



Advancing LHC Superconducting Circuit Models through Closed-Loop Multi-Objective Optimization

In collaboration with: European Organization for Nuclear Research (CERN)

Submitted by

Max GORENFLO

Matr.-Nr.: 3574737

From Faculty XIII - Physics
of Philipps University Marburg
for the attainment of the academic degree

Master of Science

- M.Sc. -

First Reviewer and Advisor: Dr. Emmanuele RAVAIOLI

Second Reviewer: Prof. Dr. Jan Christoph GOLDSCHMIDT

July 2024, Geneva



Erklärung

Hiermit erkläre ich, dass ich meine Masterarbeit mit dem Thema:

Advancing LHC Superconducting Circuit Models through Closed-Loop Multi-Objective
Optimization

selbständig verfasst sowie alle wesentlichen Quellen und Hilfsmittel angegeben habe.

Name, Vorname: Gorenflo, Max

.....

.....

Ort, Datum

Unterschrift

Advancing LHC Superconducting Circuit Models through Closed-Loop Multi-Objective Optimization

Abstract: The **L**arge **H**adron **C**ollider (LHC) comprises eight individual main dipole circuits (RB), each containing 154 superconducting magnets. Simulations of the electrothermal behavior of these circuits during transients are an essential tool to verify the circuit performance and for the analysis of unexpected events occurring in these circuits. For example they are used, to analyze **F**ast **P**ower **A**bort (FPA) events, which are triggered in case a disturbance (such as a power-supply trip, equipment malfunctioning, or a magnet quench) in the main dipole circuit is detected. Since the physical mechanisms determining the behavior of superconducting magnet circuits span different scales, it is not practical to exclusively simulate them using conventional approaches like finite element methods, as the simulation would become computationally too expensive. However, these challenging simulations can be carried out with the **S**imulation of **T**ransient **E**ffects in **A**ccelerator **M**agnets (STEAM) framework, developed at CERN. The accuracy of these simulations is highly influenced by the models of two different types of circuit subcomponents: the energy extraction resistors and the magnets' by-pass diodes. The electrical behavior of these two components is temperature-dependent which was not taken into account in their previous STEAM models. In some cases, this led to significant deviations between simulated and measured circuit signals.

The aim of this thesis was therefore to model the electrothermal behavior of these components and to implement and use a new method to automatically fit the models' free parameters to measurement data. For this purpose an already existing software interface between STEAM and the parameter optimization toolkit Dakota (developed at Sandia Labs) was extended for conducting closed-loop multi-objective optimizations. In this approach, simulation results are iteratively compared with measurement data in a closed-loop fashion. Since several signals are optimized simultaneously, the method applied in this thesis is referred to as closed-loop multi-objective optimization. The application of this method made it possible to fit the free parameters of the component models, whose influence on simulation accuracy is complex and non-linear.

By following this approach, an improvement in simulation accuracy of the energy extraction resistance of up to 14 % could be achieved (errors are normalized to the maximum absolute value of the measured quantity). Thus, the accuracy of the simulated circuit current, a crucial quantity of the superconducting magnet circuit, improved by up to 2 %. The newly implemented diode model allowed to qualitatively explain the occurrence of post-quench recooling-related peaks in the forward voltage of the diodes which was not possible before. These were observed to be accompanied by a sharp drop in the diodes' temperature and a highly increased resistance of the diodes.

The models are based on a more complete set of physical effects, and are thus expected to be more robust for simulating unexpected events for which experimental data are currently not available. The good transferability of the applied method to other models ensures that closed-loop optimization will be a promising option for further improving circuit simulation accuracy in future works. In the future, the STEAM framework will thus be more accurate and reliable for research and development related tasks like hardware commissioning, the analysis of FPA events and the further development of the LHC superconducting magnet circuits.

