant Design	47
7.2	Irradiation Campaigns	50
7.2.1	Motivation for Irradiation Campaigns	50
7.2.2	Procedures of Irradiation Campaigns	50
7.2.3	Irradiation Campaign Guide	57
8	Development of a new 600 A quench detection system	60
8.1	State of the previous QDS System	60
8.2	Requirements for the new QDS System	61
8.3	Design of the new System	61
8.4	Irradiation Campaigns on the ADS1281	63
8.4.1	Setup of the first campaign	63
8.4.2	Results of the first irradiation campaign	64
8.4.3	Conclusions from the first irradiation campaign	66
8.4.4	Setup of the second irradiation campaign	66
8.4.5	Results from the second irradiation campaign	69
8.5	Failure mitigation	73
8.5.1	Failure analysis	73
8.5.2	Failure mitigation	75
8.5.3	Mitigation results	81
9	Irradiation campaigns in support of further radiation-tolerant designs	83
9.1	The MAX1162 and MAX11100 ADCs	83
9.1.1	Description of the Component	83
9.1.2	Irradiation Campaign on the ADCs MAX1162/11100	84
9.1.3	Conclusion	91
9.2	Voltage References MAX6250, ADR421 and ADR435	91
9.2.1	Description of the Components	91
9.2.2	Irradiation Campaign on the Voltage References MAX6250, ADR421 and ADR435	92
9.2.3	Conclusion	96
9.3	The AT25DF041A Flash Memory	96
9.3.1	Description of the Component	96
9.3.2	Irradiation Campaign on the AT25DF041A Flash Memory	97
9.3.3	Conclusion	98
9.4	The ADUM3402 Digital insulator	98
9.4.1	Description of the Component	98
9.4.2	Irradiation Campaign on the ADUM3402 Digital insulator	98

9.4.3	Conclusion	98
9.5	The INA128 Instrumentation Amplifier	99
9.5.1	Description of the Component	99
9.5.2	Irradiation Campaign on the INA128 Instrumentation Amplifier	99
9.5.3	Conclusion	103
9.6	The AQW210EHA PhotoMOS	103
9.6.1	Description of the Component	103
9.6.2	Irradiation Campaign on the AQW210EHA PhotoMOS	103
9.6.3	Conclusion	103
9.7	The AQV251A PhotoMOS	103
9.7.1	Description of the Component	103
9.7.2	Irradiation Campaign on the AQV251A PhotoMOS	104
9.7.3	Conclusion	108
9.8	Conclusion and outlook	108
10	Irradiation campaigns on the SmartFusion2/Iglloo2 FPGA	109
10.1	Irradiation campaign on the SmartFusion2	110
10.2	Results	111
10.3	Results from other research teams	112
10.4	Conclusion	113
11	Conclusion	115
11.1	Conclusion of the Project	115
11.2	Outlook	118
	Declaration – Erklärung	127
	Danksagungen	129

List of Tables

3.1	Superconducting magnets of the LHC	24
3.2	600 A Magnets	25
4.1	Protection measures used in the various quench protection systems	27
6.1	Radiation levels past, present and future in comparison to the annual integrated luminosity. The values of 2016 were extrapolated from a set of data from the first 6 months of the year and the remaining days of beam time. The values for the HL-LHC phase were predicted using the data from 2016 and the annual integrated luminosity that is the goal during the HL phase.	41
7.1	Facilities that can provide a mono particle beam for irradiation campaigns.	54
8.1	ADC stop length distribution	71
8.2	Error cross sections of all discovered error modes	72
8.3	Results of the mitigation measures during the second irradiation campaign. Two multi-sample and two long ADC stops could not be mitigated. Mitigation against configuration corruption was not yet in place.	81
9.1	Overview of irradiated DUTs	84
10.1	Results of the ESRAM irradiation campaign	112
10.2	SEU cross sections of all tested systems	113

List of Figures

2.1	LHC layout	13
2.2	Availability Goals	15
2.3	Stable beam duration in 2015	16
2.4	Downtime of the accelerator due to the various systems in 2015	17
3.1	Overview of the LHC	21
3.2	General LHC Layout	22
3.3	Layout of an ARC cell of the LHC.	22
3.4	Layout of a short straight section	23
4.1	Protection system for the Main Dipoles	27
4.2	QDS for main dipoles and lattice quadrupoles	30
4.3	QDS for the individually powered magnets in the insertion region	31
4.4	Schematic of a 600 A Magnet circuit	32
5.1	Ionizing Radiation strikes a Transistor	35
5.2	Occurrence of SEUs and SETs	37
5.3	Schematic of a SRAM Cell	38
6.1	Positions of the 600 A QPS	40
7.1	Relation between radiation level and usable system types . . .	45
7.2	Triple Modular Redundancy	49
7.3	Irradiation Campaign Test Procedure	53
7.4	Proton Irradiation Facility at PSI	56
7.5	The CHARM irradiation facility	58
8.1	Circuit board of the previous 600 A QDS system(10x20 cm) .	60
8.2	Building blocks of the ADS1281	62
8.3	Schematic of the baseboard used during the first irradiation campaign	63
8.4	Setup of the first irradiation campaign	64

8.5	Oscillation measured during the first irradiation campaign . . .	65
8.6	Example ADC stop found during all irradiation campaigns . . .	66
8.7	Version of the 600 A QDS board used during the second Irradiation Campaign	67
8.8	Test Setup used during the second Irradiation Campaign . . .	68
8.9	Maximum drift ADS1281	70
8.10	Minimum drift ADS1281	70
8.11	ADC output signal during irradiation	71
8.12	Permanent change of the signal gain.	72
8.13	Filter kernels of the four FIR filter stages	74
8.14	Comparison between a simulated multi-sample error and a measured one	74
8.15	Filter chain to mitigate single and multi-sample errors	76
8.16	Two decimation results of I_{Dect} error	78
9.1	Building blocks of the MAX1162/11100	84
9.2	Baseboard	86
9.3	Schematics of the adapter boards containing the DUTs	87
9.4	MAX11100 Current Consumption	88
9.5	MAX1162 Signal Maximum Drift	88
9.6	MAX1162 Signal Minimum Drift	88
9.7	MAX11100 Signal Maximum Drift	89
9.8	MAX11100 Signal Minimum Drift	89
9.9	Integrated error distribution in relation to the total dose for MAX1162 and MAX11100	90
9.10	Output Voltage of the MAX6250 over TID	92
9.11	Current Consumption of the MAX6250	93
9.12	Output Voltage of the ADR421 over TID	93
9.13	Current Consumption of the ADR421	94
9.14	Output Voltage of the ADR435 over TID	95
9.15	Current Consumption of the ADR435	96
9.16	Flash error counters	97
9.17	Output Voltage of the INA128	100
9.18	Bias Current of the INA128	101
9.19	Current Consumption of the INA128	102
9.20	AQV251 test schematic	105
9.21	Output voltage of the AQV251	106
9.22	On-state resistance of the AQV251	107
10.1	Schematic of the SmartFusion2	110
10.2	Development board of the SmartFusion2 by Microsemi	111

11.1 Fault time comparison	116
--------------------------------------	-----

Acronyms

- **A** Availability
- **AMBA** Advanced micro controller Bus Architecture
- **APB** Advanced Peripheral Bus
- **ARM** Advanced RISC Machines
- **CHARM** Cern High-energy AcceleRator Mixed Field/facility
- **COTS** Commercials of the Shelf
- **DAC** Digital to Analogue Converter
- **DD** Displacement Damage
- **DS** Dispersion Suppressors
- **DSP** Digital Signal Processor
- **DUT** Device Under Test
- **FIR** Finite Impulse Response
- **FPGA** Field Programmable Gate Array
- **HL-LHC** HL-LHC
- **HEH** High Energy Hadron Fluence
- **HEHeq** Equivalent High Energy Hadron Fluence
- **IP** Interaction Point
- **IR** Insertion Region
- **LEP** Large Electron-Positron Collider

- **LHC** Large Hadron Collider
- **MOSFET** Metal Oxide Semiconductor Field-Effect Transistor
- **PE** Physics Efficiency
- **PIF** Proton Irradiation Facility
- **PS** Proton-Synchrotron
- **PSI** Paul-Scherrer Institute
- **QDS** Quench Detection System
- **QPS** Quench Protection System
- **RMS** Root-Mean-Square
- **SAR** Successive Approximation Register
- **SEB** Single-event Burnout
- **SEE** Single-Event Effect
- **SEFI** Single-Event Functional Interrupt
- **SEGR** Single-event Gate Rupture
- **SEL** Single-event-Latch-up
- **SER** Single-Event Rate
- **SET** Single-Event Transient
- **SEU** Single-Event Upset
- **SPS** Super-Proton-Synchrotron
- **SPT** Scheduled Proton Physics Time
- **SSS** Short Straight Section
- **TID** Total Ionizing Dose
- **TMR** Triple Modular Redundancy

Chapter 1

Introduction

Large particle accelerators, like the Large Hadron Collider (LHC), are used for collision experiments to create new particles. Particle accelerators can also be used to create an environment similar to the early time of the universe, enabling us to take a look into this mostly unknown period of time. But every reaction that one wants to observe requires a certain collision energy to be possible. Therefore to discover unknown reactions and particles a machine with collision energies higher than what previously existed is required. Furthermore, each reaction only has a very low probability of occurring for every given collision at a sufficient energy level. Therefore, to increase the chance of observing an unknown reaction it is necessary to increase the number of collisions. The number of collisions is proportional to the integrated luminosity which is the integral of the luminosity of the machine over the time it is able to sustain this luminosity (see chap. 2.1). The future upgrades to the LHC will increase the luminosity to the highest level where the produced amounts of data can still be processed by the high luminosity experiments (ATLAS and CMS). Therefore, further increase of the integrated luminosity can only be achieved by improving the availability of the accelerator (see chap. 2.2). Unfortunately, an increase in luminosity also produces an increase of the radiation level in the vicinity of the accelerator. This level of radiation can cause the superconducting magnets of the accelerator to turn resistive (which is called a Quench) which has the potential to damage the accelerator (see chap. 3). To protect the machine from such events a dedicated protection system to detect and to respond to Quenches is required (see chap. 4). But like all electronic systems this very protection system is not immune to the radiation levels at its locations (see chap. 5 and 6). Due to the, for accelerator facilities, unprecedented level of radiation present around the accelerator, shielding cannot completely prevent the irradiation of the protection systems. Most other accelerators have radiation

levels low enough to only cause few problems. The only similar challenging environment is outer space where cosmic radiation is strong enough to also damage electronic equipment. Fortunately, there exist several methods for designing and protecting electronic systems in order for them to be able to operate in such a radiation field (see chap. 7). In the scope of this work, a new protection system for the 600 A magnets was developed (see chap. 8). This includes irradiation campaigns to determine the radiation tolerance of the base components, analysis of the found failure modes, development simulation and testing of mitigation measures and final tests of the system under radiation. Furthermore, several other irradiation campaigns were conducted to collect data about components which are either already used or planned parts of future projects (see chap. 9 and 10). The results of this work were very good and the components developed will help to increase the overall accelerator availability of the LHC (see chap. 11).

Chapter 2

Introduction to the LHC

2.1 Physics at the Large Hadron Collider

The Large Hadron Collider (LHC) in Meyrin Switzerland at CERN is currently the highest-energy proton/hadron collider [1]. First officially mentioned in 1984 at a workshop in Lausanne and inaugurated in 2008, the LHC was built to address several fundamental questions of particle physics:

- Search for the Higgs boson to understand the mechanism creating the masses of elementary particles
- Search for supersymmetric particles and other predictions of approaches beyond the standard model of particle physics and identify candidates for dark matter
- Understand the mechanism creating the asymmetry between matter and antimatter in the universe
- Characterize the quark-gluon plasma

The Higgs boson was found in 2012, but the others questions remain open [2][3].

The LHC is a 27 km long proton/hadron collider built in the tunnel previously occupied by the Large Electron-Positron Collider (LEP). Using the Proton-Synchrotron (PS) and the Super-Proton-Synchrotron (SPS) as pre-accelerators the LHC can reach beam energies in the order of 7 TeV. To keep the beam on a circular orbit the LHC comprises of powerful superconducting magnets with field strengths of up to 8.33 T. This machine allows for proton-proton collisions with center of mass energies up to 14 TeV.




Figure 2.1: LHC layout [4]

In order to reveal physics beyond the standard model, it is necessary to maximize the number of collisions inside the experiments. The amount of events of a specific type is given by

$$N = L \cdot \sigma_{Event}. \quad (2.1)$$

The cross section σ_{Event} is an effective area that describes the probability of a specific scattering event occurring as a result of the crossing of two beams of discrete particles. For the collision of two identical beams in a circular collider, the luminosity of the machine L is given by

$$L = \frac{N_b^2 \cdot n_b \cdot f_{rev} \cdot \gamma_r}{4\pi \cdot \epsilon_n \cdot \beta^*} F. \quad (2.2)$$

Here N_b is the number of particles per bunch, n_b the number of bunches per beam, f_{rev} the revolution frequency, γ_r the relativistic gamma factor, ϵ_n the normalized transverse beam emittance, β^* the beta function at the collision point and F the geometric luminosity reduction factor due to the crossing angle at the interaction point (IP)

$$F = (1 + (\frac{\theta_c \cdot \sigma_z}{2 \cdot \sigma^*})^2)^{-1/2}. \quad (2.3)$$

Here θ_c is the full crossing angle at the IP, σ_z is the root-mean-square (RMS) bunch length and σ^* the transverse RMS beam size at the IP [5]. The LHC's two high luminosity experiments ATLAS [6] and CMS [7] can process events created by a luminosity up to $L = 10^{34} \text{ cm}^{-2}\text{s}^{-1}$, which is necessary to observe events with low cross section. Current upgrades will lead to an increase of the luminosity that these experiments can process by a factor of three [8]. Besides a high luminosity, it is also necessary to achieve a high percentage of time where a beam with this luminosity is available for physics experiments. During a successful run in the accelerator, the target luminosity can be achieved for some time followed by a phase of slowly degrading luminosity due to beam losses and other beam effects [9]. Therefore there is a maximum integrated luminosity that can be gained during one accelerator cycle.

2.2 Availability and Physics Efficiency

CERN is constantly in the process of upgrading the LHC to reach higher integrated luminosity. The goal for the High Luminosity phase of the LHC (HL-LHC) is an annual integrated luminosity of 250-300 fb^{-1} and 400 fb^{-1}

during ultimate conditions. With a baseline of 160 days per year with physics runs this would require about 1.9 fb^{-1} per day. For comparison the annual integrated luminosity in 2012 was only 23 fb^{-1} and 23 fb^{-1} in the first 7 months of 2015. Because this integrated luminosity goal can not be achieved by only increasing the luminosity the accelerator availability has to be increased significantly. Limiting factors are the frequency of faults and the length of downtime due to faults. Faults that force a physics beam to be aborted prematurely are an important limiting factor for the integrated luminosity.

Figure 2.2 shows the relation between the ratio of prematurely aborted beam, the average fault time and the achievable integrated luminosity [10].




Figure 2.2: Relation between the ratio of prematurely aborted beam, the average fault time and the achievable integrated luminosity. The stars show the situation in 2012, the achievable results after applying mitigation measures against radiation-induced faults and the final goal during HL-LHC [10].

Figure 2.3 shows the distribution of the duration of the periods with stable beams for the year 2015. In a large parts of the fills the beams were aborted in the first few hours.

Besides actual faults in the system, a significant part of premature beam dumps is caused by radiation-induced errors in the systems. In 2012 a total of 585 premature dumps were counted and 56 could be attributed to faults caused by the influence of ionizing radiation. 31 of those premature dumps were triggered by the Quench Detection System (QDS). The total amount of beam aborts caused by the QDS is only about double that number, making radiation effects responsible for 50% of all QDS aborts in 2015 [10]. Following a fault, measures have to be taken to return the accelerator to working




Figure 2.3: Stable beam duration in 2015 [11]

conditions. These measures include:

- Diagnosis of the fault by control room and remotely by experts
- In some cases intervention in the machine by an expert
- Recovery from the impact of a fault (for example cool down of magnets after a quench)
- Return to operation conditions (recovery from access, precycle, injection etc.)

Depending on fault type fault recovery can take several hours or even days in case of serious faults. Figure 2.4 shows the respective fault times of the LHC systems for the year 2015 [12].

Following the recovery from a fault, the accelerator has to return to stable beam conditions. The Turnaround consists of the following phases:

- Ramp down/pre-cycle
- Pre-injection checks and preparation
- Checks with set-up beam
- Nominal injection sequence
- Ramp preparation
- Ramp
- Adjust/collisions

LHC Impact Fault Times by System




Figure 2.4: Downtime of the accelerator due to the various systems in 2015 [12]. As can be seen the Quench Protection system (QPS) has been a major source of accelerator downtime.

The turnaround time, which averaged to about 5.5 hours in 2012, is another source of downtime which is lost for physics operation. All of these factors influence the availability of the accelerator. The availability (A) of an accelerator is defined as the time planned during which a proton beam should be available for physics experiments (**S**cheduled proton **P**hysics **T**ime **SPT**) minus the time assigned to faults and fault recovery expressed as a percentage of the SPT. Another important characteristic is the physics efficiency (PE). It describes the duration while the accelerator is available and can supply a stable beam for experiments. The relations between availability A and physics efficiency PE with N_f being the number of accelerator fills in the measurement frame and T_{fill} and T_{around} being the time spent either during a stable beam or during turnaround of the beam [13], are given by

$$PE = A \cdot SPT - N_f \cdot T_{around}, \quad (2.4)$$

$$N_f = \frac{A \cdot SPT}{T_{fill} + T_{around}}, \quad (2.5)$$

and hence

$$PE = A \cdot SPT \cdot \left(1 - \frac{T_{around}}{T_{fill} + T_{around}}\right). \quad (2.6)$$

In order to achieve the goals of the LHC, it is necessary to optimize availability and PE of the accelerator further. In the year 2015, the availability of the LHC was 68% for runs with a 25 ns bunch spacing. Long term goals for the LHC high luminosity phase (HL) target an availability of around 90% [14]. The whole system will have to mature considerably to reach this goal. To increase the availability there are several factors that can be modified like the frequency of faults and the speed of recovery after a fault. Reducing the frequency of faults that require the abort of the beam would also increase the PE because the production of a new stable beam requires time (T_{around}). The last years have seen a reduction in the turnaround time by measures like combining the energy ramp with squeezing (reducing beam circumference in preparation for collision). Still, the average turnaround time in 2015 was about 8.8 hours or 6.8 hours if only fills without serious faults are counted [11]. Faults in a system can also be categorized as those that have their origin in a real problem of the machine, like a quenched magnet (see chap. 3), and those that originate from errors in the various protection systems. Protection systems are required for the safe operation of accelerator equipment (e.g. superconducting magnets), as well as for operating with beams that have significant damage potential to prevent beam induced damage. Their detection systems are vulnerable to the effects of ionizing radiation which can induce signals that mimic critical conditions inside the accelerator and lead to an abort of the fill (see chap. 5). This work presents several efforts made in support of the update the Quench Detection System (QDS). The goal is to protect the system from radiation-induced errors in order to increase the overall accelerator availability and physics efficiency. A special focus is put on the development of a new QDS for the 600 A corrector magnets (see chap. 3.2). Further the results of several irradiation campaigns are discussed which provide valuable data concerning components that can be used in the future development of radiation-tolerant systems

Chapter 3

Superconducting magnets of the LHC

3.1 Superconducting magnets

Because the LHC was built in the tunnel previously occupied by LEP, the radius of the accelerator was fixed by the already existing tunnel radius. A synchrotron like the LHC uses a magnetic field to keep charged particles traveling close to the speed of light on a circular orbit. Due to the fixed radius and the higher mass of the accelerated particles, the LHC requires a field strength about forty times higher than LEP. At energies of 7 TeV, a magnetic field of 8.33 T is required [5]. Achieving such a field strength using normal conducting magnets is possible but impractical. The current consumption of such magnets would be very high and the cooling system would not be able to cope with the ohmic heating of the magnet. Therefore most magnets of the LHC are superconducting. Superconductivity is a state where some materials lose all electric resistance if cooled below a critical temperature [15]. This eliminates the losses due to ohmic heating in the magnets. The superconducting magnets used for the LHC use proven NbTi technology. Previous machines using such magnets are the Tevatron [16], HERA [17] and RHIC [18] which cooled the magnets down to temperatures of 4.2 K reaching magnetic fields around 5 T. Increasing the magnetic field further is not trivial because every superconductor has a maximum magnetic field, depending on material, geometry, current density and temperature, after which it transitions to normal conductivity [19]. To enable fields of 8-8.5 T, superfluid helium is used to cool the magnets below 2 K. But lowering the temperature by more than a factor of two has the detrimental effect of reducing the heat capacity of the superconductor by almost an order of magnitude [15][5].

Therefore the energy necessary to raise the temperature per degree is lower than at higher temperatures. This makes it more probable for a small energy deposition into the magnet to raise the local temperature above the critical temperature which causes that spot of the magnet to transition to normal conductivity. Ohmic heating of that point increases the temperature in that spot further leading to a loss of superconductivity in the magnet. This process is called a quench. A Quench has the potential to heavily damage a magnet due to the large amount of energy stored in the magnets. Hotspot temperatures inside the magnet during a quench can reach temperatures to melt the magnet material. It is not possible to prevent the occurrence of quenches perfectly. Therefore several protection measures are in place to protect magnets during a quench by extracting the energy from the magnet circuit and preventing extreme hot-spot temperatures using quench heaters (see chap. 4). Even if the protection is successful the affected magnets have to be cooled back to superconductivity before the operation can resume. This process takes a considerable amount of time which impacts the availability of the accelerator.

3.2 Magnet layout of the LHC

The LHC comprises eight sectors (see fig. 3.1) each reaching from one interaction point to the next. Each Sector contains one ARC, two Dispersion Suppressors (DS), two Matching Sectors, two inner triplets and one Interaction Point (see fig. 3.2).

The ARCs are the curved sections of the accelerator containing the main bending dipole magnets (MBA and MBB). Each ARC comprises of 23 lattice cells that can each be separated into two half cells (see fig. 3.3). Each half cell contains three bending dipoles and a Short Straight Section (SSS).

To achieve the required field quality to operate the LHC, the dipole magnets are equipped with Sextupole (MCS) and Decapole (MCD) spool corrector magnets. Each decapole corrector will, in addition, have an Octupole insert (MCO) and these together are designated as MCDO [20]. These spool piece correctors mitigate systematic sextupole and decapole field errors in the dipole field [21]. The SSS contain the sets of Focusing/Defocusing Quadrupoles (MQ) that focus the beam to counteract the natural angular divergence of the beam. Further, the SSS contain several corrector magnets that compensate imperfections in the field of a dipole, quadrupole, and other corrector magnets (see tab. 3.1).

Before and after each arc a dispersion suppressor (DS) is used to reduce the machine dispersion (momentum spread of beam particles). Each LHC




Figure 3.1: Overview of the LHC [5]

dispersion suppressor consists of four individually powered quadrupole magnets (MQ, MQML, MQML+MQM) which are separated by two dipole magnets (MBA and MBB [22]). The DS also contain further corrector magnets for Skew, Trim, and orbit correction.

The DS is then followed by a matching section whose magnetic setup is unique for the specific Insertion Region (IR). The matching sections contain 3 types of matching quadrupoles and further orbit and trim correctors.

Several interaction points are flanked by a pair of dipole separation magnets. For the experiment points these magnets (D1/D2) bring the two beams onto a collision orbit and the separate them again beyond the collision point. Currently, the D1 magnets of the high luminosity experiments are normal conducting magnets due to the high radiation levels. During the HL-LHC phase, those will be replaced with superconducting magnets. At the RF point, a pair (D3/D4) separates the beam so that individual RF cavities can be installed for each beam.




Figure 3.4: Example layout of one type of several types of short straight sections used in the LHC [5]

3.3 The 600 ampere magnet system

All magnets listed in the previous chapters have their own protection system to detect losses of superconductivity and execute protection measures to prevent damage. Depending on the magnet type and their current level different protection methods are necessary (see chap. 4.1). This work focuses mainly on all magnets that are operated in a range of ± 600 A and their Quench Detection System (QDS). The 600 A magnets are corrector magnets that are used for orbit correction and to mitigate higher order field errors of the main dipoles and quadrupoles. They can be found in all parts of the accelerator (ARC, Short-Straight-Sections, Dispersion Suppressor, Matching section and the inner triplet). Table 3.2 lists the magnet families belonging to the 600 A system, their function and their respective position in the accelerator.

Family Name	Description	System	Nummber
MB	Dipole Bending Magnet	ARC DS	1232
MCS	Sextupole Spool Corrector	ARC	2464
MCD	Decapole Spool Corrector	ARC	1232
MCO	Octupole Spool Corrector	ARC	1232
MQ	Focusing/Defocusing Quadrupole	SSS DS	392
MO	Landau dampening Octupoles	SSS	336
MQT	Tuning Quadrupoles	SSS	320
MQS	Skew Quadrupoles	SSS	64
MSCB	Skew Sextupole + Dipole Orbit Correctors	SSS DS	688
MQTL	Trim Quadrupole	DS	120
		Matching Section	
MQML		DS	36
		Matching Section	
MCBC	Closed Orbit Corrector	DS	908
		Matching Section	
MQMC		DS	12
MQM	Matching Quadrupole	DS	38
		Matching Section	
MQY		Matching Section	24
MCBY	Closed Orbit Corrector	Matching Section	88
MBX (D1)	Separation Dipole	Before Inner Triplet	4
MBRC (D2)	Separation Dipole	Before Inner Triplet	8
MBRS (D3)	Separation Dipole	Before Inner Triplet	4
MBRB (D4)	Separation Dipole	Before Inner Triplet	4
MQXA	Triplet Quadrupole	Inner Triplet	16
MQXB	Triplet Quadrupole	Inner Triplet	16
MCBX	Orbit Corrector	Inner Triplet	48
MCSTX	Sextupole/Decapole Corrector Package	Inner Triplet	8
MCSOX	Skew Sextupole-Octopole Corrector	Inner Triplet	8
MQSX	Skew Quadrupole	Inner Triplet	8

Table 3.1: Superconducting magnets of the LHC

Magnet Family	Description	Position
MO	Main Octupole	SSS
MQT	Tuning Quadrupole	SSS
MQS	Skew Quadrupole	SSS
MSCB	Skew Sextupole + Dipole Orbit Correctors	SSS DS
MCS	Spool piece Sextupole Corrector	ARC
MCO	Spool piece Octopole Corrector	ARC
MCD	Spool piece Decapole Corrector	ARC
MQTL	Trim Quadrupole	DS
		Matching section
MCBX	Orbit Corrector	Inner Triplet
MQSX	Skew Quadrupole	Inner Triplet
MCSOX	Skew Sextupole-Octopole Corrector	Inner Triplet

Table 3.2: 600 A Magnets

Chapter 4

Magnet protection system

4.1 Quench protection methods

A transition away from superconductivity poses two different threats to a magnet: high temperatures due to ohmic heating and high voltages due to a growing resistance at high currents [23] [24]. High temperatures can destroy the insulation material or even result in a meltdown of the cable. High voltages can cause electric discharges that could melt holes into the insulation of the cables. Furthermore, high current density and temperature can irreversibly degrade the superconducting material reducing its current carrying capability [25]. In the LHC different protection methods are used according to magnet type. Magnets with higher current or lower critical temperature require more protection than less fragile magnets. The protection system can contain the following components:

- cold bypass diodes
- quench heaters
- energy extraction
- cold parallel resistor

Table 4.1 shows the components that are used to protect magnets in a certain current range.

Not all magnets require the full set of protection measures [26]. Most magnets in one family are connected in series and depending on their operation current store large amounts of energy. The higher the energy stored in the chain the more protection measures are required. Due to technical limits and to limit maximum energy and voltage in a magnet chain, the chains do

Magnet Current	≤ 600 A	6 kA + 13 kA Individ. powered	13 kA
Cold bypass diodes	–	X	X
Dump resistors	X (most)	–	X
Quench heaters	–	X	X
Parallel resistor	X (most)	–	–

Table 4.1: Protection measures used in the various quench protection systems

not extend beyond one sector. If no precautions were taken, all the energy in such an electrical circuit would be deposited in the quenching magnet. Quench protection uses both passive and active methods. The passive methods include choosing cable material and design techniques that provide more operational margin. The active methods used at the LHC are described in the following chapters. Figure 4.1 shows the full protection system for the main dipoles.




Figure 4.1: Protection system for the Main Dipoles [27]

4.1.1 Quench Heater Strips

As soon as a quench in a magnet circuit is detected, the power converter of that circuit is deactivated leading to a decay of the current. The speed of this decay is dependent on the resistance of the magnet chain. This process is sped up by heating large parts of the magnets. By turning the whole magnet resistive the stored energy does not dissipate at a single spot which would cause destructive hot-spot temperatures. This is done by activating Quench heaters, heater strips that cover the outside of the magnet coils. Heating only parts of the magnet is enough because quenches propagate into the rest of the magnet with velocities of 15 m/s to 20 m/s. A 400 mm long section would be fully quenched in less than 10 ms [28]. The increase of resistance

of the magnet also increases the voltage over the magnet above the turn-on voltage of the Cold Bypass Diode.

4.1.2 Cold Bypass Diode

The current in the magnet chains decays usually over a time in the order of a few 100 ms. Letting a high current flow through the now resistive magnet could damage it due to resistive heating. To prevent this, a cold bypass diode that is connected in parallel to each magnet is used to route most of the current around the quenching magnet. During normal operation, the cold diodes have a turn-on voltage of 6 V. The turn-on voltage is high enough for the diode to remain inactive during normal magnet ramps (about 1 V for main dipoles with 10 A/s). During a quench, the rising resistance of the magnet and the corresponding voltage activate the diode which starts to conduct. The following warming of the diode further decreases the turn-on voltage allowing most of the current to flow through the diode instead of the magnet. To speed up the activation of the diode, quench heaters are activated as soon as a quench is detected [25].

4.1.3 Energy Extraction

The decay of the current using only the resistance of the quenching magnets can take hours depending on the stored energy in the circuit. The high current flowing through the quenching magnet during this time would damage the magnet. To speed up the process in a magnet circuit it is possible to use a dump resistor to extract the energy. The dump resistor is connected in series to the magnet circuit and can be switched on with by opening a switch (see fig. 4.1). This is necessary for all magnet circuits that store energy equal or higher than 2 kJ like the 13 kA main dipole and main quadrupole magnets and several of the 600 A magnet systems [5]. The dimension of the dump resistor is chosen according to how long the rest of the circuit can withstand the high voltage during a quench [25].

4.2 Quench Detection System

All of the superconducting components of the LHC need protection in case of a loss of superconductivity. To ensure the activation of proper mitigation measures for quenches it is necessary to detect them as fast as possible. This is performed by dedicated quench detection systems (QDS). Each type of magnet family has a QDS system customized for its operation environment

and instrumentation. While all quenches have to be detected to ensure the protection of the magnets it is also necessary to prevent erroneous activation of the magnet protection systems. Every activation of the quench protection measures caused the beam to be dumped and quench heaters to heat up the affected magnets. The re-cooling of those magnets below critical temperature requires minutes to hours depending on circuit type. Therefore errors in the QDS can strongly impact the availability of the accelerator. Furthermore, all QDS systems are installed close to the accelerator which subjects them to, depending on their exact location, varying levels of ionizing radiation. Radiation is an additional source of errors for electronic systems and decreases their lifetime (see chap. 5).

4.2.1 Main Dipoles and Lattice Quadrupoles

In case of the main dipoles and lattice quadrupoles the detection of quenches is accomplished by measuring the resistance in the two apertures using an analogue Wheatstone bridge like setup (see fig. 4.2). The voltage over both apertures is constantly measured and compared with each other. In case of a transition away from superconductivity in one of the apertures, the voltage will increase significantly. If the voltage difference between both apertures increases above a system-specific threshold, the magnet protection system is triggered. To prevent erroneous triggers due to the influence of noise, the trigger signal is routed through a time discriminator [29]. To protect the magnets against aperture symmetric quenches a second type of quench detectors was installed. The symmetric quench protection system compares the total magnet voltage of four electrically adjacent dipoles. In case the difference of one of the four voltages with respect to the others exceeds the threshold, the heaters of this magnet are triggered.

4.2.2 Insertion Magnets

Magnets in the insertion regions are mostly individually powered and therefore can not be compared to adjacent magnets for quench detection. Therefore the voltage of the first and second half of the same magnet is compared with each other (see fig. 4.3). For this purpose, the magnet has a voltage tap at its midpoint. A difference between both magnet halves higher than a magnet-specific threshold is considered as a quench.




Figure 4.2: QDS for main dipoles and lattice quadrupoles. The setup forms a Wheatstone bridge allowing for the calculation of the resistance in the coils if the values of one coil and both resistors are known.

4.2.3 Corrector Magnets

Instead of local quench detection using mid-point voltage taps, the 600 A circuits use a global system monitoring the whole circuit. By utilizing current sensors at the ends of the current leads, the expected voltage of the circuit in a superconducting state can be calculated. A comparison between the measured and the expected voltage in a 600 A magnet circuit (see fig. 4.4) is given by

$$U_{Res} = U_{Diff} - L(I_{Dcct}) \cdot \left(\frac{\delta I_{Dcct}}{\delta t} + \frac{\delta U_{Diff}}{\delta t} \cdot \frac{1}{R_{Par}} \right). \quad (4.1)$$

Here U_{Diff} is the voltage measured over the whole circuit and I_{Dcct} is the current flowing through the circuit. U_{Res} is the difference between measured voltage and an expected voltage calculated under the assumption that the circuit is superconductive. If the difference between expected and measured




Figure 4.3: QDS for the individually powered magnets in the insertion region. The voltage over both magnet halves is measured and compared using two ADCs and an FPGA.

voltage passes a threshold of 100 mV this is considered as a quench. The calculation of a derivative of both current and voltage requires a high precision measurement of both values. Derivation of low-resolution measurements would easily produce spikes sufficiently high to trigger the QDS.




Figure 4.4: A simplified schematic of a 600 A magnet circuit. The whole circuit is in a state of superconductivity. During normal operation R_{Par} is the only resistance in the circuit and is a constant value. The inductance inductance of the coil L is taken from a table that correlates current with inductance. By measuring and deriving I_{dcct} it is possible to compute the expected voltage to compare with the measured voltage U_{diff}

Chapter 5

Radiation effects

5.1 Effects of radiation on electronics

Before 1978, radiation-induced errors were considered mainly as a problem for space applications. Cosmic rays consisting mostly out of high-energy protons, neutrons, electrons and alpha particles were known to cause issues with electronic components in satellites and space crafts. At sea level, radiation-induced errors were not considered to be a source of problems due to the shielding effect of the atmosphere. 1978 this proved to be a misconception with the paper by T.C.May and M.H.Woods detailing radiation-induced errors in processor cores due to uranium-contaminated packaging material [30]. There are two general types of influence radiation can have on electronic systems: single-event effects (SEE) and cumulative damage. Cumulative damage can further be separated into the effects of total ionizing dose (TID) and displacement damage (DD). Cumulative damage is caused by prolonged exposure to high-energy radiation which causes damage to the semiconductor lattice and surrounding layers. SEEs are caused by single ionizing particles passing a semiconductor lattice and creating additional free charge carriers. If the free charge reaches a critical level Q_{Crit} [30], an SEE is produced. The type of SEE produced depends on the circuit type and specific location of the critical charge in the circuit. Furthermore, SEEs can be separated into soft errors causing a temporary error in a system, which can be fixed, and hard errors that permanently damage or destroy a system. The physics of these effects is further discussed in the following sections. In 1979, J.F. Ziegler and W.A. Lanford predicted SEEs at sea level [31], which was confirmed by IBM in 1984 [32]. Later in 1995 a new source of soft errors was discovered in electronic devices containing borophosphosilicate glass. Upon capturing a thermal neutron boron-10 isotopes fissure $^{10}B(n, \alpha)^7Li$ and the created α

and Li recoil particles were shown to cause soft errors [33]. In general, collecting historical data about soft errors is challenging because it is difficult to trace back occurring errors to radiation. Furthermore, companies are usually unwilling to reveal information about problems with their equipment. Two other known events which could be attributed to soft errors happened 2004 at Cypress Semiconductor [32] and 2005 at the Los Alamos National Laboratory [34].

5.1.1 Displacement Damage

Interaction of ionizing particles with matter can both influence the electrons and the nucleus of a lattice atom. An interaction with the nucleus can change the position of the atom within the lattice of materials like silicon. Such an interaction creates a vacancy in the lattice and an interstitial atom called a Frenkel pair. Such disturbances of the lattice are called displacement damage. Electrons and low energy protons create single pairs while neutrons and higher energy hadrons can create large cascades of Frenkel pairs. Most of those pairs ($\approx 90\%$) recombine quickly, but a fraction remains stable [35]. The displaced atoms and vacancies form additional energy levels in semiconductor materials which leads to a change in the electrical properties of the material. The additional energy levels serve as recombination centers for minority carriers and therefore reduce their lifetime and diffusion length. This can, as an example, lead to an increased forward threshold in components like PIN diodes. Sufficient radiation fluence can even cause an inversion of the electrical properties of a n-doped semiconductor into a p-doped [36]. The unit of reference for this type of damage is 1-MeV equivalent neutron fluence.

5.1.2 Total ionizing dose effects

A much larger part of the energy lost by ionizing radiation is in the form of ionization. Valence electrons are excited to the conduction band and create electron-hole pairs in the material. For a short time every solid, even an insulator, will have a higher concentration of mobile charge carriers. Only a small amount of energy is necessary to create a single electron-hole pair. Therefore the energy of the radiation is less important for total dose effects than for atomic displacement. The creation of electron-hole pairs causes damage to MOS and bipolar devices through the accumulation of charges in oxide films. Electrons have a chance to escape non-conductive oxide films due to their high mobility preventing the recombination with another hole. This leads to a continuous increasing positive charge in the oxide layers. This charge alters the conductivity of all surrounding semiconducting layers

altering the electrical properties of the component. All effects caused by the build-up of charge in oxide layers are collectively called total ionizing dose effects [36]. The unit of reference for this type of damage is total ionizing dose (TID) measured in Gray (Gy).

5.1.3 Single-event effects

The source of soft errors in electronics is the energy deposited by ionizing radiation on their way through a semiconductor. For this purpose, it is unimportant if the radiation stems from cosmic rays, radioactive decay or stray radiation from a particle accelerator. By passing through the semiconductor lattice such radiation will inject energy into the material producing electron-hole pairs. Furthermore by collision with lattice atoms they can create a cascade of secondary ionizing particles which in turn create even more electron-hole pairs. The type of effect that is caused by this depends on what type of circuit is hit. But generally, if the produced charge is bigger than a device specific critical charge Q_{Crit} and the charge is collected by source or drain then the behavior of the device is influenced. A MOSFET transistor, as shown in figure 5.1, is one typical example for a component vulnerable to single-event Effects (SEE).




Figure 5.1: A transistor like the one shown in this figure is controlled by the voltage between source and gate. The electric field created by this voltage is used to either enrich or deplete the charge carriers in the area between source and drain. This switches the gate between a conducting and non-conducting state. A strike by an ionizing particle (proton, neutron or alpha particle) injects energy into the semiconductor along its path thus creating many additional free charge carriers. These charge carriers weaken the depletion zone separating the source from the drain contact. If the accumulated charge is greater than Q_{Crit} it leads to a radiation induced signal in the circuit.

How many charge carriers are liberated by a passage depend on the stopping power of the striking particle and the cross section for inelastic interaction with lattice atoms. A 10 MeV alpha particle has a stopping power of about $10 \text{ keV}/\mu\text{m}$. The energy to liberate a electron-hole pair is 3.6 eV so it can produce $2.8 \cdot 10^4$ electron-hole pairs/ μm equaling a charge of $4.5 \text{ fC}/\mu\text{m}$. With Q_{Crit} values reaching from 1 – 4 fC [37] this would be enough to cause a SEE if most of the charge would be collected. The representative unit for SEEs is the integrated hadron fluence above 20 MeV high energy hadron fluence (HEH). Hadrons of lower energy are not considered to be able to reach the sensitive areas of the device with enough energy to induce an SEE [38][39]. This energy threshold does not apply to neutrons [40][41]. While charged particles can directly liberate electron-hole pairs, uncharged neutrons cannot interact electromagnetically with the lattice. Instead, they collide with nuclei in the semiconductor creating several secondary nuclear fragments. Many of those are charged particles able to create ionization tracks in the semiconductor [42]. The amount of charge produced is often enough to surpass Q_{Crit} , but the chance of a collision of the neutron with semiconductor nuclei is very small. Therefore the amount of neutrons necessary to create the same amount of soft errors as with protons/alphas is 5 times as big [43]. Therefore a contribution from 0.2-20 MeV neutrons is added to the HEH forming the equivalent HEH fluence (HEHeq). Furthermore contributions from thermal neutrons [44] [45] and the increasing cross sections for hadrons in the GeV range [42] have to be taken into account.