Acknowledgements

First, I would like to thank Emmanuele Ravaioli, for the excellent supervision of my work. Furthermore, I would like to thank Prof. Dr. Jan Christoph Goldschmidt, who supervised my thesis as a second reviewer. I would like to thank Arjan Verweij for his valuable inputs on the topics of my thesis. I am particularly thankful to the members of the STEAM and SIGMON teams for their continuous support and helpfulness with questions of all kinds, which contributed significantly to the success of my thesis. Furthermore, I am grateful for the help of Dieter Hagedorn and Daniel Wollmann who shared their knowledge about diodes with me. Additionally, I want to thank Bozhidar Ivanov Panev, for providing me with valuable insights about the principles of operation of the circuits' energy extraction system.

Contents

List of symbols	iii
List of abbreviations	v
1 Introduction	1
1.1 Large Hadron Collider and High Luminosity LHC project	1
1.2 Superconducting magnet circuits	2
1.2.1 Foundations of superconductivity	2
1.2.2 Superconducting cables and magnets	5
1.2.3 Losses in superconducting cables	7
1.2.4 Quench	8
1.3 Quench protections systems	8
1.4 LHC main dipole circuit and FPA events	11
1.5 Simulation of superconducting magnet circuits	13
2 Methods and tools for optimizing superconducting magnet circuit models	15
2.1 STEAM	15
2.1.1 Objective of STEAM	15
2.1.2 Structure of the STEAM-SDK	15
2.1.3 File conventions	16
2.2 Automated simulation and validation of magnet circuits	17
2.2.1 Netlist simulations	17
2.2.2 Analog behavior modelling	17
2.2.3 Structure and sequence of the analysis	18
2.2.4 Acquisition of measurement data	19
2.3 Closed-loop multi-objective optimization	21
2.3.1 STEAM-Dakota interface	21
2.3.2 Algorithms	22
3 Model of the LHC main dipole circuit's energy extraction resistors	25
3.1 Model	25
3.1.1 Thermodynamic equivalence model	25
3.1.2 Temperature-dependent material properties	29
3.1.3 Full equivalence model of an energy extraction unit	30
3.2 Setup of the optimization	31
3.2.1 Selection of suitable events	31
3.2.2 Selection of a suitable optimization algorithm	33
3.3 Results	34
3.3.1 Old energy extraction model (OM)	35
3.3.2 New energy extraction model (NM)	35
3.3.3 New energy extraction model after fit (NMF)	41

3.3.4	New energy extraction model after fit considering variable inductance (NMFI)	41
3.3.5	Meta analysis of the parameter fits	50
3.4	Summary	52
4	Model of the LHC main dipole magnet's by-pass diodes	57
4.1	Model	57
4.1.1	Modelling of the diodes' temperature-dependent I-V characteristic	59
4.1.2	Thermodynamic modelling of the diodes	65
4.2	Results	68
4.3	Summary	74
5	Conclusion and outlook	77
A	Appendix	81

List of symbols

Sign	Description	Unit
C_v	Heat capacity at constant volume	W K^{-1}
ΔT_{calc}	Temperature increase which would be necessary to explain the measured increase in resistance (assuming homogeneous resistivity, heat capacity and temperature across the resistor)	K
ΔT	Increase in temperature	K
I_0	Circuit current at the time the fast power abort is triggered	A
I_D	Current through the diode	A
I_{MEAS}	Circuit current	A
I_S	Reverse-bias saturation current	A
I_{res}	Current through the energy extraction resistor	A
P_{Cv}	Power absorbed or released by the systems's heat capacity	W
P_{exchange}	Power associated with heat exchange	W
P_{ohmic}	Ohmic heating power introduced by a current flow through the system	W
R_{EE}	Resistance of the energy extraction unit	Ω
T_D	Temperature of the diode's wafer	K
T_{HS}	Temperature of the diode's heat sink	K
T_{He}	Temperature of superfluid helium	K
T_{amb}	Ambient temperature	K
T_{res}	Temperature of the energy extraction resistor	K
T	Temperature	K
V_{EE}	Voltage across the energy extraction	V
V_D	Voltage across the diode	V
V_T	Thermal voltage of the diode	V
c_v	Volumetric heat capacity at constant volume	$\text{W K}^{-1} \text{m}^3$
f_{Cv_D}	Scaling factor accounting for imperfections in the assumed heat capacity of the diode	-