The value of Q_{Crit} is usually determined empirically by injecting current pulses of different strength into the device until it malfunctions. Using Q_{Crit} , the single-event rate (SER) can be described by:

$$N_{SER} = C \cdot \Phi \cdot A \cdot e^{-\frac{Q_{Crit}}{Q_{Coll}}} \quad (5.1)$$

The constant C is dependent on several factors like process technology and circuit design style. The Flux Φ is the ionizing radiation in the environment of the circuit and area A the vulnerable area of the circuit. Q_{Coll} is an efficiency ratio of produced charge carriers to collected charge carriers.

A radiation-induced signal in a circuit does not always cause an error in the system. An induced signal inside the combinatorial part of a circuit is called a single-event transient (SET). To actually have some influence on the system, a SET has to be captured by a latch or flip-flop. For this to happen the SET has to reach a flip-flop element at the same time as a clock pulse

5.2.

Another possible SEE is a memory element that is directly struck by ionizing radiation. A SRAM cell like the one shown in figure 5.3 is a typical




Figure 5.2: Depending on which part of a digital circuit(e.g. an FPGA) is hit by radiation the type of error occurring changes. A hit in the combinatorial parts causes a momentary spike in current which propagates further into the circuit. A long-lasting change in the state of the system only happens if this spike gets clocked into a synchronous part of the FPGA. Most SETs do not transform into SEUs this way because the SET has to reach the synchronous element without attenuating too much and during the setup-hold time. Therefore faster systems with a higher clock frequency have a higher chance of creating SEUs this way.

memory element often used for storing data. It has two stable states and does not change its state as long as transistors M_5 and M_6 are closed. An ionizing ray producing more than Q_{Crit} in free charge carriers in the bulk of one of the internal transistor $M_1 - M_4$ can cause the SRAM cell to flip from one stable state into the other one and remaining there. While this does not cause any permanent damage to the cell the data stored within is lost and could cause serious malfunctions of the whole system depending on the importance of the bit.

According to their influence on the system, SEEs can be separated into soft and hard errors. Soft errors do not permanently influence or damage the system. The following SEEs are soft errors [36]:

- Single-Event Transient (SET): Current pulse due to free charges. Can cause non-transient effects if it is captured by a register.
- Single-Event Upset (SEU): Injected charge changes the state in a bi-stable element like an SRAM cell. The change is permanent until the cell is rewritten.
- Single-Event Functional Interrupt (SEFI): Single-event possibly related to an upset in a control register that temporarily locks the device.

Hard error have a permanent effect on the system usually by destroying it unless specific countermeasures are in place. Examples for hard errors are:




Figure 5.3: This figure shows a basic schematic of an SRAM cell. As long as the transistors at M_5 and M_6 are closed, the system should remain static. But ionizing radiation can induce a voltage in the base of one of the internal transistors $M_1 - M_4$. This leads to a permanent switch from one of the two stable states into the other which then will remain stable again. Such an event is called an SEU (Single-event Upset).

- Single-Event Latch-up (SEL): Due to the placement of several transistors in the shared substrate of an integrated circuit, ionizing radiation can activate an additional parasitic thyristor. This short circuit between transistors can destroy the device if no protection measures in place.
- Single-Event Burnout (SEB): Increased current between drain and source of a MOSFET destroys the component.
- Single-Event Gate Rupture (SEGR): Destruction of the gate dielectric of a MOSFET.

Chapter 6

Radiation levels at CERN

There are three primary sources of the radiation in the vicinity of the LHC:

- The beam interactions at the collision points of the accelerator experiments. This radiation field mainly affects the experiments and the DS sections. The effects scale with the produced luminosity.
- Beam losses at the collimators. The affected areas are the vicinity of the collimators. The strength of this effect is highly dependent on the collimator settings and the beam parameters.
- Interaction of the beam with residual gas atoms. This can occur at any point of the accelerator. The effect is proportional to the beam current and the pressure of the residual gas.

The impact of these three sources depends on the location of the electronic systems. Figure 6.1 shows the location of all installed 600 A QPS which are the focus point of this project. The most critical locations of those shown are the RRs 13,17,53,57,73 and 77 (areas with increased tunnel radius and shielded compartments for electronics) which are either in the vicinity to the two high-luminosity experiments Atlas and CMS or close to a beam cleaning section (Collimators). Table 6.1 shows the high-energy hadron fluence (HEH) and the TID measured in 2015 and predictions for the year 2016 and the HL-LHC phase. There is already a significant increase in radiation level at some parts of the accelerator between 2015 and 2016. The HL-LHC phase will operate with even higher luminosity and also increased beam current. Therefore, electronic systems will have to be designed to withstand a radiation level 10-20 times higher than the current one.




Figure 6.1: The red circles in this figure mark the positions of the 600 A QPS. Positions with the highest radiation level are the RRs next to the two high luminosity experiments Atlas and CMS and to the collimators in point 7. [46]

Critical Areas	Measured 2015	Prediction 2016	Prediction HL-LHC
	HEH/cm ²	HEH/cm ²	HEH/cm ²
	Gy	Gy	Gy
RR13	$1.44 \cdot 10^7$	$7.15 \cdot 10^7$	$58.4 \cdot 10^7$
	0.029	0.14	1.14
RR17	$1.05 \cdot 10^7$	$3.3 \cdot 10^7$	$2.7 \cdot 10^7$
	0.021	0.066	0.54
RR53	$1.44 \cdot 10^7$	$25.35 \cdot 10^7$	$207.2 \cdot 10^7$
	0.029	0.509	4.16
RR57	$2.1 \cdot 10^7$	$32.5 \cdot 10^7$	$265.7 \cdot 10^7$
	0.042	0.652	5.33
RR73	$1.31 \cdot 10^7$	$1.65 \cdot 10^7$	$13.5 \cdot 10^7$
	0.0262	0.11	0.9
RR77	$0.52 \cdot 10^7$	$1.65 \cdot 10^7$	$13.49 \cdot 10^7$
	0.0105	0.033	0.27
Integrated Luminosity	$4.2 \text{ fb}^{-1} y^{-1}$	$\approx 36.7 \text{ fb}^{-1} y^{-1}$	$\approx 300 \text{ fb}^{-1} y^{-1}$

Table 6.1: Radiation levels past, present and future in comparison to the annual integrated luminosity. The values of 2016 were extrapolated from a set of data from the first 6 months of the year and the remaining days of beam time. The values for the HL-LHC phase were predicted using the data from 2016 and the annual integrated luminosity that is the goal during the HL phase.

Chapter 7

Development of radiation-tolerant systems

In areas with radiation, it is necessary to use electronic systems that have been prepared for such an environment. Two different ways to improve the radiation tolerance of a design are explored in this work. One is to build a system only from components whose radiation-tolerance is known and sufficient for the specific application. This method increases the tolerance of the complete system against cumulative type damage and hard SEEs. The other is to protect a system against soft SEEs with specific design and coding techniques. The following sections will focus on discussing both methods in detail. Furthermore, the importance of radiation tests to verify the radiation tolerance of systems and single components will be explained. Guidelines and proper procedure for radiation tests are presented at the end of this chapter.

7.1 Development and Protection Strategies

As discussed in chapter 5 there are two categories of damage radiation can cause to electronic systems, cumulative damage (DD and TID) and stochastic damage (SEEs). The vulnerability of existing devices with regard to cumulative damage is dependent on their internal design and can therefore not be improved further. In order to protect against this damage type, it is necessary to choose hardware with high inherent radiation tolerance. Section 7.1.1 will give an overview of known radiation-tolerant technologies and the general design rules for radiation-tolerant systems. The severity of stochastic errors is heavily influenced by the choice of hardware. Also for stochastic errors, it is possible to prevent or mask the influence of soft errors with certain design and coding choices. The techniques that are able to mitigate SEEs

are discussed in section 7.1.2.

7.1.1 Radiation-Tolerant Technologies

When radiation damage to electronics was first recognized as a problem, it was mostly focused on space applications. The high-energy radiation present in space was known to cause degradation of semiconductor circuits and functional disturbances due to SEEs. Today the effects of radiation damage is a problem that has to be taken into consideration even at sea level. The decrease in feature size due to Moore's law has led to a continuous decrease of Q_{Crit} in all electronics systems. This is partially counteracted by the decrease in the area that is used for charge collection, but the SER for electronics systems are constantly on the rise. For consumer electronics, this is generally of low importance, but in the case of safety-critical applications it has to be taken into consideration. Especially avionic systems that receive a higher radiation load have to be protected. The radiation at sea level and in flight does not cause permanent damage to electronics due to cumulative damage. Therefore cumulative damage is only a problem for areas with artificial radiation fields like nuclear reactors and particle accelerators or outer space. Before designing a radiation-tolerant system the type of hardware used has to be selected. The three possible systems types are defined as following:

- Commercials of the Shelf (COTS) are commercially available solutions for a certain task. They range from single components like an ADC to extensive systems like a ventilation system or an oxygen detector.
- Custom systems based on COTS are systems designed specifically designed for a purpose with radiation-tolerance in mind. Their base components are COTS that have been chosen for their above-average radiation-tolerance.
- Radiation-Hardened Components are components or full systems that were specially designed to withstand radiation by their manufacturer.

The usage of COTS saves the development time and cost necessary to develop a custom solution. Therefore COTS are advantageous to use as long as the project specifications allow them to be used. In case of the electronics needs to be faster or more radiation tolerant, COTS cannot be used. COTS systems are usually feature rich and come with some level of support by the manufacturer. Their disadvantage is that their behavior under radiation is unknown because they were not designed with radiation in mind. Most manufacturers do not provide any information or offer any support in concerning

radiation-tolerance. Therefore to enable their usage in critical systems it is mandatory to conduct irradiation campaigns to determine their radiation tolerance. The goal is to determine their tolerance in regards to cumulative damage and the existence of any occurring SEE error modes. The existence of any hard SEEs usually prevents their usage in areas with radiation above those at sea level. Assuming no such SEEs exist, a COTS system can be used in areas around the accelerator that are considered radiation free. This entails an environment where the annual fluence of high-energy hadrons and the annual TID is not higher than $10^6 \text{ cm}^{-2}\text{y}^{-1}$ and $10^{-3} \text{ Gy} \cdot \text{y}^{-1}$. This is a radiation level comparable to the avionics level. Additionally, there are several methods of system error mitigation that can be used to limit the influence of soft SEEs. Remote resets of a system allow to remedy most effects of soft SEEs and hardware redundancy allows for a high availability of the system.

In areas with a radiation level higher than the avionics level, COTS systems can no longer be used. In environments like this, the available solutions are the usage of radiation-hardened systems or custom systems based on COTS. Radiation hardened systems are designed especially for operation in environments like outer space. Therefore they are highly resistant to cumulative damage and have few or no SEE error modes. But due to the small size of the market for such components, their price is usually several orders of magnitude higher than those of comparable non-radiation hardened components. The market size also limits the selection of available components. This makes it difficult to acquire all required components for a project as radiation hardened. Custom Systems based on COTS are systems specifically designed for a purpose and are made up of normal COTS components. Designing the systems yourself allows for the introduction of hardware mitigation measures to increase error tolerance. Furthermore a system designed for a specific purpose can have a much higher reliability and speed than a COTS. All main LHC systems include both COTS as well as custom designed electronics. Some of the system, e.g. the experiments, were developed with radiation tolerant or even radiation hard components. While COTS are not designed to be radiation-tolerant, there exist large differences in radiation-tolerance between otherwise comparable components. By using COTS with high radiation tolerance, the resulting custom system will also have a high radiation tolerance. Because the components are COTS the price of the whole system remains low. As mentioned the manufacturers of COTS do not offer any information on the radiation tolerance of their products. Therefore it is necessary to qualify all used components and the final custom system using an irradiation campaign. Figure 7.1 presents an overview of the radiation levels




Figure 7.1: Relation between radiation level and usable system types. The radiation levels are given in the three units that are related to the different types of radiation damage (High-energy hadron fluence ($\text{cm}^{-2}\text{y}^{-1}$) for SEEs, total ionizing dose (Gy y^{-1}) for dose effects and 1 MeV neutron equivalent fluence ($\text{cm}^{-2}\text{y}^{-1}$) for displacement damage). Radiation levels higher than sea level can be separated into three levels of severity. In the lowest level, electronic systems are usually only affected by SEEs. COTS can be used if mitigation measures are in place and the radiation is in the lower half of this level. Using custom systems based on COTS is advisable to reduce the chance of SEEs influencing the system. In the next level, the radiation is high enough to cause cumulative damage to electronic systems. Custom systems that have been qualified to withstand higher doses of TID and DD can still be used. In the higher half of this level, the usage of radiation hardened components might be necessary. In the final level, the material of the electronic systems will incur damage. Electronic systems in such an environment have to be replaced in certain intervals. [42]

present in the LHC accelerator complex and the environments the different types of systems can operate in.

Field Programmable Gate Arrays

The centerpiece of many projects for digital signal processing and protection systems are Field Programmable Gate Arrays (FPGA). FPGAs are both vulnerable to cumulative damage and SEEs. Especially the memory cells that store the configuration of the FPGA can become the critical point of failure. SEUs changing the state of a memory cell in the configuration memory would change the algorithm running on the FPGA causing any number of possible errors in the system. FPGAs can be separated into three categories by how they store their configuration.

- Anti-fuse FPGAs: FPGAs using this technology use anti-fuses to fix their programming permanently. Therefore they can be only programmed once, but are also immune to any radiation-induced changes to their configuration.
- SRAM-based FPGAs: Most FPGAs currently sold use SRAM cells to store their configuration. They have a higher gate density than other technologies and lower prices compared to other technologies.
- Flash-based FPGAs: FPGAs that store their configuration on flash cells make up only a small percentage of the market. Flash cells have a high resistance against radiation

Anti-fuse technology loses some of the advantages of FPGAs to achieve greater robustness with respect to radiation. During programming of an Anti-fuse FPGA the unwanted connections of the interconnects are permanently altered. Because it is missing the configuration memory it is, therefore, immune to the corruption of the algorithm due to radiation-induced SEUs. Such a configuration is non-volatile, eliminating the need for a reprogramming of the FPGA after a power-cycle. The omission of an external storage device and the on-board configuration memory saves space both on the chip and on the circuit board. The drawback of an Anti-fuse FPGA is that it can only be programmed once. After that further algorithm changes are not possible and would require an exchange of the FPGA. While the configuration of an anti-fuse FPGA is immune to SEUs its logic cells are based on SRAM technology. Therefore the combinatorial elements and registers that form the algorithm are vulnerable to SEUs.

SRAM-based FPGAs are the most frequently used type of FPGA. They utilize SRAM cells to store the configuration of the FPGA. Because SRAM cells are also used in many other electronics systems, FPGAs using them

profit from a large amount of research being done on them. Thanks to being on the forefront of technology, they offer the highest density of logic blocks. A drawback is that their configuration is volatile. Every time during activation of the system the configuration data has to be re-sent from an external device. Most of the time this will be an onboard flash memory which takes up additional space in the system. Therefore such an FPGA is not immediately available after power-up. Furthermore, all SRAM-based technology is inherently vulnerable to data corruption by radiation. In the case of the configuration memory of an FPGA, an SEU would modify the interconnections between the logic blocks. This would alter the function of the FPGA until the configuration is rewritten. Newer generations of FPGAs include dedicated error correction measures to prevent the corruption of configuration memory. Furthermore, the decreasing feature size reduces the cross section of critical SEEs [47]. In devices that do not require a 100% availability, it is also possible to repair corruptions in the configuration by reprogramming the FPGA (Scrubbing). Newer generation FPGAs offer several methods from full reconfiguration to partial reconfiguration of only the affected parts [48].

Flash-based FPGAs use flash cells to store their configuration. This gives them an inherent immunity to SEUs, and their configuration is non-volatile. Due to less research being done on Flash FPGAs, they are 1-2 generations behind SRAM technology. Furthermore, their size and power consumption are considerably higher than their SRAM counterparts. Nevertheless, the inherent immunity of flash cells to SEUs is beneficial for all projects where radiation-tolerance is an issue. As with anti-fuse FPGAs, the immunity to SEUs does not cover the logic cells of the FPGA. A flash cell has a limit how many times it can be programmed. Using them as registers for fast calculations or similar tasks can quickly exhaust this limit. Because of this and due to their lower speed they are unsuitable for the use in logic cells. The high voltage required to erase the state of a flash cell damages the oxide layer between floating gate and the data line. Depending on the specific flash type after 100,000 to 1,000,000 write operations a flash cell will cease functioning. Therefore flash FPGAs use SRAM technology as well for the flip-flops of their logic cells.

7.1.2 Radiation-tolerant Design

The utilization of radiation-tolerant components while designing a system can already increase its radiation tolerance to a high degree. Depending on the radiation level in the area where the system will be installed, the danger of cumulative damage might not be an issue. If all components have been

qualified using irradiation campaigns and their dose and fluence limit are above the expected level reached during their operation time or lifetime, no further action has to be taken. If the expected dose during the lifetime of the system surpasses the dose limit, determined during the irradiation campaign, measures have to be taken to replace the system before it fails. But besides cumulative damage, it is usually not possible to completely remove the chance of SEEs influencing a system. Therefore mitigation measures have to be included into a system to mask the effects of any occurring SEEs. A basic method of detecting and masking error in a system is redundancy. An example would be reading out data from a memory element. Redundancy, in this case, would be either reading the required data from different addresses of the same memory element or from several independent memory elements. A voting logic then compares all read data with each other. By reading data two times it is possible to detect a read error if the read samples do not match. By reading data three times it is possible to determine the erroneous read and mask it by transmitting the majority result. In the case of memory elements like flash cells which are immune to SEUs it is also possible to use time redundancy. This entails reading the data from the same memory element three times after each other and comparing the result with a majority voter.

Redundancy is often used for FPGA algorithms. Triple modular redundancy (TMR) is a common method to protect the logic cells of an FPGA against SETs and SEUs. TMR creates three instances of every element of the algorithm followed by a voter that compares the three results (see fig. 7.2). As long as only one instance is faulty, the error is completely masked and does not interfere with the operation of the system further downstream. This technique is not perfect because there is a chance of an SEU or SET occurring simultaneously in two instances and the existence of the voter as a single point of failure. Fortunately, the chance of a simultaneous occurrence of identical errors is quite low, and the voter has due to its small and simple structure a very good radiation tolerance. Other methods include protecting memory elements (configuration memory, block RAM or fabric memory) with error correction codes (e.g. Hamming codes). Such correction methods can fix a single changed bit similar to TMR, but can only detect multiple changed bits. Recent FPGAs have advanced this protection method for their configuration memory by protecting single bits from several data words with one error correction code. This reduces the chance of unrecoverable errors caused by one ionizing particle hitting to adjacent memory bits of one data word.




Figure 7.2: TMR triplicates every part of the original algorithm creating 3 independent instances. The results of the three instances are then evaluated by a majority voter and the result with at least 2 votes gets passed on. If every bit is stored in 3 places and every calculation is done 3 times, then a single SEU cannot influence the result. The majority voter is still a vulnerable point, but because its structure is both small and simple, its radiation resistance is usually very good. Furthermore, an ion hitting two TMR parts is not impossible so TMR can not completely prevent SEUs.

7.2 Irradiation Campaigns

7.2.1 Motivation for Irradiation Campaigns

As explained in chapters 6 and 7, many electronic systems of the LHC will operate in environments with a radiation level higher than sea level. Some areas are shielded enough to allow the usage of COTS systems. Other areas will require custom systems based on COTS and radiation hardened systems. In all cases it is required to know how a system will operate in the environment it will be used in. In the case of radiation hardened components, the manufacturer provides information about the tolerance of the components against one specific radiation field. This information might not be comparable to the radiation field the system will operate in. Therefore in most cases, it is required to launch an irradiation campaign to quantify the radiation tolerance of those systems. To design a custom system based on COTS it is required to have information about the radiation tolerance of the required components. Though COTS are not designed for radiation tolerance, irradiation campaigns have shown big differences in the tolerance of two otherwise comparable components. While certain design features of components like high resistance to thermal influence or electric noise can hint at a higher radiation tolerance it is principally impossible to determine their radiation-tolerance without an irradiation campaign. Therefore frequent irradiation campaigns are necessary to acquire knowledge about components that can be used to design custom systems. The qualification of every system needs to be investigated before introduction into the accelerator.

7.2.2 Procedures of Irradiation Campaigns

The goal of an irradiation campaign is to understand how a component acts when subjected to radiation. This can be separated into its reaction to cumulative damage (total ionizing dose and displacement damage) and SEEs. For cumulative damage, it is necessary to determine how the electrical parameters of a device change in accordance with the increase in dose and neutron equivalent fluence. Depending on the device some may even break down after a certain dose. In this case the specifics of the breakdown and in which final state the component remains is interesting. The level of dose and fluence that is still acceptable depends on the target project specifications because requirement can vary greatly. For SEEs the target is to find all types of SEEs occurring and calculate cross sections for every type. A special focus is on finding any hard SEEs which would usually preclude a component from being used in an irradiated environment.

Irradiation campaigns acquire this information by irradiating the target component and monitoring its behavior. Due to the insufficient knowledge about the manufacturing process of most components, an in-depth analysis of the influence of radiation on specific layers of a component is very difficult. Therefore radiation tests monitor as many key values as possible during irradiation to try to gain understanding on the state of the component. The key values that are measured depend heavily on the type of component. In most cases, power consumption can provide valuable information about any changes in the component, and it reacts strongly in the case of a breakdown. Other examples for monitoring values are the gain and the forward threshold voltage of an amplifier and the on-state resistance and output voltage of a photoMOS. Depending on the type of component there will be other values that are interesting to measure.

The search of SEEs is significantly more complicated because they can alter the behavior of any part of a component. Destructive SEEs like SEL, SEB and SEGR are usually detected by measuring the current consumption. SEFIs are interrupts of the function of the device, so depending on the device, its operation has to be closely monitored to detect such interrupts. SEUs, and to a lesser degree, SETs can alter the state of bits in a device. In the case of a memory device, such alterations can be detected by writing a known pattern on the device and reading it out several times. SEUs and SETs can be found by comparing the read-out pattern with the one that was written onto the memory elements. Any changed bits would be the effects of an SEE. Such deviations from the pattern may even occur for flash cells that are immune to SEUs. Those are caused via influencing the read-out electronic producing an altered result. Data alterations can also occur if the component in question is a data transmitting device like an ADC. Data alterations can take the form of a single altered bit in a data word, but, depending on the inner workings of the component, it could also lead to a greater number of altered bits and data words. Again, this is detected by transmitting a known pattern and comparing the results.

Setup of an Irradiation Campaign

During the irradiation campaigns of this project, the device under test (DUT) was usually mounted on a special circuit board designed for testing. Examples for such circuit boards are shown in chapter 8. For efficiency reasons more than one DUT is usually tested during one time-slot at a test facility (test facilities are described in 7.2.2). To allow for the easy testing of several DUTs each was mounted on an adapter board that was connected to a base

board during irradiation. While the adapter board contains the DUT and any peripherals unique to this DUT the base board handles the power supply and communication of data. The baseboard also contains an FPGA which allows for additional configuration options during testing, handles the communication with the user control system, the data acquisition and provides any required input signals for the DUT. Due to the failure modes of a device being unknown before the test, it is required to acquire as many data about the device behavior as possible. This allows for an easier understanding of the behavior of the device under radiation after the analysis of the acquired data. Special care has to be applied to the data readout system. Usually, during the irradiation of a component, all readout electronics is positioned, depending on the irradiation facility, 20-30 meters away from the component. For digital data, this poses a problem because many typical transfer protocols like SPI/I²C can not be used over that distance. Electronic components that translate these protocols into those that exhibit a larger range of transmission like Ethernet, in turn, are vulnerable to radiation and can therefore not be used safely. Possible solutions for this problem are either to shield the vulnerable electronic systems or to use a transfer method that is tolerant against radiation. More details about data transmission can be found in the chapters detailing the irradiation campaigns that are part of this project.

Methodology of Irradiation Campaigns

Knowing the behavior of a component under radiation is necessary to use it for radiation-tolerant projects. Unfortunately, data about this behavior can easily become obsolete. The radiation tolerance may vary substantially even for one specific component if it was produced using different manufacturing processes or at a different manufacturing site. Furthermore, manufacturers may produce newer revisions of a component using different methods which could also greatly influence the radiation tolerance. Such changes are usually not published, which can lead to unexpected failures of radiation-tolerant equipment. In order to avoid such failures, radiation tests should be conducted before and in parallel with the development of a system. Figure 7.3 shows how the design of a system and the irradiation tests should optimally be interleaved.

Concerning the choice of the type of radiation for the irradiation, there are two methods for testing, using a mono-particle beam (one type of particle at a specific energy level) and using a mixed radiation field (a wide variety of particles of different energies). Mono particle beams have the advantage that they are more easily available. Testing facilities that can provide access to such a beam are discussed further in the next paragraph. Mono parti-




Figure 7.3: In order to take new revisions of a component and other changes of a component in consideration several irradiation campaigns should be conducted during development of the full system. If an adequate component for a project has been found, there should be an initial irradiation campaign testing a small sample of the target component. If the radiation-tolerance has been verified, a large batch (enough to last for the full project) should be bought. Sometimes, depending on the relation to the manufacturer it is possible to acquire the components from a single production batch which helps with keeping the tolerance levels equal. Following the initial irradiation test, it should be repeated with a sample from the large batch to account for possible changes in the time between both tests. If the available time for development is too short then the initial step can be omitted. In any case any large batches that are deemed unusable for radiation-tolerant projects can be utilized for systems in radiation free environments.

cle beams are especially preferable if the information required is focused on one type of radiation effects. A radiation field of gamma particles is the most easily available(usually produced by a ^{60}Co -source) and can be used to test for TID effects. Neutron beams can be used for DD effects and synchrotrons/cyclotrons providing a proton beam are best for testing for SEEs. A disadvantage of tests in mono-particle beam facilities is that the provided radiation is very different from the environment close to the LHC. The much wider energy range of radiation in the vicinity of the accelerator and a large number of occurring particles make this environment significantly more harmful for electronic components. Therefore the results of mono-particle beams have to be adjusted by a safety margin to account for the difference between test and real environment. Irradiation tests using mixed field are far more realistic, but there are significantly fewer facilities where such tests can be conducted than for mono-particle tests.

Testing Facilities

There are several external facilities that can provide a mono-particle beam for irradiation campaigns. Table 7.1 list some of them together with the beam type they can provide.

Facility	Beam Type
Fraunhofer INT	Gammas
	Neutrons
Paul-Scherrer Institute (PSI)	Protons
European Space Agency	Gammas
Centre Energie Atomique (CEA)	Neutrons

Table 7.1: Facilities that can provide a mono particle beam for irradiation campaigns.

Most irradiation campaigns for this project were carried out at the Paul Scherrer Institute (PSI) Proton Irradiation Facility (PIF) (see fig. 7.4) in Switzerland [49] [50]. The initial proton beam for PIF is delivered from the COMET cyclotron with the help of the primary energy degrader, which allows setting a few discrete initial beam energies in the range from 200 MeV down to 30 MeV. The beam is subsequently guided to the experimental area, where PIF is located. A further set of energy degraders is placed before the DUT to set a lower energy. For all conducted tests, the primary energy was set to 200 MeV. Beam particles with higher energy create more secondary ionizing particles and therefore produce more free charge carriers

in the material. Below a threshold of about 20 MeV, SEEs are usually not observable. Therefore all irradiation campaigns were conducted using the highest beam energy available. As mentioned the radiation field of the LHC contains protons with energies in the GeV range. In the RR sections (see chap. 6) which are most relevant for this work about 10% of all particles have energies above 690 MeV and 1% have energies above 2.8 GeV. The interaction cross section of high-Z materials like tungsten and hafnium does not saturate at energies of 200 MeV. Therefore, particles with energies higher than 200 MeV produce more secondary particles leading to an increased SEE rate. Cross section measurements at PSI will, therefore, produce results that are lower than comparable measurement in a more realistic radiation field. Simulations have calculated that a safety margin of 2 has to be applied to results from PSI test to account for high energy particles [42]. The typical flux of the PSI beam is about $1.7 \cdot 10^8 \text{ p} \cdot \text{s}^{-1} \cdot \text{cm}^{-2}$. This high flux allows for accelerated tests where the fluence of several years can be achieved in hours. The total integrated dose and fluence are provided by the facility. Figure 7.4 shows the PIF testing site. By removing the circular collimator at the end of the beam line, the radius of the beam can be increased to a level allowing the irradiation of a whole circuit board. Systems bigger than about 10 cm cannot be tested at PSI.

Facilities that enable mix field irradiation campaigns with similar environments to the LHC vicinity are not readily available. Furthermore, sometimes there is a need to test larger systems with a volume that cannot be tested in most external facilities. Therefore CERN has put continuous effort into providing a testing facility at CERN that allows both for large volume testing and testing in a realistic environment. The first testing facility of this type was CNRAD. In 2007 the gallery close to the CNGS neutrino target was cleared from electronic equipment due to the intense radiation showers produced by the fast extraction of the SPS beam. After installing radiation monitors this gallery could be used for irradiation tests. It provided a HEH flux between $2 \cdot 10^{10}$ - $2 \cdot 10^{11} \text{ cm}^{-2}$ per week and a dose between 3-30 Gy per week. CNRAD was able to reproduce the energy spectrum of the shielded areas as they exist in the LHC tunnel. The flux could be varied by changing the distance from the collision target. Furthermore, the radiation field produced by hitting a target with the primary proton beam could be used to test for all three interactions of radiation with electronic systems. The disadvantage of CNRAD was that it was purely a parasitic test site to the CNGS experiment. Access was only possible during the few stops of the experiment. Furthermore, the control of irradiation campaigns has to be placed in the CNGS control room leading to cable lengths of 1 km. Therefore CNRAD could not be a final solution and was closed in 2012.




Figure 7.4: Top: Schematic of the Proton Irradiation Facility at PSI. Bottom: Irradiation site of PIF.

Because previous mixed-field irradiation facilities at CERN had several disadvantages, the dedicated irradiation CHARM (Cern High-energy Accelerator Mixed field/facility) site was designed and was commissioned at the end of 2014. CHARM is not only able to reproduce the radiation field occurring in the vicinity of the LHC but also several atmospheric and space environments. Furthermore, the test site is big enough to house whole accelerator control systems or parts of airplanes and cars.

Inside the CHARM irradiation area (see fig. 7.5) a 24 GeV proton beam hits a target, and the produced secondary radiation scatters through the whole room. By changing the target material, choosing a specific test position and extending up to 4 movable shields it is possible to create a wider variety of radiation field. The flux CHARM can produce reaches between 10^7 and $10^{12} \frac{n}{\text{cm}^2 \cdot \text{s}}$. Test setups are brought inside the test facility using a remote handling to limit radiation exposure of the users. Because the mixed radiation field is close to the real radiation conditions at the LHC, no safety margins have to be applied to the results.

7.2.3 Irradiation Campaign Guide

An irradiation campaign is a complex process requiring a multitude of choices to be made ahead of the testing. Failures in planning can hinder the success of the project. The following section is meant as a collection of advice to help circumvent some of the most common problems.

The adapter board with the DUT, the baseboard, all cable connections between user control and test setup and the test algorithm running on the FPGA on the baseboard have to be designed precisely to ensure that all necessary properties of the DUT are measured. Most of these components can be tested by replicating the test setup in a lab and conducting a dry-run of the whole irradiation campaign. Special care should be taken to verify that the cable connections and all communication protocols passing them can operate over a distance of up to 30 meters.

The user control for the test should be made simple to operate. Testing conditions (length and time of day) tend to be not optimal so an easy to use system helps to prevent mishaps due to difficult conditions. Furthermore, time slots for irradiation campaigns are fairly limited so retesting due to a failed first test might not be possible.

The user control and data acquisition usually runs on a PC that is connected by cable to the test setup. For creating this software most programming languages are possible options as long as they support the necessary communication interfaces to the test setup. LabViewTM is a tool that can be used




Figure 7.5: (Top left): CHARM irradiation chamber 18 positions, 4 shield layers and different target materials allow the production of varied radiation fields. (Top right and bottom): Different perspectives on the CHARM installation.

for creating both a control and a data acquisition interface. LabView™ allows creating programs very quickly by offering pre-made functions for most interfaces. During this work LabView™ was used for both control and data-acquisition. As an end result LabView™ was deemed suitable for those tasks unless the data-rate to be acquired is too high. To account for data rates in the range of MBaud future irradiation campaigns will use a data-acquisition written in Python.

While creating a format for the data that is acquired from the DUT, it is important to include reliable time stamps. Because the data about the radiation level is supplied by the testing facility, is necessary to relate both sets of data with each other. Absolute timestamps are preferred against time stamps relative to the start of a test to avoid possible confusion.

Depending on the type of irradiation campaign the requirements for the user control differ. A test with a high radiation flux, like it can be conducted at PSI, can be finished in a few hours. Therefore personal supervision of the test is possible during the whole test. In this case, the user control can profit from a multitude of indicators indicating the state of the measurement and in the case of problems the possibility to reconfigure the test to fix problems. In the case of a test with lower flux, like at CHARM, where the test takes a whole week of continuous irradiation, the whole test setup requires the ability to operate without the user. Still, the user needs to have the ability to check upon the test state regularly. This includes the ability to reset the test remotely. Manual control of the test hardware is usually not possible during tests at CHARM. Besides data acquired from test modules like ADCs or flash cells, it is helpful to log as many additional signals that can give an understanding about the state of the DUT. After a successful irradiation campaign, all data has to be combined and evaluated to categorize the DUT's behavior under radiation. For the evaluation of data samples LabView™ is again a useful tool. In the case of high amounts of data (>100 MByte) LabView™ loses most advantages due to the effort required to optimize the project to be able to analyze such an amount of data in acceptable time. In such cases, a self-written program in a language of choice might be more advantageous.

Chapter 8

Development of a new 600 A quench detection system

8.1 State of the previous QDS System

The previous QDS system is structured as described in section 4.2.3. The circuit board used in this system (see fig. 8.1) is equipped with a digital signal processor (DSP) implementing the QDS algorithm.




Figure 8.1: Circuit board of the previous 600 A QDS system(10x20 cm)

During the initial development of the quench detection electronics, it was assumed that there would be hardly any radiation at the location of the electronics. Therefore no protection measures are included into the design. Especially the DSP, which is based on code-execution from SRAM cells, is highly vulnerable to SEUs. In 2012 the previous system produced 16

errors that could be traced back to radiation effects. Using the fluence data from 2015 (see section 6), one can extrapolate that during the HL-LHC phase of the experiment the previous system would produce about 70 errors every year. Under such a condition the availability requirements for HL-LHC cannot be fulfilled. Therefore a new system has to be designed that can operate under the expected levels of HEH fluence.

8.2 Requirements for the new QDS System

Chapter 2.2 described the requirements on the availability of the whole system for the HL-LHC phase. The limits for premature beam aborts and downtime of the individual systems making up the LHC (QPS, cryogenics, power converters vacuum and others) are given in [10]. The whole QPS system has a radiation error budget of nine faults that force a beam dump every year. A generally agreed target during the development of the new 600 A QDS system was to limit radiation-induced errors at a maximum of 2 annually for the full system. The new QDS system was installed beginning 2016 in all locations where problematic radiation levels are to be expected (see fig. 6.1). About 200 circuit boards were used to replace the previous system in these locations. For an average HL-LHC HEH fluence of $6 \cdot 10^8 \text{cm}^{-2}$ this would require an error cross section of less than $3.3 \cdot 10^{-9} \text{cm}^2$ for the full system. Furthermore, even though the annual TID in the shielded RR sections even during HL-LHC is not very high, all components should have a TID limit of several hundred Gy to preclude any possible trouble due to cumulative damage.

8.3 Design of the new System

During the start of this work, the new circuit board was already in the later stages of development. The new system still retains the quench detection method of the previous system. But to increase the radiation-tolerance all used components were such whose radiation-tolerances were already known or would be tested in the scope of this work. The first important step for increasing the radiation-tolerance was to replace the DSP by a flash-based FPGA from MicrosemiTM. The ProAsic3TM is an FPGA based on flash technology that had shown high radiation tolerance in tests [51] and was already part of many radiation-tolerant systems at CERN. The usage of flash cells means that the configuration of the FPGA is secure, while the fabric flip-flops can be protected using TMR.

The new circuit board also featured the AT25DF041A Flash Memory, the AQW210EHA PhotoMOS and the ADUM3402 digital isolator whose irradiation campaigns are discussed in later chapters. Another important component of the new system was a radiation-tolerant high-resolution ADC. The used quench detection method requires high-resolution measurement of voltage and current over the magnet circuit because bigger step sizes in their signal could easily lead to large spikes in the derived signal (see Eq. 4.1). Such spikes can produce results higher than the voltage threshold and trigger the QDS logic. The ADS1281 [52] was chosen for this project because there existed already encouraging results from two previous irradiation campaigns [53][54]. It comprises a $\Delta\Sigma$ -Modulator, a programmable filter composed of a Sinc filter and four Finite Impulse Response (FIR) filters, a calibration block and a serial interface for communication with the rest of the system (see fig. 8.2).




Figure 8.2: Building blocks of the ADS1281

The first irradiation campaign was conducted at the H4IRRAD facility, a mixed field test site at CERN [55], while the second was conducted using the proton beam at the PSI [49]. In the first test at H4IRRAD [53], the ADC was stimulated by a test signal. The output of the $\Delta\Sigma$ -modulator was monitored for SEUs in the bit stream, bypassing the digital filter block. Furthermore, the power consumption was measured to search for SELs. No SEUs or SELs were found during the measurement. Therefore the cross sections for these events cannot be high enough to result in a full event when multiplied with the fluence. From this the upper limits for the cross sections of both error modes could be determined as $\sigma_{SEL} < 8.06 \cdot 10^{-13} \text{ cm}^2/\text{device}$ and $\sigma_{SEU} < 1.31 \cdot 10^{-12} \text{ cm}^2/\text{device}$.

The second irradiation campaign [54] searched further for SELs. Moreover, the modulator output stream was monitored for breaks in the stream (single-event functional interrupts; SEFIs). No SELs were detected. The cross section for SEFIs was calculated to be $\sigma_{SEFI} = 9.38 \cdot 10^{-13} \text{ cm}^2/\text{device}$.

Due to the cross sections of all detected error modes being low, the ADS1281 was a suitable candidate for the 600 A QDS system. But the previous tests focused only on the $\Delta\Sigma$ -modulator while omitting the other blocks. During

usage as part of the 600 A QDS, these blocks would not be bypassed. Therefore additional irradiation campaigns were required to analyze their behavior and find any additional error failure modes.

8.4 Irradiation Campaigns on the ADS1281

8.4.1 Setup of the first campaign

The first irradiation campaign on the ADS1281 was conducted at PSI. The test hardware used during the campaign comprises a small circuit board with two ADCs and all required peripherals and a baseboard. During the irradiation, the board housing the DUTs is attached to the baseboard. This modular approach was chosen to enable testing multiple instances of a DUT without replacing the baseboard. The baseboard provided the necessary peripherals for data acquisition and communication to the controlling computer.




Figure 8.3: The baseboard is equipped with a radiation-tolerant FPGA running the necessary algorithms for data acquisition from the DUT and communication with the rest of the test setup. Configuration from and data transmission to the controlling computer are transmitted via an I²C data line.

The baseboard was equipped with a Microsemi[™] ProAsic3 FPGA which houses the necessary algorithms for providing stimulus data to the DUTs. The acquisition of data was done using a self-written module also running on the FPGA which used a combined I²C[™]/Ethernet data line to connect to a PC in the control room (see fig. 8.4). This PC connected to this data line using a self-written LabView[™] control and data acquisition program. The

FPGA on the baseboard would also receive a certain amount of radiation during the irradiation campaign so the usage of a flash-based FPGA with high tolerance was required. The algorithm on the FPGA was protected with TMR to prevent radiation-induced errors in the SRAM-based logic elements of the FPGA.




Figure 8.4: To protect most instruments from radiation only the DUT and the necessary I²C™ to USB and USB to Ethernet converter are inside the beam chamber. The beam was collimated to limit radiation exposure of the baseboard excluding the area of the DUTs. The stimulus signal for the DUT and the power supply are both provided from the radiation free control room. A LabView™ program is used to store the data acquired via the Ethernet data line.

During irradiation, the ADC was stimulated with a slow low-amplitude sine wave. The ADC sampled the stimulus signal with a frequency of 2 kHz (due to usage of pin-mode the highest value available), and the resulting data stream was fully acquired and stored for later analysis.

8.4.2 Results of the first irradiation campaign

Two ADCs were irradiated with a TID of 300 Gy. Until the end, none of the two ADCs experienced a functional breakdown. Unfortunately during this measurement the LabView™ data-acquisition produced several critical

errors, leading to complete data loss for many of the irradiation phases. These errors were most likely due to radiation-induced errors in either master (the I²C[™] to ethernet adapter) or slave (the test board) of the I²C[™] connection inside the beam chamber. The I²C[™] protocol allows for states where both master and slave wait for a response from each other. Normally unreachable, such states can occur if either master or slave is switched to a different state due to an SEU in the state register. Only data during irradiation with a TID of about 120 Gy was collected, which was deemed too small a sample for a proper statistical analysis. In a qualitative analysis of the data, three distinct error modes were found inside the ADC output signal.

In 28 cases the ADC signal showed single samples with vastly different values than the surrounding samples. These δ -function-like errors ranged from small to maximum signal amplitude spikes in both directions.

Second most frequently with twelve detected events were multi-sample errors as shown in figure 8.5. These errors also ranged from small to full signal amplitude. Their width remained fairly constant with a length of 40-50 samples.




Figure 8.5: Example for one of the twelve multi-sample errors from the first irradiation campaign. Their average length was between 40-50 samples.

Least frequently with only six detected events were temporary stops in the




Figure 8.6: Example for a temporary stop of the ADC data flow. Five similar events were detected during the irradiation campaign. The average length of these stops were 64 samples with a spread of ± 3 samples

data stream produced by the ADC. One of the ADCs would stop producing new samples and restart its function after 32 milliseconds (64 samples with a sampling frequency of 2 kHz, see fig. 8.6).

8.4.3 Conclusions from the first irradiation campaign

The first irradiation campaign on the ADS1281 was a partial success. It could be determined that the ADC was able to withstand a TID of about 300 Gy without functional breakdown. This confirmed the results of a previous test and the amount of dose tolerance is sufficient for QDS projects. Several failure modes were discovered in the behavior of the ADC under radiation which would prove critical for a component of the QDS. Furthermore, the radiation vulnerability of the data acquisition leads to an insufficient data sample being collected. Therefore, to acquire enough data to be able to calculate statistical sound cross sections for the failure modes, a second irradiation campaign would be necessary. To that end, the data acquisition had to be revised to prevent the loss of data due to irradiation.

8.4.4 Setup of the second irradiation campaign

While it was not possible to calculate meaningful cross sections due to the insufficient amount of data the observed failure modes would be critical for




Figure 8.7: This is the version of the 600 A QDS board used during the second irradiation campaign (22*10 cm). The red circle shows the approximate dimension of the collimated beam used in most of the tests.

the QDS system. Therefore it was necessary to conduct a second campaign to acquire more data about those failure modes and search for any other that were overlooked due to the insufficient amount of data. Furthermore, several problems had to be solved first before conducting a second irradiation campaign on the ADC. In the previous campaign, only a small part of the necessary amount of data was acquired due to serious failings of the data acquisition system. The most probable source of these errors was the radiation vulnerability of the I²CsTM data-line. Therefore parallel to the I²CsTM line, which would still be used for configuration of the device, a dedicated SPITM data readout line was created. Furthermore, this campaign was also used to test the prototype of mitigation measures against the failure modes discovered during the first campaign (see chap. 8.5.2).

The second irradiation campaign was also conducted at PSI 7.2.2. Instead of a dedicated test board, the current version of the 600 A QDS board was chosen as a DUT. This would allow for a test of the whole system instead of only the ADC. The 600 A QDS board is equipped with two ADCs and a MicrosemiTM ProAsic3 FPGA. Due to the radiation tolerance of the FPGA, it was possible to use the uncollimated beam, for some of the irradiation steps, to test the rest of the board for unexpected vulnerability to radiation.

The setup for this second irradiation campaign was similar to the first one. Instead of one data line handling both data acquisition and configuration two dedicated data lines were used. The I²CTM data-line failures only occurred during data acquisition at high data rates so using it solely for the configuration of the FPGA was possible. Using the SPITM new data line a vector containing 13 bytes of data was transmitted. 12 of these bytes were




Figure 8.8: Schematic representation of the setup of the electronics for the second irradiation campaign. Data and configuration line have been split to increase the stability of data acquisition. Furthermore, besides the previously used sine-wave generator, for stimulus signals a generator producing close-to-real magnet signals was included. These magnet signals mimic the behavior of voltage and current inside a magnet for a normal accelerator cycle and allow for realistic stimulus signals.

used to transmit four three-byte signals containing either raw data from both FPGAs, the final voltage U_{Res} (see eq. 4.1) and several intermediate signals from stages of the calculation and the mitigation measures. The last byte was used to transmit several flags providing information about the state of the algorithm. The transmission happened in bursts of a high number of these vectors instead of a continuous transmission. Between data bursts, the produced data was stored in several TMR-protected FIFOs on the FPGA. The general throughput was high enough to read out all data produced by the ADC running at a sampling frequency of 2 kHz.

During irradiation the ADCs were stimulated with two types of signals, a slow sine wave like during the first irradiation campaign and simulated magnet ramps that mimic voltage and current behavior during an accelerator cycle.

8.4.5 Results from the second irradiation campaign

During this irradiation campaign, the new data-acquisition performed without any problems. Data during a total dose of 3378 Gy of irradiation was acquired. This data was collected over the irradiation of three DUTs with a total of six ADCs. The three DUTs were irradiated with 650, 628 and 640 Gy respectively. During irradiation with a dose of 1898 Gy the ADC was stimulated with a sine wave and during 1480 Gy with the simulated magnet ramps.

Besides verification of the previous results, the campaign also focused on the effect of cumulative damage on the ADC. To evaluate the changing behavior of the ADC with increasing dose, the acquired data was first searched for drift in the signal level. The position of the peaks of the stimulating sine function in parts of the campaign was measured, and the resulting plot can be seen in figures 8.9 and 8.10. Due to changes in the setup between measurements, some data sets had a different offset and were removed from this figure. As can be seen, both the maximum and minimum show a distinct drift downwards starting from a dose of about 250-300 Gy. For most applications in the QDS, a dose budget of 250 Gy is more than sufficient due to annual doses usually being not higher than 10 Gy. Furthermore, the signal measured (a ± 100 mV sine wave) had a maximum range of ± 10 V, whereby the signal only drifts by 0.002% of this range. This is more than the maximum error from the specifications of $200\text{ }\mu\text{V}$, but holds no consequence for the 600 A quench detector.

The measurement of current consumption of the device alone was not possible during these tests because the full QDS board for the 600 A magnets was used as a test board. The small current consumption of the ADCs and any possible drift would be too small in comparison to the current consumption of the entire board.

Besides measuring the signal drift, the data was again searched for any distortions in the signal and any other possible error modes of the data acquisition. Figure 8.11 shows one of the measurements where the stimulus signal was a magnet ramp. The many distortions here are mainly single-sample errors with a few multi-sample errors.

53 ADC stops were also found in the data. The higher amount of data as compared to the first irradiation campaign allowed getting a clearer picture of the length distribution of the stops. Table 8.1 shows the number of stops in relation to their length. It is clear that besides some aberrants the stop length is fairly fixed around 62 samples.

Table 8.2 shows the cross sections of all failure modes, the amount of counted failures and the size of the data sample used to calculate the mea-




Figure 8.9: Drift of the maximum value of a measured sine wave during irradiation for both ADCs. The red lines are polynomial fits to the data.




Figure 8.10: Drift of the minimum value of a measured sine wave during irradiation for both ADCs. The red lines are polynomial fits to the data.

sured in HEH fluence that passed the DUT during acquisition. During this campaign, a fourth failure mode was discovered. Two times during the irradiation the gain of a signal increased significantly (see fig. 8.12). This change persisted until a restart of the ADC.

Unmitigated all of these failure modes would trigger the quench protection




Figure 8.11: ADC output signal during irradiation while being stimulated with a simulated magnet signal. The stimulus signal is a 19-minute long simulated magnet ramp. The measured signal is distorted in many places due to radiation errors. Most distortions of the signal are only one sample wide while some span over multiple consecutive samples.

Stop Length in Samples	Qty. of Stops
59-66	29
32	2
≈ 100	2
$\approx 10^5$	2

Table 8.1: Distribution of the lengths of stops observed during irradiation of the ADC. Most stop lengths were inside a narrow area between 59 and 66 samples while a few aberrant were observed with a length up to 10^5 .

logic leading to a beam abort. According to the radiation levels shown in chapter 6 and with a system comprising 400 ADCs, these failure cross sections




Figure 8.12: Permanent change of the signal gain.

	Qty. Errors	Fluence	Cross section in $\text{cm}^2/\text{Device}$
Single-sample errors	191	$1.85 \pm 0.09 \cdot 10^{12}$	$1.03 \pm 0.05 \cdot 10^{-10}$
Multi-sample errors	97	$1.85 \pm 0.09 \cdot 10^{12}$	$5.24 \pm 0.03 \cdot 10^{-11}$
ADC Stops	35	$6.24 \pm 0.31 \cdot 10^{12}$	$5.60 \pm 0.28 \cdot 10^{-12}$
Gain Jumps	2	$6.24 \pm 0.31 \cdot 10^{12}$	$3.20 \pm 0.16 \cdot 10^{-13}$
Combined			$1.61 \pm 0.05 \cdot 10^{-10}$

Table 8.2: Error cross sections of all discovered error modes. ADC stops and gain jumps are based on a higher amount of fluence due to the possibility to find them in the data set collected with magnet ramps as stimulus. Single and multi-sample errors could only be detected in the data sample produced with a sine wave as stimulus.

would cause a significant amount of beam aborts every year. A calculation using an average of the radiation levels given in table 6.1 predicts about one failure every year for 2015, about 8 during 2016 and during the HL-LHC phase about 60 failures every year. Failure rates this high would negatively impact the overall availability of the accelerator and limit the integrated luminosity that can be achieved. Therefore a dedicated mitigation system is required to limit the impact of those failures to an acceptable level.

8.5 Failure mitigation

8.5.1 Failure analysis

To design a mitigation system for all discovered failure modes, it is necessary to understand how they occur. Because none of the failures occur outside of radiation it should be possible to trace all of them back to radiation. The observed drift of the signal (see figs. 8.9 and 8.10) is clearly an effect of cumulative damage which can not be mitigated. All other failure modes are singular events only lasting a short amount of time. Therefore they most likely stem from SEEs in one of the functional blocks seen in figure 8.2. Among the four failure modes discovered single-sample errors were the easiest to be understood. The output of the delta-sigma modulator of the ADC is a bit stream that is converted into a data word by the sinc-filter block. This data word is then further transmitted downstream through the ADC undergoing several operations in the other ADC blocks. Between the different functional blocks, the data word is stored in SRAM output registers. An SEU flipping the value of one of the bits in the data word would change the value of the data word and produce a short spike in a stream of otherwise continuous values. Therefore it can be concluded that single-sample errors are the product of SEUs in the data word in positions of the ADC downstream of the filtering block.

Multi-sample errors are closely related to single-sample errors due to their origin being SEUs as well. Digital Finite Impulse Response (FIR) filter like those used in this ADC will upon stimulation by a δ -function-like signal output their filter kernel. An SEU switching one bit of a data word produces a signal very similar to the δ -function. The analysis of multi-sample errors showed a clear similarity between the most common form of those errors and the kernels of the four-stage FIR filter of the ADC. Figure 8.13 shows the 4 kernels (available from the data sheet [52]), although the first one is not visible due to its low amplitude. Therefore it is very likely that multi-sample errors are created by a superposition of the filter kernels of the ADC after stimulation by a δ -pulse created through an SEU.

To verify this theory the FIR filter stage was simulated using LabView™ as a four-step FIR filter using the kernels from the datasheet. Using this simulation artificial SEUs were induced in the stages before the FIR filters and their transformation were observed. Figure 8.14 shows a comparison between a measured multi-sample error and one produced by simulation. The similarity between both curves suggests the plausibility of the theory. A definite proof is unfortunately not possible because the inner functionality of the ADC was not made public.




Figure 8.13: Filter kernel used in the ADC to reduce unwanted high frequencies. The first kernel has an amplitude too low to be visible next to the other three.




Figure 8.14: Comparison between a simulated multi-sample error and a measured one. The red curve was created by replicating the ADC FIR filter stages using LabView™. This filter was then stimulated with δ -like pulses to produce this multi-sample error. The shift between both curves was included for better visibility of the individual curves.

The short stops of the data flow of the ADC are difficult to analyze. From their nature, they can be clearly categorized as single-event functional interrupts. One of the previous campaigns on this ADC already found SEFIs in the output of the delta-sigma modulator [54]. The cross section deduced

from the previous experiment is 5-6 times smaller than the one found in this work, but it is still highly probable that both failures are described by the same event. The implementation of this specific delta-sigma modulator is not public, so further investigation of this failure mode is not possible.


The change in the gain of the signal can be easily traced back to the calibration block of the ADC. According to the configuration of the ADC offset, gain and other values of the signal can be manipulated in this block. The ADC has two configuration modes. In register mode, all configuration is stored in registers in the ADC. Those values can be modified using the serial interface of the ADC. In the second configuration mode, pin-mode, all configuration is done by tying most of the ADC pins to a fixed signal level. This mode allows access to only a few of the configuration options of the ADC. The registers of all other options are set to standard values. During both irradiation campaigns, the ADC was set to pin mode in order to increase the radiation tolerance by fixing the values of important configurations. While in pin mode options like gain and offset are not accessible by the user, but SEUs can still change the values in the correlating registers which may lead to gain changes as observed.

8.5.2 Failure mitigation

Digital filtering

To limit the effects of the found failure modes, dedicated mitigation measures for each failure mode have to be included into the design. The observed failure modes can be separated into two categories: failures that can be mitigated by filtering the ADC output signal and those that have an effect too strong to mitigate it by filtering. Single- and multi-sample errors are of the former type. Due to their origin being SEUs in the ADC, their occurrence cannot be prevented. Therefore the only feasible method is to remove them from the ADC output signal before it is used for quench detection. To this end a chain of digital filters used to reject electrical noise was improved and extended to also mitigate those two types of failures. Figure 8.15 shows the chain of digital filters after it had been extended. The decimation of the signal was included as a new first step and the length of all other filters was chosen for optimized filtering of radiation failures.

Figure 8.15 shows starting from the left the filter chain that combines error mitigation with noise suppression of the ADC signals and the calculations necessary for quench detection (see. eq. 4.1). The main filters used are median and moving-average filters together with a decimation of the ADC signal. The decimation of a signal takes several consecutive samples and



$$U_{Res} = U_{Diff} - L(I_{Dcct}) \cdot \left(\frac{\delta I_{dcct}}{\delta t} + \frac{\delta U_{Diff}}{\delta t} \cdot \frac{1}{R_{Par}} \right)$$

Figure 8.15: Filter chain to mitigate single and multi-sample errors. The input signals are the Voltage (U_{Diff}) and Current (I_{Dcct}) measured at the magnet circuit coming from the left direction. Those measured signals are then filtered using several steps of decimation, median and moving-average filters. Following this, the difference between measured and calculated voltage is outputted as U_{Res} . With the parallel resistor (R_{Par}) being the only resistive part this value should be close to zero.

computes their average. Averaging over multiple samples has the advantage of lowering the amplitude of oscillation type errors like multi-sample errors. This works especially well if the error contains points both above and lower than the undisturbed signal level. Furthermore, by combining 8 samples into 1 the length of disturbances like multi-sample errors is also divided by 8 which is essential for the median filter. The disadvantage of decimation is that by reducing the sampling frequency of the signal the system response time is slightly increased.

The median filter is similar to a shift register wherein each step the oldest sample is replaced by a new one but within the median filter, the contents of the register are sorted according to value. Each cycle the output of a median filter is the sample in the middle of the shift register, the median,

which blocks extreme values because they are never the median during the time they are inside the register. The ability of a median filter to eliminate spikes from a signal depends on the width of its register and the widths of the spike. A filter with a width of n can only eliminate spikes with a length of $\frac{n-1}{2}$ perfectly. Any longer spike can not be eliminated completely because the $\frac{n}{2}$ sample of the spike will become the median. Single-sample errors can be removed from the signal easily but multi-sample errors have a length around 54 samples. Even after decimation, they will be around 6 samples long which would require a median filter with a length of 13 to filter. Due to a lack of resources on the FPGA and due to the requirement of keeping the response time as low as possible the median filter in the final version was shorter than that (length 11 for U_{Res} and 9 for I_{Dcct}). The ability of such a short filter to eliminate multi-sample errors longer than half its width depends on the presence of error samples on the other side of the baseline of the signal. In a median filter, those samples would be sorted into the other half of the median filter and therefore not occupy the same space most of the error samples would. Because a multi-sample error oscillates between high and low values the existence of samples on both sides of the baseline is probable but also heavily dependent on the position of the decimation windows. As can be seen in figure 8.16 a shift of the decimation window by one sample can result in one error spike that can be filtered and one that can not be completely eliminated.

The last component used during filtering is a moving-average filter. The main ability of a moving-average filter is to eliminate the noise of a signal by averaging over a number of consecutive samples. This ability is crucial because any noise would increase after derivation of the signal. Remaining error signals that could not be fully eliminated by the median filter can only be decreased in amplitude by the moving-average filter. This also reduces the amplitude of the resulting error pulse in U_{Res} to keep it below the threshold of ± 100 mV.

Filter Resource and Delay Cost

During the design of the filter the resource cost and the delay each filter component causes have to be accounted for. Because decimation decreases the data rate by a factor of 8 the delay it causes is equal to 8 times the original interval between samples $t_{\text{Delay,Decimation}} = 8 \cdot 250 \text{ us} = 2 \text{ ms}$.

The delay of a median filter is proportional to $\frac{n-1}{2}$ with n being the width of the filter. A strong change in the signal only gets propagated after filling one-half of the median filter with new samples. In case of the filter system the delay of the U_{Diff} median is 10 ms and for I_{Dcct} 12 ms. The delay of a




Figure 8.16: The two spikes are the results of the decimation of the same error found in the I_{Dcct} signal but with the decimation window shifted by one sample. While very similar the crucial difference between the two spikes is that the left one consists of 4 samples higher than the baseline while the right spike consists of 5 samples of that kind. The rest of the samples of the error pulse is of a value lower than the baseline. A median filter with a length 9, like it is used for the I_{Dcct} signal, will be able to eliminate the left spike completely while not being able to eliminate the first point of the right spike. The result will be a 5 sample plateau with the value of the least extreme of the 5 spike samples.

moving-average filter is dependent on the width of the internal shift register, the input signal and in the case of trigger applications the target threshold. A moving-average filter propagates every change in its input signal without any delay, but because the first changed sample is averaged together with the rest of the yet unchanged samples in the shift register, the output signal rises slowly until the shift register is completely filled with new samples. Therefore, depending on register length, the amplitude of the new samples and threshold level it takes several sample cycles before a change higher than the threshold gets propagated. Due to these many influencing factors, only a worst-case delay can be calculated which is the length of the filter multiplied by the decimated sample interval of 2 ms. For long moving average filters like the one in the I_{Dcct} line (width 256) this can theoretically create a delay in the range of half a second, but because magnet quenches produce large changes in I_{Dcct} , the real delay is much shorter. While not a digital filter, the

derivation of the signals in each data line also causes a delay which is the δt of the derivation. Both lines use a δt of 16 ms.

As can be seen in figure 8.15, the delay length in the three branches of the filter system is different for every branch. This leads to a phase difference when computing U_{Res} which will cause some distortions in the signal. But the amplitude of these distortions is not high enough to cause spurious triggers. Concerning resource costs of the filter components, the decimation and the moving-average are very cost efficient. Both filters require a shift register and a division for averaging the contents of the register. By choosing a register length which is a value of 2^n , the division can be done by a simple bit-shift operation. The necessary registers can be implemented using RAM memory which is available in much higher amount than registers in the FPGA fabric. The implementation of the median filter used is tuned for fast execution and therefore requires an above average amount of hardware resources. In the current implementation of the filter system, the median filters utilize 18.8% of the logic elements of the FPGA. Previous versions of the filter system decimated the signals only by a factor of four which, in turn, required longer median filters with a length of 18-20. Median filter of this dimension already utilized about 28% of the FPGA. Therefore, we decided to double the decimation factor.

Filter Simulation

To verify the ability of the described filter system to mitigate all distortions in the signal a LabView[™] test bench was created. This test bench mimicked all elements present in the filter chain. As input for the two data channels U_{Diff} and I_{Dcct} simulated magnet signals measured during the second irradiation campaign were used. This data sample was taken while irradiating the test board with a dose of 790 Gy.

Focus during the simulation was also the influence of a shift of the decimation window on the ability to mitigate one specific signal distortion. Therefore all data samples were used eight times while shifting the decimation window by one each time. The results of this simulation showed that the filter system is sufficiently able to mitigate signal distortions. The simulation showed that due to the much lower dynamic range of the U_{Diff} no observed distortions in this line could cause spikes in the U_{Res} signal surpassing the threshold. Distortions in I_{Dcct} can cause such spikes depending on the position of the decimation window.

ADC stop mitigation

If the signal of one ADC stops this will, in most cases, lead to fulfilling the trigger condition of the QDS. Unless the magnet current is constant for a time, a stop would lead quickly to a misalignment of the two signals. Fortunately, the data about ADC stops has shown that most of them last only about 62 samples. The latest versions of the system use a sampling frequency of 4 kHz after switching the ADC configuration method from pin mode to register mode (see. chap. 8.5.2). Therefore a 62 sample stop only lasts 15 ms. With the required reaction time of the QDS being 100 ms, an ADC stop of this length is not inherently critical. But with one ADC stopped the still changing data from the other ADC would quickly create a trigger condition for the QDS. Therefore, all triggers from the QDS have to be prevented during the duration of the stop. To achieve this, the result of the derivation of the stopped signal is held constant until the stop recovers. Furthermore, once a stop recovers, the new samples could have a large difference to the ones remaining from before the stop. This difference would cause a huge spike in the derivated signal causing a QDS trigger. To mitigate this effect the derivation is held constant not only until the stop recovers but until all samples from before the stop have been cleared from the system. The derivation uses a derivation interval ∂t of 16 ms which, even together with the stop length, is well below the reaction time.

In order to mitigate the few stops longer than the standard length, a system was included in the ADC VHDL interface to restart the ADC in the case of a stop. For this purpose, the flow of samples from the ADC is monitored to detect missing ones. In the case of a missing sample, either the reset or power cycle pin of the ADC is toggled to initiate a restart. The restart time after toggling one of these pins is the same as the standard stop length, so no additional time is lost in this procedure. If the ADC does not resume function after this time, the restart is attempted a second time. If both tries fail, the interface sends a critical failure notice to the QDS leading to a beam abort.

Mitigation of configuration corruption

While corruption of the configuration of the ADC was the rarest of the discovered error modes, it has the highest probability of causing a QDS trigger. Even small changes in gain or offset could trigger the QDS at some point. To mitigate this failure mode it is necessary to detect any corruption in the configuration and overwrite it with the proper settings. This required a switch from the pin mode of the ADC to the register mode due to the register not

being accessible in pin mode. The switch to register mode allows access to the registers using a serial interface and also allows to run the ADC at its highest sampling rate of 4 kHz which is not usable in pin mode. This increase in sampling frequency was helpful for the digital filtering of multi-sample errors because the decimation factor could be increased without losing any more reaction speed. The read-out of the whole configuration of the ADC is a fast process and can be done between the readout of each 4 kHz sample. In the case of a corruption, the register is completely rewritten. For this process, the ADC has to be interrupted, but the total time for rewriting and restarting the ADC is lower than the necessary reaction time of the QDS.

8.5.3 Mitigation results

To varying degrees, all mitigation measures were tested upon their efficiency. A first prototype of the digital filter chain was already in use during the second irradiation campaign. The results of this chain were already very good mitigating all single and most of the multi-sample errors. ADC stops were also partially mitigated during the second campaign by including the masking feature to prevent any triggering during a normal-length ADC stop. Table 8.3 shows the results of the mitigation measures used during that campaign. As can be seen, most failures could be mitigated leading to a combined error cross sections well below the one required.

Error Type	Qty.	Qty. after mitigation	Annual Fluence in cm^{-2}	Cross section in $\text{cm}^2/\text{Device}$
Single-sample	191	0	$1.85 \pm 0.09 \cdot 10^{12}$	0
Multi-sample	97	2	$1.85 \pm 0.09 \cdot 10^{12}$	$1.08 \pm 0.54 \cdot 10^{-12}$
ADC stops	35	2	$6.24 \pm 0.31 \cdot 10^{12}$	$3.21 \pm 0.16 \cdot 10^{-13}$
Configuration	2	2	$6.24 \pm 0.31 \cdot 10^{12}$	$3.205 \pm 0.16 \cdot 10^{-13}$
Combined				$1.72 \pm 0.57 \cdot 10^{-12}$

Table 8.3: Results of the mitigation measures during the second irradiation campaign. Two multi-sample and two long ADC stops could not be mitigated. Mitigation against configuration corruption was not yet in place.

Further simulation of the filter chain revealed the relation between the possibility to mitigate a multi-sample error and the position of the decimation window. Simulations with the data from the second campaign and the newest version of the filter chain lowered the cross section of multi-sample errors for a QDS circuit board (not a single ADC) down to $1.54 \pm 0.25 \cdot 10^{-12} \text{ cm}^2/\text{Circuit Board}$. This includes the findings that due to the lower dynamic range of the

signal no observed multi-sample error can pass the mitigation measures if it occurred in U_{Diff} .

A third irradiation campaign in the mixed field facility CHARM was conducted to verify the remaining mitigation measures. Unfortunately, this campaign suffered from severe mistakes during preparation and produced few results. Due to missing configuration on the DUT the mitigation measures for multi- and single-sample errors were not active. Therefore no additional data about the effectiveness of these mitigation measures could be collected. What could be observed was that the technique to restart the ADC using either the reset or power cycle pin as soon as a stop is detected proved ineffective. Most likely only an external power cycle has the ability to interrupt an ADC stop. Unfortunately, this is not possible with the current hardware setup.

Corruption of the ADC configuration could not be observed during the third campaign. The general ability of the interface to change ADC configuration was verified outside of radiation. Furthermore, a detailed simulation using a self-written hardware model of the ADC was used to verify the proper timing of the mitigation measure during ADC operation.

Combining all the results from the latest version of all mitigation measures produces a cross-section of $2.179 \pm 0.282 \cdot 10^{-12} \text{cm}^2$ for one QDS board and $4.358 \pm 0.564 \cdot 10^{-10} \text{cm}^2$ for the whole QDS system comprising 200 boards. This is an order of magnitude lower than the requirements for the system. These results are derived from data acquired during irradiation campaigns using a mono-energetic proton beam which usually causes less severe radiation errors than a mixed field with energies well into the GeV range. But even after applying the safety margin for such a beam (a factor of two) the requirements of the system are still fulfilled by a wide margin.

Chapter 9

Irradiation campaigns in support of further radiation-tolerant designs

As discussed earlier to retain the ability to create radiation-tolerant systems it is required to have access to a database listing the behavior of components under radiation. The following sections list the results of the irradiation campaigns of several components that either have been used in projects (including the 600 A QDS) or are planned to be used in future projects. The following devices were all irradiated at PSI using a 200 MeV proton beam. Most campaigns used a new data-acquisition system and a new test-board which are described in 9.1.2.

9.1 The MAX1162 and MAX11100 ADCs

9.1.1 Description of the Component

The MAX1162 [56] is 16 bit ADC based on the successive approximation method (SAR) with a maximum sampling frequency of 200 kSPS. Since 2009 it is part of the symmetric quench detector, but it was never tested for its radiation tolerance in detail. The MAX11100 [57] is an ADC with very similar specifications to the MAX1162, albeit with different pin-out and a lower price. It will be used for the new main quadrupole QDS which will be installed during long shutdown 2 (2018-2019). The successive approximation method that both ADCs use and its generally simpler structure (see fig. 9.1) was also hoped to be resistant to single event effects (SEE).

Device Name	Device Type	Irradiated Dose [Gy]
MAX1162	ADC	500
MAX11100	ADC	500
MAX6250	Voltage Reference	500
ADR421	Voltage Reference	380
ADR435	Voltage Reference	500
AT25DF041A	Flash Memory	520
ADUM3402	Digital isolator	580
INA128	Instrumentation Amplifier	360-500
AQW210EHA	PhotoMOS	160
AQV251A	PhotoMOS	300

Table 9.1: Overview of irradiated DUTs




Figure 9.1: Building blocks of the MAX1162/11100. The track-and-hold block samples the input signal and holds its output signal constant for a specific minimum time. The SAR ADC block contains the successive approximation algorithm. The control block enables or disables the ADC and the serial data link transmits the digital output value to the outside.

9.1.2 Irradiation Campaign on the ADCs MAX1162/11100

The irradiation campaign of the MAX1162 and the MAX11100 was conducted at PSI. The test utilized the new universal radiation tester baseboard (see fig. 9.2). The board contains a ProAsic3E™ FPGA [58] for control and data acquisition and two RS485 ports [59]. One is dedicated to the acquisition of data while the other can be used for communication between user control and FPGA. RS485 has the advantage of being much more stable

in irradiation environments. The DUT only requires a very simple adapter component, which was presumed to be very radiation-tolerant, to receive and send RS485 communication. Furthermore, the usage of a differential signal for communication allows a high noise resistance and enables communication over a distance of 30 m. In comparison to the previous method of using SPI/I²C this eliminates the need for a converter module to Ethernet close to the beam which was a possible cause for problems during previous irradiation campaigns. Furthermore, the maximum data throughput of 2.5 MBaud allows to acquire a large amount of data from the DUT. All used components of the test-board are radiation tolerant components.

Figure 9.3 shows two schematics (top MAX1162, bottom Max11100) of the adapter boards holding the respective DUTs. The adapter boards are connected to the backside of the baseboard. Due to the maximum diameter of the PSI beam of a few centimeters, the electronic components on the baseboard receives a small dose during an irradiation campaign only. The radiation tolerance for most of the components, like the operation amplifier used, are known, or they are even immune to radiation.

As with the ADS1281, the test was conducted by stimulating the ADCs with a sinusoid wave and reading out all sampled data for later analysis. Several ADCs of both types were tested together. Additionally, the power consumption of the ADCs was monitored for latch-ups and dose-related drift. As shown in figure 9.4, the current consumption remains constant for a dose of up to 200 Gy, which again is fully sufficient for operation in most of the irradiated areas in the LHC. The spike around 100 Gy could not be categorized because the ADC data showed no reaction at that point in time.

As with the ADS1281, the drift of the maximum and minimum of the measured sine wave was analysed for both the MAX1162 and MAX11100. Figures 9.5 to 9.8 show the results of maximum and minimum drift for both ADCs and a close-up view on the area between 0 and 220 Gy. These figures were created by separating the ADC data in blocks equaling one second of data and determining the maximum and minimum values of these blocks. Both ADCs behave very similarly regarding their signal drift. In the area between 0 to 160 Gy for the maximum and 0 to 120 Gy for the minimum, the signal level remains constant. Until a dose of 220 Gy, the signal drops by an average of 0.2% and further to 1.5% until the end of the measurement. The drift is significantly larger than the offset and gain error specified in the data sheet at a maximum of 1 mV. Nonetheless, the drift in the first 100 Gy is sufficiently low and a dose budget of over 100 Gy remains sufficient for QDS projects where dose-rates even during HL-LHC will be well below 10 Gy a year in most areas.

Analysis of the ADC data found only one failure mode in the form of




Figure 9.2: The baseboard used during the irradiation campaign is equipped with a ProASIC3E FPGA [51] [58] which is known for its radiation-tolerance (25*27 cm). An adapter board carrying the DUTs is attached to a connector on the back side of the board (location of the hole). The beam is collimated to a maximum diameter of 5 cm focused on the position of the DUTs. With the collimator, the radiation level in other parts of the base board remain low. Furthermore, the components were selected for their radiation tolerance. The two connectors at the bottom left allow for transmission of measurement data and control of the FPGA to and from a PC using an RS485/UART interface.




Figure 9.3: Schematics of the adapter boards containing the DUTs ((A) MAX1162, (B) MAX11100). These boards are connected to the baseboard 9.2 during irradiation.

single-sample spikes in the output signal. The cause of such a spike are either SEUs in the SAR (Successive Approximation Register) or a register inside the serial block (see fig. 9.1). Figure 9.9 shows the total number of errors in relation to the dose. These figures were created by counting all errors found in the data and plotting their sum in relation to the total dose. Errors within an interval of 5 mGy were summed up and for better visibility in the plot, the total number of single-sample errors was capped at 193,520. All devices under test (DUT) showed similar behavior by having a low error rate up to a total integrated dose of 220 Gy. At some point beyond that, the error rate increases exponentially until most measured samples are faulty. All DUTs experienced this increase in error rate at a total integrated dose of 500 Gy, although none of them had a complete functional breakdown. Because the only error mode is single-sample spikes, they can be easily filtered by using a median filter. This would allow for the usage of the ADCs without additional redundancy measures in areas with annual radiation doses of 1-10 Gy.




Figure 9.4: Current consumption of the circuit containing ADCs. This change occurred while irradiating the MAX1100 up to a dose of 360 Gy.





Figure 9.5: Drift of the maximum of the MAX1162 ADC data (top) and a close up view of the range until 220 Gy. This drift is caused by to the influence of cumulative damage to the ADC silicon fabric.





Figure 9.6: Drift of the minimum of the MAX1162 ADC data (top) and a close up view of the range until 220 Gy. This drift is caused by to the influence of cumulative damage to the ADC silicon fabric.




Figure 9.7: Drift of the maximum of the MAX11100 ADC data (top) and a close up view of the range until 220 Gy. This drift is caused by to the influence of cumulative damage to the ADC silicon fabric. The distortion at 130 Gy was caused by the test-board.




Figure 9.8: Drift of the minimum of the MAX11100 ADC data (top) and a close up view of the range until 220 Gy. This drift is caused by to the influence of cumulative damage to the ADC silicon fabric.




Figure 9.9: Integrated error distribution in relation to the total dose for MAX1162 (left) and MAX11100 (right). These figures were created by counting all errors found in the data and plotting their sum in relation to the total dose. Errors inside an interval of 5 mGy were counted together and the total amount of errors was capped at 193,520.

9.1.3 Conclusion

The results of the previous chapter show that both ADCs display sufficient radiation-tolerance to be used for their respective projects (symmetric quench detector and main quadrupole QDS). The current consumption of the MAX1162 and the MAX11100 remained constant until about 200 Gy, and the current consumption of the entire board used during the test of the ADS1281 showed no significant increase or drop. No latch-ups were detected for all DUTs. Both ADCs experienced a drift of their output signal with increasing dose. Starting at 100 Gy the growth rate of the deviation increases significantly. A maximum deviation of 1.5% is reached until the end of the measurement at 420 Gy. During the first 100 Gy the deviation of the signal is very low so the component can be used for QDS projects. Besides dose tolerance, the QDS system requires a very low error rate for the full system. The QDS system aims for maximal two beam aborts a year due to radiation-induced errors. Critical errors are those that force an abort of the beam, blind the system or require intervention in the accelerator. None of the ADCs experienced a complete functional breakdown requiring an intervention to replace it. The MAX1162 and the MAX11100 showed a very low vulnerability to SEEs. The only error mode observed was single-sample spikes, which can be easily mitigated. Therefore, it can operate as part of the QDS without significantly increasing the total error rate of the system. The ADCs showed a high tolerance to TID up to a limit of 300 Gy. These are very good values because the expected total dose for the QDS system is much lower.

With all this considered, it can be concluded that both ADCs are sufficiently radiation-tolerant to be used as part of the QDS.

9.2 Voltage References MAX6250, ADR421 and ADR435

9.2.1 Description of the Components

Voltage references are required for high-precision voltage measurements. But usually, the precision of references degrades due to the influence of cumulative radiation damage. In order to find a voltage reference with the least loss in precision, an irradiation campaign targeting three voltage references (MAX6250 with ± 1 mV precision [60], ADR421 ± 3 mV precision [61] and ADR435 ± 3 mV precision [62]) was conducted.

9.2.2 Irradiation Campaign on the Voltage References MAX6250, ADR421 and ADR435

The irradiation took place at PSI and used the new universal radiation tester baseboard (see fig. 9.2). Each adapter board contained 8 DUTs of one type. During irradiation the output voltage of each DUT was measured separately as well as the current consumption for the whole board. Figures 9.10 to 9.15 show the results of these measurements and compare the measured output voltage with the specifications given in the data sheet.




Figure 9.10: Measured output voltages of the eight MAX6250 DUTs in relation to the increasing TID. The red dashed line shows the maximal allowed deviation from the reference voltage to be expected due to production [60]. As can be seen the cumulative damage due to TID causes the output voltage to drop linearly until about 320 Gy.




Figure 9.11: The current consumption of the MAX6250 was not measured for each device individually but for all eight voltage references together. Similar to the output voltage the current consumption drops linearly during irradiation until about 400 Gy and remains constant until end of measurement.




Figure 9.12: Measured output voltage of the ADR421 in relation to the increasing TID. Due to beam stops during the measurement the TID is plotted as separate graph over the measurement time. The red dotted line shows the tolerance to be expected due to production [61]. As can be seen the cumulative damage due to TID causes the output voltage to increase.




Figure 9.13: The current consumption of the ADR421 was not measured for each device individually, but for all eight voltage references combined. Following the output voltage, the current consumption increases with the TID.




Figure 9.14: Measured output voltage of the ADR435 in relation to the increasing TID. Due to beam stops (in this figure the beam started only after about 100 sec of the measurement) during the measurement the TID is plotted as separate graph over the measurement time. The red dotted line shows the acceptable voltage error caused during production listed in the datasheet [62]. This voltage reference is sold in two accuracy categories with different output spread which are both plotted as dotted red lines. For this voltage reference two drift processes caused by TID oppose each other which allows the reference to remain within the specifications (A spec ± 6 mV and B spec ± 2 mV) for more than 200 Gy.




Figure 9.15: The current consumption of the ADR435 was not measured for each device individually, but for all eight voltage references together. The behavior of the current consumption mimics the overall behavior of the output voltage, staying fairly constant while the upward and downward drift negate each other and rising when the upward drift increases.

9.2.3 Conclusion

All tested voltage references show a clear drift of their output voltage as a function of the TID. But due to their design, the drift varies both in strength and sign. From the three voltage references, the ADR435 breaks its specifications only at 200 Gy which is the best value for all tested voltage references.

9.3 The AT25DF041A Flash Memory

9.3.1 Description of the Component

Many QDS systems require a non-volatile storage medium. The 600 A QDS system as an example uses three flash memory chips for the storage of threshold data. Inherently, flash memory cells are immune to bit flips by SEUs, but errors can still happen during the read process or the flash can stop to function due to cumulative damage. The AT25DF041A 6 Mbit flash [63] was expected to be compatible with many systems already in place because flash devices from the same manufacturer were already in use.

9.3.2 Irradiation Campaign on the AT25DF041A Flash Memory

The irradiation took place at PSI. To verify its radiation-tolerance, the flash memory was, in several steps, irradiated up to a TID of 560 Gy. Before irradiation and in between irradiation steps every byte of the flash memory was programmed with the LSB of its address which was read out at maximum speed during irradiation phases. The readout data was then compared to the pattern to search for read errors. The results are shown in figure 9.16




Figure 9.16: All six flash memories showed only very low error rates. During irradiation up to a TID of 520 Gy a total of $1.2 \cdot 10^9$ read cycles were conducted. A maximum of 20 reads contained at least one wrong value.

Up to a TID of 520 Gy a total of $1.2 \cdot 10^9$ read operations on each device were conducted, and the total error rate remained low. Shortly after passing the TID of 520 Gy, one device became completely inoperable. In the following irradiation step all remaining flashes produced large amounts of read-out errors. Consecutive read-outs showed a fixed relation between failed and successful read-outs. Such a scenario could be explained by sectors of the flash memory losing their ability to be read-out. Any readout of the whole flash memory would therefore result in errors for those sectors and successful reads for the rest, hence a fixed relation between errors and successes.

9.3.3 Conclusion

A dose budget of 520 Gy is fully sufficient for QDS projects, and the low amount of read errors can be mitigated by using three memory chips and voting over the read result. This is already done successfully for the 600 A QDS project. Another mitigation option, if time constraints allow for it, would be reading the data three times from one flash and voting for the correct result. Data corruption by SEUs is usually not possible with flash cells. Therefore, a temporal triplication can mitigate all errors during the read process. Unfortunately the component is no longer produced but a sufficient supply was acquired to finish all current projects utilizing it. A successor model exists (the AT25SF041 flash memory) but was produced using a different process. Therefore no predictions can be made how this model would behave under radiation.

9.4 The ADUM3402 Digital insulator

9.4.1 Description of the Component

The ADUM3402 is a 4-channel low power digital insulator [64]. Digital insulators are basic electronic components used to transmit signals between otherwise isolated circuits. It was tested to search for upsets in the output signal and to test the general dose tolerance.

9.4.2 Irradiation Campaign on the ADUM3402 Digital insulator

The campaign was conducted at PSI. During irradiation the insulator was stimulated with a 8 Hz square signal, while the output signal was sampled with 40 MHz and the result was compared with the expectation value. No upsets were found during irradiation to a TID of 580 Gy, and all DUTs remained functional to the end of the irradiation procedure. No detailed measurements on the current consumption were conducted, but the initial consumption of 50 mA for the whole test board increased to 140 mA after irradiation.

9.4.3 Conclusion

No upsets were found during the irradiation, and the total dose tolerance was shown to be very high. Therefore the component can be used in environments featuring ionizing radiation without additional mitigation measures.

9.5 The INA128 Instrumentation Amplifier

9.5.1 Description of the Component

The INA128 is a low-power general-purpose high-precision instrumentation amplifier. There are currently no projects in planning that will use it, but as a general-purpose component, it can be of use in many future projects.

9.5.2 Irradiation Campaign on the INA128 Instrumentation Amplifier

The irradiation campaign on the INA128 [65] instrumentation amplifier focused on determining the influence of TID on the output voltage and the input bias current (a automatic current in one direction present in bi-polar devices). The irradiation took place at PSI and used the new universal radiation tester baseboard (see fig. 9.2). Two boards each equipped with eight instrumentation amplifiers were irradiated with doses up to 500 Gy, and four of the amplifiers on each board were used to measure bias current or output voltage. Furthermore, the current consumption of each board was measured during irradiation. The results of the measurements are presented in figures 9.17 to 9.19.




Figure 9.17: Measurement of the output voltage of the INA128. Of the 16 total instrumentation amplifiers distributed over two test boards four on each board were used to measure the output voltage and four to measure the bias current. As can be seen in both figures (top board 1 and bottom board 2) there is a general upwards drift of output during irradiation. The devices exceed the specified limits (± 1.2 mV with a gain of 1, see [65]) at doses above 50 Gy. In this figure the total ionizing dose (TID) is given by the blue line and the right y-axis. All other curves depict the output voltage of the DUT measured in Volt (left y-axis).




Figure 9.18: Measurement of the bias current of the INA128. Of the 16 total instrumentation amplifiers distributed over two test boards four on each board were used to measure the output voltage and four to measure the bias current (top board 1 and bottom board 2). While there is a slight increase of the current even until 490 Gy the specifications (± 5 nA [65]) are met. The TID is given by the blue line and the right y-axis while all others curves and the left y-axis describe the bias voltage.




Figure 9.19: The current consumption of all eight instrumentation amplifiers on both boards was measured together (top board 1 and bottom board 2). A clear downward drift in relation to the TID can be seen. The TID is given by the blue line and the right y-axis while all others curves and the left y-axis describe the current consumption.

9.5.3 Conclusion

After a TID of 50 Gy, the DUT exceeded the limits defined in the specifications. This is a low dose tolerance in comparison with other tested devices. Nonetheless, in shielded areas, as described in section 6, this level of radiation tolerance is still sufficient. In areas with higher radiation level, a different choice of the device would be advisable.

9.6 The AQW210EHA PhotoMOS

9.6.1 Description of the Component

The PhotoMOS tested was the AQW210EHA [66], which was already used in QDS projects. The test focused on finding transients and to determine the total dose tolerance.

9.6.2 Irradiation Campaign on the AQW210EHA PhotoMOS

The irradiation campaign was conducted at PSI. During irradiation the PhotoMOS was stimulated with a slow square signal while reading its state with 40 MHz frequency and comparing the result with the expected state. A total of twelve PhotoMOS were tested in parallel and were irradiated up to a dose of 160 Gy. No transients were found, but the PhotoMOS started to experience functional breakdown starting at a dose of 120 Gy. At 160 Gy only three PhotoMOS remained intact.

9.6.3 Conclusion

The immunity to upset due to radiation up to 120 Gy allows for the component to remain part of current QDS projects. While the total dose tolerance is not very high, it is sufficient for most shielded areas in the vicinity of the LHC.

9.7 The AQV251A PhotoMOS

9.7.1 Description of the Component

Due to the high on-state resistance and the low dose tolerance of the AQW210EHA a second PhotoMOS, the AQV251A [67], was tested. Some planned systems

use long lines of several PhotoMOS. Such a line would have a too high resistance with a on-state resistance like the previous PhotoMOS. Besides finding possible transients, as with the AQW210EHA, this test also focused on the effects of TID on the on-state resistance and in which state the PhotoMOS would remain after a functional breakdown. The failure mode is an important information concerning the QDS safety.

9.7.2 Irradiation Campaign on the AQV251A PhotoMOS

The irradiation took place at PSI and used the new universal radiation tester baseboard (see fig. 9.2). For this test eight PhotoMOS were irradiated and their output voltage was measured (see fig. 9.20).

During irradiation the PhotoMOS were stimulated with a 8 Hz square-wave signal. Therefore, the output voltage would switch between high and low values as the PhotoMOS opens and closes. To calculate the on-state resistance only voltage values from the closed state are required. For this purpose the voltage data was filtered by replacing any points from the off stage with points from the previous on-state and eliminating any remaining points from flanks with a short median filter. The resulting data is shown in figure 9.21. Figure 9.22 shows the behavior of the on-state resistance during irradiation calculated from the filtered voltage data.




Figure 9.20: Schematic of the test setup used during the irradiation campaign on the AQV251. The PhotoMOS is connected in pattern A according to the data-sheet [68]. The voltage drop over the PhotoMOS is amplified by a factor of 100 before measurement.




Figure 9.21: Eight AQV251 PhotoMOS were tested in parallel and their output was measured to observe the influence of TID. The data for this plot was created by filtering the eight oscillating voltage signals and removing all samples no belonging to the on-state of the PhotoMOS. All PhotoMOS experience slow downward drift up to a critical dose above 170 Gy they experience functional breakdown in an permanently open condition. This state would activate the beam interlocks and is therefore a safe state.




Figure 9.22: The on-state resistance of the eight photoMOS was calculated from the filtered data in figure 9.21. The calculated on-state resistance is in good agreement with the values given in [68].

As can be seen from figure 9.21 and 9.22 the influence of TID on the output voltage and on-state resistance is low until a critical dose of about 175 Gy is reached. The AQV251 therefore shows a higher dose tolerance than the AQW210EHA which could only tolerate 120 Gy. By counting flanks in the unfiltered voltage data it was also proven that until critical dose the PhotoMOS never failed to switch state following the stimulus signal. When failing the AQV251 remains in an open state which is the preferred state for the QDS system.

9.7.3 Conclusion

The results of the AQV251 are better than those of the AQW210EHA with a higher dose tolerance and no upsets. The on-state resistance is also lower and stays constant until 175 Gy. This qualifies it for usage in a long interlock loop. The breakdown of the device leaves it in an open state which makes it a secure component for a beam interlock.

9.8 Conclusion and outlook

The chapters 8 and 9 describe all irradiation campaigns conducted by the team responsible for the 600 A QDS at CERN. These campaigns were necessary to acquire information about the radiation tolerance of components planned to be part of future projects or already part of some systems already at work in the machine. The information gained allowed for the successful design of the new 600 A QDS system and several other components with good radiation tolerance were found. These new components will be a valuable asset in the further development of radiation tolerant systems at CERN. Furthermore, after every irradiation campaign the tools and methods for data acquisition, planning and execution of campaigns, error handling and after campaign data analysis have been improved considerably. The experience gained during these 3 years will be an asset in further efforts to understand the behavior of yet untested components under radiation.

Chapter 10

Irradiation campaigns on the SmartFusion2/Igloo2 FPGA

The new 600 A QDS system uses the ProAsic3TM [69], an FPGA based on 130 nm flash technology as one of its main components. This FPGA has a long history of being used in radiation-tolerant projects and its behavior under radiation is well known [70](e.g. symmetric quench detector for main dipoles and quadrupoles DQQDS and a quench heater supervision unit DQHSU). Unfortunately, projects like the 600 A QDS already use more than 90% of the logic cells of the ProAsic3TM, which can cause problems during implementation. For future projects that could have an even bigger scope it is therefore required to find a new radiation-tolerant FPGA with higher capacity. MicrosemiTM, also producing the ProAsic3TM, has recently released the next generation FPGA SmartFusion2TM based on 65 nm flash technology. Using this FPGA allows to realize more complex and larger functions. The SmartFusion2TM is an FPGA that combines the logic elements of an FPGA with an integrated hardcore Advanced RISC Machines (ARM) processor. This combination would allow to combine the parallel processing of an FPGA with the ability to run C-code on the processor. Unfortunately, a processor is very vulnerable to SEUs and therefore not usable for radiation-tolerant projects. MicrosemiTM also offers the Igloo2, an FPGA based on the same layout as the SmartFusion2TM but without the ARM processor. Due to the lower prices the Igloo2TM would be the preferred component for radiation-tolerant projects. Irradiation campaigns were conducted using the SmartFusion2TM, but due to the similarity of their design all results should be valid for the Igloo2TM as well.

10.1 Irradiation campaign on the SmartFusion2

To verify its radiation tolerance, an irradiation campaign targeting the SmartFusion2[™] was conducted at PSI in collaboration with another team from CERN. The SmartFusion2 is a possible future component for radiation tolerant equipment. Besides the normal FPGA fabric the SmartFusion2 also contains a microcontroller subsystem which can be accessed from the fabric using an Advanced-High-performance-bus-lite (AHBL) (see fig.10.1). This subsystem contains additional memory elements like an SRAM and a Flash storage. As part of this irradiation campaign these two storages were tested for their vulnerability to radiation, focusing mostly on the effects of SEEs.




Figure 10.1: Schematic of the SmartFusion2 focusing on the division between FPGA fabric and micro-controller subsystem (HPMS). All user created algorithms are located in the FPGA fabric. Direct read and write access to elements of the HPMS from the FPGA fabric are not possible. All commands to elements like the flash memory (eNVM) and the eSRAM have to be passed through the Advanced-High-performance-bus-lite (AHBL). Therefore, an AHBL master has to be created in the FPGA fabric to connect to this bus.

The test board used during this campaign was the SmartFusion2 development-kit produced by Microsemi[™] (see fig. 10.2). This board was used to save the development time for a dedicated test board. The development kits originally ship with an outdated version of the FPGA. This version was already known to suffer from micro-latch-ups. Microsemi[™] claimed to have fixed

these latch-ups in the newest version of the FPGA. Therefore, the FPGA on the development kit was replaced with a current version SmartFusion2.




Figure 10.2: Development board of the SmartFusion2 by Microsemi (about 20 cm x 10 cm) [71]

Before irradiation both SRAM and Flash were written with a known pattern which was then read out during irradiation. The read-out data was compared with the pattern using a triplicated algorithm running in the FPGA fabric. The triplication protects the read-out algorithm in order to prevent errors in the data transmission to appear as errors in the target system. The results of the comparison were then sent to the user using a RS485/UART data line. For the SRAM the data transmitted to the user contained both the error counter as well as the full data read-out from the eSRAM. In case of the flash memory only the error counters were transmitted to the user. The radiation targeted the whole FPGA but during each test only one component of the FPGA was tested at once. Its not possible to eliminate the possibility that some errors during one measurement steam from components that were not the target of that measurement, but if any this interference should be small.

10.2 Results

Table 10.1 shows the results of the irradiation of the eSRAM of the SmartFusion2. During irradiation, the data was read out from a partition of

Run Nr.	Dose in Gy	Fluence in cm^{-2}	Events	Cross section in $\text{cm}^2/\text{Device}$	Cross section in cm^2/Bit
1	82	$1.409 \cdot 10^{11}$	467	$3.314 \pm 0.165 \cdot 10^{-9}$	$2.529 \pm 0.126 \cdot 10^{-14}$
2	29	$4.983 \cdot 10^{10}$	131	$2.629 \pm 0.131 \cdot 10^{-9}$	$2.006 \pm 0.1 \cdot 10^{-14}$
3	59	$1.014 \cdot 10^{11}$	315	$3.107 \pm 0.155 \cdot 10^{-9}$	$2.371 \pm 0.118 \cdot 10^{-14}$
4	64	$1.1 \cdot 10^{11}$	290	$2.636 \pm 0.132 \cdot 10^{-9}$	$2.012 \pm 0.1 \cdot 10^{-14}$
Comb.	234	$4.021 \cdot 10^{11}$	1203	$2.992 \pm 0.15 \cdot 10^{-9}$	$2.283 \pm 0.114 \cdot 10^{-14}$

Table 10.1: Results of the ESRAM irradiation campaign

the eSRAM with a size of 131040 bit. The per bit cross section of this SRAM of about $2.28 \pm 0.15 \cdot 10^{-14} \text{cm}^2/\text{Bit}$ is a good result comparable to other radiation-tolerant SRAMs (eg. Cypress CY62157EV30 memory with a cross-section of $2 \cdot 10^{-13} \text{cm}^2/\text{Bit}$).

The results from the eNVM flash were very unsatisfying. Already during the first irradiation at a dose of only 60-70 Gy a constant large number of mismatches (1023 of 65024) appeared during the read out of the flash. Until a dose of 120 Gy this number returned to zero for a short time and then climbed to the maximum number (65024) of bits read from the flash. The following runs including several rewrites of the flash content produced similar or worse results.

10.3 Results from other research teams

In parallel to this test of the SmartFusion2 another research team was measuring the SEE cross sections of the main FPGA fabric [72]. This tests included tests on the flip-flops of the logic cells, the fabric SRAM, the arithmetic blocks/multipliers, counters, reprogrammability and the CoreABC bus controller responsible for the Advanced microcontroller Bus Architecture (AMBA) and Advanced Peripheral Bus (APB). Furthermore the current consumption was measured during all tests to search for SEL.

The cross sections found during the test of the fabric flip-flops range between $2.59 \cdot 10^{-14} \text{cm}^2$ and $6.49 \cdot 10^{-15} \text{cm}^2$. A surprising result is that TMR protected flip-flops did not show a significant decrease in cross section with values between $3.79 \cdot 10^{-15} \text{cm}^2$ and $3.51 \cdot 10^{-15} \text{cm}^2$. This issue is currently under further investigation. The whole test suffered from a low total amount of detected events, but shows a clear trend in the behavior of the flip-flops. Furthermore a steep increase in the error rate up to twice the value until then was detected after a total dose of 200 Gy. The device started failing at a dose above 580 Gy.

Test Target	Cross Section in cm^2/Bit
ESRAM	$2.283 \pm 0.15 \cdot 10^{-14}$
Flip-Flops	$6.49 \cdot 10^{-15} - 2.59 \cdot 10^{-14}$
Flip-Flops TMR	$3.79 \cdot 10^{-15} - 3.51 \cdot 10^{-15}$
Fabric SRAM	$2.36 \cdot 10^{-14}$
Multiplier	$1.5 \cdot 10^{-11}$
Multiplier Accumulator	$6.5 \cdot 10^{-12}$
Combination of 32 single byte counters	$2.46 \cdot 10^{-14}$
521 bit counter	$1.7 \cdot 10^{-14}$
512 bit counter TMR	$4.73 \cdot 10^{-16}$

Table 10.2: SEU cross sections of all tested systems

The fabric SRAM showed a very good cross section of $2.36 \cdot 10^{-14} \text{cm}^2/\text{Bit}$ which is very similar to the one found for the ESRAM.

For the arithmetic blocks multipliers and multiplier accumulators were tested. The results were a cross section of about $1.5 \cdot 10^{-11} \text{cm}^2/\text{Bit}$ for multipliers and $6.5 \cdot 10^{-12} \text{cm}^2/\text{Bit}$ for the multiplier accumulators. The failure dose during this measurement was 480 Gy.

For counters both a combination of 32 single-byte counters and one 512 bit counter together with a TMR 512 bit counter were tested. The resulting cross sections were $2.46 \cdot 10^{-14} \text{cm}^2/\text{Bit}$ for the multi counter and $1.7 \cdot 10^{-14} \text{cm}^2/\text{Bit}$ for the 512 bit counter. The TMR-protected 512 bit counters showed a very low cross section with $4.73 \cdot 10^{-16} \text{cm}^2/\text{Bit}$. This is very interesting because this cross section is lower than the TMR-protected flip-flops. The board failed at 610 Gy during this test.

The CoreABC was tested up to a dose of 580 Gy, and no upsets or transmission interruptions were detected.

By irradiating the FPGA with dose steps of 5-10 Gy between reprogramming it the maximal dose for reprogramming was measured. A total of 72 Gy could be withstood before losing the ability to reprogram the FPGA. Table 10.2 shows all cross sections calculated from the different irradiation tests.

10.4 Conclusion

Overall the results from the SmartFusion2/Igloo2 irradiation campaign are very good. The resulting cross sections are comparable to results gained from testing the ProAsic3 [70]. The test of the eNVM in the microcontroller subsystem revealed some serious weakness to radiation. The observed high and

constant error numbers hint at the unresponsiveness of a sector of the eNVM. The SEU cross section of the SRAM is satisfying for radiation-tolerant components. Both eNVM and eSRAM can only be reached using the internal AHB bus of the SmartFusion2 which is, partially due to the toolchain and documentation of Microsemi, quite challenging. The usage of external flash and SRAM might in many cases be the more practical choice. Therefore, the unsatisfying results from the eNVM are not critical for the utilization of the FPGA.

Chapter 11

Conclusion

11.1 Conclusion of the Project

Premature beam dumps due to radiation-induced errors in LHC systems have been one of the important sources of lost integrated luminosity in the last few years. In order to fulfill the project goals of the HL-LHC phase the number of such errors has to be reduced significantly even though the radiation levels will increase, depending on location, by factors between 10-100 times the current levels. Therefore a concentrated effort has been made to increase the radiation tolerance of all relevant systems. The main goal of this project was the increase of the radiation tolerance of a new quench detection system for the 600 A magnets of the LHC. The new system was aimed to produce only two radiation errors every year during the HL-LHC phase. The steps required to achieve this were to verify the behavior of all used electric components under radiation and if necessary to develop mitigation measures against radiation induced faults. Unfortunately, the behavior of electronic components under radiation is neither known to the manufacturer nor possible to be predicted by simple means. Therefore, experiments in which components are artificially irradiated with a high dose while acquiring all possible output parameters, are necessary to determine the radiation tolerance of a component. In the scope of this project, several irradiation campaigns were conducted, testing all components required for the project itself and several others which were parts of other projects or considered useful for future projects. Over the course of these campaigns a lot of experience and techniques were acquired how to conduct efficient irradiation campaigns, which will be useful for future similar endeavors. The results of the irradiation campaign revealed the radiation vulnerability of one of the core components of the new QDS, the ADS1281 ADC. Several failure modes




Figure 11.1: The fault times compared here are from a time interval between the first January and the 15. of August of the years 2015 and 2016. Many groups including the QPS group were able to significantly reduce the failure time their system produced.

were found which would cause a beam dump in most cases and whose combined cross section would make the two errors a year target impossible. In order to understand the origin of these error modes, the behavior of the ADC was simulated using the LabView environment. After understanding the behavior of the failure modes several mitigation measures were developed. The measures were in turn verified during irradiation campaigns and using the LabView simulation and further improved upon. The resulting system is predicted to have an error cross section an order of magnitude lower than the project target leading to only 0.2 errors every year during the HL-LHC phase. In comparison, the previous QDS system was estimated to produce between 70-160 errors every year during HL-LHC so a significant improvement was achieved. The new system was introduced into the machine beginning 2016 and has since then produced not a single error that could be attributed to the effects of radiation. Starting with the year 2016 several other upgraded systems, including the new 600 A QDS have been started operation and the first results are very good. Figures 11.2 and 11.1 show comparisons of stable beam duration and system fault time between the years of 2015 and 2016.

As can be seen the fault times of equipment groups have decreased significantly. In 2016 the fault time produced by the QPS is only about 15% of the fault time in 2015. This reduction in failures has also lead to a signifi-




Figure 11.2: Between the years 2015 and 2016 the average duration of a stable beam changes significantly. This was achieved by reducing the amount of faults that cause a premature beam abort. The average lifetime increased from about 5 hours in 2015 to about 10 hours in 2016.

cant increase in average lifetime of stable beams. The average beam lifetime increased from about 5 hours in 2015 to about 10 hours in 2016. The overall equipment upgrades of the LHC protection systems in general and specifically the upgrade of the 600 A QDS can be called a full success.

11.2 Outlook

While already a high increase in accelerator availability was achieved in the time frame of this work development of this system will not cease. While the approach of using low cost electronic components to create systems with high radiation tolerance is fairly unique to CERN there are several other equipments groups working on similar projects. All these groups are connected within the Radiation to Electronics group at CERN and will gain something from the results produced during this work. One example is the MAX11100 which was qualified in this work and will be used as part of the new main quadrupole QDS that is currently in development. From the test results of the SmartFusion2/Igloo2TM it can be concluded that this device can be used for several future projects. One important candidate will be the development of the QDS for the high luminosity magnets that will be included into the LHC during the change to the HL-LHC phase. The continuous effort to prepare the LHC for the HL phase will keep this topic interesting for the foreseeable future. Depending on the layout of the future FCC accelerator this type of development of system will persist into this era. A parallel equipment tunnel to the accelerator tunnel would allow to place all electronic systems in safe zones without damaging levels of radiation. But such a parallel tunnel would increase the cost of the FCC project considerably. Therefore, there is a high probability that the FCC will have a need for low cost electronic systems with a high radiation tolerance as they are described in this work.

Bibliography

- [1] A. Breskin and R. Voss, *The CERN Large Hadron Collider: Accelerator and Experiments*. Geneva: CERN, 2009. [Online]. Available: <https://cds.cern.ch/record/1244506>
- [2] CMS Collaboration, “Observation of a new boson at a mass of 125 GeV with the CMS experiment at the LHC,” *Physics Letters B, issue 716*, 2012. [Online]. Available: <https://cds.cern.ch/record/1471016/files/plb.716.30.pdf>
- [3] 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, issue 716*, 2012. [Online]. Available: http://ac.els-cdn.com/S037026931200857X/1-s2.0-S037026931200857X-main.pdf?_tid=c0191090-07cd-11e6-be93-00000aab0f01&acdnat=1461249282_5482b7f74d20cfa45c2688d73871ea60
- [4] “LHC Images.” [Online]. Available: https://lhc-machine-outreach.web.cern.ch/lhc-machine-outreach/lhc_in_pictures.htm
- [5] O. S. Bruening, P. Collier, P. Lebrun, S. Myers, R. Ostojic, J. Poole, and P. Proudlock, *LHC Design Report*. Geneva: CERN, 2004. [Online]. Available: <https://cds.cern.ch/record/782076>
- [6] *ATLAS: technical proposal for a general-purpose pp experiment at the Large Hadron Collider at CERN*, ser. LHC Tech. Proposal. Geneva: CERN, 1994. [Online]. Available: <http://cds.cern.ch/record/290968>
- [7] *Technical proposal*, ser. LHC Tech. Proposal. Geneva: CERN, 1994, cover title : CMS, the Compact Muon Solenoid : technical proposal. [Online]. Available: <http://cds.cern.ch/record/290969>
- [8] “Technical Design Report for the Phase-I Upgrade of the ATLAS TDAQ System,” Tech. Rep. CERN-LHCC-2013-018. ATLAS-TDR-023,

- Sep 2013, final version presented to December 2013 LHCC. [Online]. Available: <http://cds.cern.ch/record/1602235>
- [9] A. Apollonio, M. Jonker, R. Schmidt, B. Todd, S. Wagner, D. Wollmann, and M. Zerlauth, “HL-LHC: Integrated Luminosity and Availability,” no. CERN-ACC-2013-0066, p. 4 p, Jul 2013. [Online]. Available: <http://cds.cern.ch/record/1574577>
 - [10] A. Apollonio, M. Brugger, L. Rossi, R. Schmidt, B. Todd, D. Wollmann, and M. Zerlauth, “Roadmap towards High Accelerator Availability for the CERN HL-LHC Era,” pp. Proceedings of IPAC2015, Richmond, VA, USA, 2015. [Online]. Available: <http://cds.cern.ch/record/2141848>
 - [11] M. S. Camillocci, “LHC operation and efficiency in 2015,” 2016, LHC Performance Workshop. [Online]. Available: https://indico.cern.ch/event/448109/contributions/1942062/attachments/1216263/1794003/LHC_operation_and_efficiency_MS.pdf
 - [12] “Accelerator Fault Tracker System.” [Online]. Available: <https://aft.cern.ch>
 - [13] M. Lamont, “How to reach the required availability in the HL-LHC era,” 2014, published by CERN in the Proceedings of RLIUP. [Online]. Available: <https://cds.cern.ch/record/1977366>
 - [14] A. Apollonio, “LHC Availability - Status and Prospects,” 2016, LHC Performance Workshop. [Online]. Available: https://indico.cern.ch/event/448109/contributions/1942032/attachments/1216937/1777608/20160126_availability_Chamonix_AA_final.pdf
 - [15] P. D. W. Buckel and P. D. R. Kleiner, “Supraleitung: Grundlagen und anwendungen 7.auflage,” 2013, ISBN: 978-3-527-66267-8. [Online]. Available: <http://onlinelibrary.wiley.com/book/10.1002/9783527668670;jsessionid=A12756FE355811B9B2F4C31D55FA07F8.f03t04>
 - [16] R. R. Wilson, “The tevatron,” *FERMILAB-TM-0763*, February 1978. [Online]. Available: <http://lss.fnal.gov/archive/test-tm/0000/fermilab-tm-0763.shtml>
 - [17] M. Klein and R. Yoshida, “Collider physics at HERA,” *Progress in Particle and Nuclear Physics*, May 2008. [Online]. Available: <https://arxiv.org/pdf/0805.3334.pdf>

- [18] M. Harrison, T. Ludlam, and S. Ozaki, "RHIC project overview," *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, vol. 499, December 2002. [Online]. Available: <http://www.sciencedirect.com/science/article/pii/S016890020201937X>
- [19] D. Dew-Hughes, "The critical current of superconductors: an historical review," *LOW TEMPERATURE PHYSICS*, *issue 27*, 2001. [Online]. Available: <http://www.w2agz.com/Library/Classic%20Papers%20in%20Superconductivity/DewHughes%20LTP%27%20Jc%20history.pdf>
- [20] M. Allitt, S. Bapna, A. Ijspeert, M. Karmarkar, M. Karppinen, J. Mazet, J. Perez, A. Puntambekar, K. Ruwali, J. Salminen, and A. Thipsay, "Further Development of the Sextupole and Decapole Spool Corrector Magnets for the LHC," *IEEE Trans. Appl. Supercond.*, vol. 10, no. LHC-Project-Report-353. CERN-LHC-Project-Report-353. 1, pp. 162–5. 5 p, Dec 1999. [Online]. Available: <https://cds.cern.ch/record/411154>
- [21] J. Salminen, A. Ijspeert, and A. Puntambekar, "Superconducting Sextupole Corrector Magnet for the LHC Main Dipoles," no. LHC-Project-Report-27. CERN-LHC-Project-Report-27, p. 3 p, June 1996, revised version number 1 submitted on 2004-02-17 09:22:46. [Online]. Available: <https://cds.cern.ch/record/307557>
- [22] "HL-LHC Glossary." [Online]. Available: <http://hilumilhc.web.cern.ch/glossary/d>
- [23] L. Coull, D. Hagedorn, G. Krainz, F. Rodriguez-Mateos, and R. Schmidt, "Electrodynamic behaviour of the LHC superconducting magnet string during a discharge," *5th European Particle Accelerator Conference EPAC 96*, no. LHC-Project-Report-353. CERN-LHC-Project-Report-353. 1, 1996. [Online]. Available: <http://cds.cern.ch/record/314682/files/lhc-project-report-61.pdf>
- [24] Y. Iwasa, "Stability and protection of superconducting magnets-a discussion," *IEEE Transactions on Applied Superconductivity*, vol. 15, no. 2, pp. 1615–1620, June 2005.
- [25] N. Schwerg, K. Dahlerup-Petersen, R. Denz, R. Schmidt, and A. Siemko, "Protection of Superconducting Accelerator Magnets, An Overview."
- [26] K. Dahlerup-Petersen, R. Denz, J. L. Gomez-Costa, D. Hagedorn, P. Proudlock, F. Rodriguez-Mateos, R. Schmidt, and F. Sonnemann,

- “The Protection System for the Superconducting Elements of the Large Hadron Collider at CERN,” no. LHC-Project-Report-283. CERN-LHC-Project-Report-283, p. 4 p, Apr 1999. [Online]. Available: <https://cds.cern.ch/record/386677>
- [27] MP3, “Figure of the Magnet protection System.” [Online]. Available: <https://edms.cern.ch/ui/file/874713/5.1/LHC-MPP-HCP-0001-5-1.pdf>
- [28] F. Rodriguez-Mateos, P. Pagnat, S. Sanfilippo, R. Schmidt, A. Siemko, and F. Sonnemann, “Quench Heater Experiments on the LHC Main Superconducting Magnets,” no. LHC-Project-Report-418. CERN-LHC-Project-Report-418, p. 4 p, Sep 2000. [Online]. Available: <http://cds.cern.ch/record/466520>
- [29] R. Denz and F. Rodriguez-Mateos, “Detection of Resistive Transitions in LHC Superconducting Components,” no. LHC-Project-Report-482. CERN-LHC-Project-Report-482, p. 4 p, Jun 2001. [Online]. Available: <https://cds.cern.ch/record/514344>
- [30] T. May and M. Woods, “Alpha-particle-induced soft errors in dynamic memories,” *IEEE Transactions on Electronic Devices*, vol. 26, no. 1, January 1979.
- [31] J. Ziegler and W. Lanford, “The effect of cosmic rays on computer memories,” *Science*, vol. 206, no. 776, January 1979.
- [32] J. Zielger and H. Puchner, *SER History Trends and Challenges*. Cypress Semiconductor Corporation, 2004.
- [33] R. Baumann, T. Hossain, E. Smith, S. Murata, and H. Kitagawa, “Boron as a primary source of radiation in high density DRAMs,” in *Digest of Technical Papers*. IEEE, 1995.
- [34] S. Michalak, K. Harris, N. Hengartner, B. Takala, , and S. Wender, “Predicting the number of fatal soft errors in Los Alamos national laboratory ASCQ supercomputer,” *IEEE Transactions on Device and Materials Reliability*, vol. 5, no. 3, September 2005.
- [35] C. Povey and G. Hopkinson, “Displacement Damage Mechanism and Effects,” 2009, ESA Presentation. [Online]. Available: http://space.epfl.ch/webdav/site/space/shared/industry_media/05%20DD%20issue3.pdf

- [36] A. G. Holmes-Siedle and L. Adams, *Handbook of Radiation Effects; 2nd ed.* Oxford: Oxford Univ. Press, 2002. [Online]. Available: <https://cds.cern.ch/record/579299>
- [37] R. Naseer, Y. Boulghassoul, J. Draper, S. DasGupta, and A. Witulski, "Critical charge characterization for soft error rate modeling in 90nm sram," in *2007 IEEE International Symposium on Circuits and Systems*, May 2007, pp. 1879–1882.
- [38] M. Huhtinen and F. Faccio, "Computational method to estimate single event upset rates in an accelerator environment," *Nucl. Instrum. Methods Phys. Res., A*, vol. 450, no. 1, pp. 155–72, 2000. [Online]. Available: <https://cds.cern.ch/record/460981>
- [39] D. M. Hiemstra and E. W. Blackmore, "Let spectra of proton energy levels from 50 to 500 mev and their effectiveness for single event effects characterization of microelectronics," *IEEE Transactions on Nuclear Science*, vol. 50, no. 6, pp. 2245–2250, Dec 2003.
- [40] K. Roed, M. Brugger, and C. Pignard, "PTB Irridiation tests of the LHC Radiation Monitor (RadMon). PTB Physikalische Technische Bundesanstalt," Feb 2011. [Online]. Available: <https://cds.cern.ch/record/1329478>
- [41] B. D. Sierawski, K. M. Warren, R. A. Reed, R. A. Weller, M. M. Mendenhall, R. D. Schrimpf, R. C. Baumann, and V. Zhu, "Contribution of low-energy (10 mev) neutrons to upset rate in a 65 nm sram," in *Reliability Physics Symposium (IRPS), 2010 IEEE International*, May 2010, pp. 395–399.
- [42] M. Brugger, R. Alia, S. Danzeca, R. Denz, J. Mekki, P. Oser, P. Peronnard, J. P. De Carvalho Saraiva, G. Spiezia, J. Steckert, and Y. T. S. Uznanski, "Radiation effects, calculation methods and radiation test challenges in accelerator mixed particle and energy environments," *Proceedings of NSREC 2014*. [Online]. Available: <http://www.nsrec.com/short-course.html>
- [43] S. Mukherjee, *Architecture Design for Soft Errors*. San Francisco, CA, USA: Morgan Kaufmann Publishers Inc., 2008.
- [44] R. Baumann, "Soft errors in advanced computer systems," *IEEE Design Test of Computers*, vol. 22, no. 3, pp. 258–266, May 2005.

- [45] J. L. Autran, S. Serre, S. Semikh, D. Munteanu, G. Gasiot, and P. Roche, “Soft-error rate induced by thermal and low energy neutrons in 40 nm srams,” *IEEE Transactions on Nuclear Science*, vol. 59, no. 6, pp. 2658–2665, Dec 2012.
- [46] Picture taken from:. [Online]. Available: http://te-epc-lpc.web.cern.ch/te-epc-lpc/context/radiations/r2e_environment.stm
- [47] X. Ken Chapman, “SEU Mitigation Techniques for SRAM based FPGAs,” Presentation TWEPP, 2015.
- [48] F. Brosser, E. Milh, V. Geijer, and P. Larsson-Edefors, “Assessing scrubbing techniques for xilinx sram-based fpgas in space applications,” in *Field-Programmable Technology (FPT), 2014 International Conference on*, Dec 2014, pp. 296–299.
- [49] PIF homepage. [Online]. Available: <http://pif.web.psi.ch/>
- [50] *Proscan project Documentation*, 2003. [Online]. Available: http://iwbs2004.web.psi.ch/Scientific_Reports/2003/PSI-Scientific_and_Technical_Report_2003-Volume_VI_Dateien/Seite113_114_2003.pdf
- [51] M. Grandjean, “Single Event Effects Test Report A3PE3000L,” *HRX/SEE/0303*, vol. 03, Aug 2003.
- [52] T. Instruments, “High-resolution analog-to-digital converter ads1281,” <http://www.ti.com/lit/ds/symlink/ads1281.pdf>.
- [53] S. Uznanski and Q. King, “H4IRRAD radiation test ADS1281 ADC,” Jun 2012.
- [54] S. Uznanski, B. Todd, A. Dinius, and M. Brugger, “ADS1281: SEL/SEFI test set-up and cross section evaluation,” Oct 2013.
- [55] B. Biskup, M. Brugger, M. Calviani, I. Efthymiopoulos, R. Kwee, J. Mekki, F. La Torre, E. Lebbos, P. Mala, G. Manessi, A. Nordt, F. Pozzi, K. Roed, C. Severino, M. Silari, and A. Thornton, “Commissioning and Operation of the H4IRRAD Mixed-Field Test Area,” Dec 2011. [Online]. Available: <http://cds.cern.ch/record/1403188>
- [56] *MAX1162 datasheet*, 2010. [Online]. Available: <https://datasheets.maximintegrated.com/en/ds/MAX1162.pdf>

- [57] *MAX11100 datasheet*, 2012. [Online]. Available: <https://datasheets.maximintegrated.com/en/ds/MAX11100.pdf>
- [58] “ProAsic3 Product page.” [Online]. Available: <http://www.microsemi.com/products/fpga-soc/radtolerant-fpgas/rt-proasic3>
- [59] “RS-485 specifications.” [Online]. Available: <http://www.ti.com/lit/an/slla070d/slla070d.pdf>
- [60] *MAX6250 datasheet*, 2014. [Online]. Available: <https://datasheets.maximintegrated.com/en/ds/MAX6225-MAX6250.pdf>
- [61] *ADR421 datasheet*, 2013. [Online]. Available: http://www.analog.com/media/en/technical-documentation/data-sheets/ADR420_421_423_425.pdf
- [62] *ADR435 datasheet*, 2014. [Online]. Available: http://www.analog.com/media/en/technical-documentation/data-sheets/ADR430_431_433_434_435.pdf
- [63] *AT25DF041A datasheet*, 2014. [Online]. Available: <https://www.adestotech.com/wp-content/uploads/doc3668.pdf>
- [64] *ADUM3402 datasheet*, 2014. [Online]. Available: http://www.analog.com/media/en/technical-documentation/data-sheets/ADUM3400_3401_3402.pdf
- [65] *INA128 datasheet*, 2014. [Online]. Available: <http://www.ti.com/lit/ds/symlink/ina128.pdf>
- [66] *AQW210EHA datasheet*, 2014. [Online]. Available: <http://www.alliedelec.com/images/products/datasheets/bm/AROMAT/70158338.pdf>
- [67] *AQV251A datasheet*, 2014. [Online]. Available: <http://www.farnell.com/datasheets/10810.pdf>
- [68] *Panasonic Catalog*, 2013. [Online]. Available: https://www.panasonic-electric-works.com/cps/rde/xbcr/pew_eu_en/ca_x61_en_semiconductor_1154.pdf
- [69] *ProAsic3 datasheet*, 2012. [Online]. Available: www.microsemi.com/document-portal/doc_download/130714-radiation-tolerant-proasic3-low-power-space-flight-fpgas-datasheet

- [70] C. Godichal, “Pro Asic3 -Radiation test at CHARM,” 2016. [Online]. Available: https://indico.cern.ch/event/545241/contributions/2213627/attachments/1303482/1948474/Charm_GEFE_results.pdf
- [71] Microsemi. [Online]. Available: <http://www.microsemi.com/products/fpga-soc/design-resources/dev-kits/smartfusion2/smartfusion2-advanced-development-kit>
- [72] S. D. Georgios Tsiligiannis, Rudy Ferraro, “Radiation Test Report â€œSmartfusion2-M2S010 FLASH-based FPGA,” 2016.

Declaration – Erklärung

I declare that I wrote this Doctoral-Thesis independently and without any other references and resources than stated in the bibliography. I have never submitted a Doctoral-Thesis before.

Hiermit erkläre ich, dass ich die vorliegende Doktorarbeit selbstständig verfasst habe und keine als die angegebenen Referenzen und Hilfsmittel verwendet habe. Ich habe nie zuvor die Promotion beantragt.

Darmstadt, Februar 2017

Lebenslauf

- Geboren : 14.07.1983 in Frankfurt am Main
- 1990-1994 Besuch der Waldschule in Maintal Bischofsheim
- 1994-2003 Besuch des Albert-Einstein Gymnasiums
- 2003-2017 Physik Studium TU-Darmstadt mit spezialisierung auf Beschleunigerphysik. Erlangung des Rangs eines Master of Science Physik.
- 2013-2016 Promotion durchgeführt am CERN mit dem Titel "‘Development of Radiation-tolerant Components for the Quench Detection System at the CERN Large Hadron Collider"’

Danksagungen

At the end I would now like to thank all the people who helped me create this Doctoral-Thesis.

I want to express my gratitude to my supervisor at CERN Dr. Reiner Denz and to my colleagues Jens Steckert and Dr. Rüdiger Schmidt. Their continuous assistance and helpful advise made this work possible.

Furthermore I want to thank Prof. Joachim Enders who supervised the writing of this thesis. His many suggestions helped tremendously to improve the quality of this work.

Lastly I want to thank the german "Bundesministerium für Bildung und Forschung" and the Gentner stipend that enabled my stay at CERN to write this thesis.