Sign	Description	Unit
$f_{\text{Cv_HS}}$	Scaling factor accounting for imperfections in the assumed heat capacity of the diode's heat sinks	-
$f_{\text{Cv_EE}}$	Scaling factor accounting for imperfections in the assumed heat capacity of the energy extraction	-
f_{res}	Scaling factor of the initial resistance of the energy extraction resistor	-
$h_{\text{th_D-HS}}$	Thermal conductance between the diode and its heat sinks	W K^{-1}
$h_{\text{th_HS-He}}$	Thermal conductance between diode's heat sinks and the helium	W K^{-1}
$h_{\text{th_EE}}$	Thermal conductance to model the energy extraction's cooling	W K^{-1}
k_{B}	Boltzmann constant	J K^{-1}
n	Ideality factor of the diode	-
q	Elementary charge of an electron	C
ρ	Resistivity	$\Omega \text{ m}$

List of abbreviations

Acronym	Description
ABM	Analog B ehavior M odel
CERN	European Organization for Nuclear Research
CLIQ	C oupling L oss I nduced Q uench
CLO	C losed- L oop O ptimization
co-simulation	C ooperative s imulation
EE	E nergy E xtraction
FEM	F inite E lement M ethod
FPA	F ast P ower A bort
HL-LHC	H igh L uminosity L arge H adron C ollider
HTS	H igh- T emperature S uperconductors
LHC	L arge H adron C ollider
LHC-SM-API	LHC-S ignal- M onitoring- API
LS1	L ong S hutdown 1
LTS	L ow- T emperature S uperconductors
MQE	M inimum Q uenching E nergy
NM	N ew energy extraction M odel
NMF	N ew energy extraction M odel after F it
NMFI	N ew energy extraction M odel after F it considering variable I nductance
NXCALS	N e X t C ERN A ccelerator L ogging S ervice
OM	O ld energy extraction M odel
PC	P ower C onverter
PM	P ost M ortem
QDS	Q uench D etection S ystem
QH	Q uench H eater
STEAM	S imulation of T ransient E ffects in A ccelerator M agnets
STEAM-DI	STEAM-D akota I nterface
STEAM-LEDET	STEAM-L umped E lement D ynamic E lectro T hermal

Acronym	Description
STEAM-SDK	STEAM -Software D evelopment K it

Introduction

1.1 Large Hadron Collider and High Luminosity LHC project

With a total length of about 27 km, the **Large Hadron Collider (LHC)** developed at the **European Organization for Nuclear Research (CERN)** is the largest and most powerful particle accelerator in the world. Through the collision of protons or heavy ions, the **LHC** generates conditions that are similar to those just seconds after the big bang [1, 2]. These enormously high energy densities achieved are ideal to study the generation, decay, and interaction processes of different particles in order to gain a better understanding of the elementary building blocks and mechanisms, shaping our universe. Inside the **LHC**'s circular tunnel, two particle beams are accelerated in opposite directions and crossed at various points, where the beams then collide. Each collision point is equipped with a different experiment, measuring the collision products using particle detectors. The elementary particles produced in the collisions, especially the more massive ones such as the Higgs boson, have a very short lifetime and are mostly not directly measurable in the detectors of the individual experiments. As these particles decay, however, they leave behind a characteristic trail of decay particles, most of which can be measured in the particle detectors at **CERN**. By analyzing this "decay fingerprint", it is possible to determine a probability that a particular particle was formed within a collision for each event in the **LHC**. However, the probability of formation of heavy elementary particles is often directly linked to the energy of the collisions. There are, for example, particles whose probability of formation at very low energies is so vanishingly small that they could either not be observed at all or could not be statistically significantly distinguished from the noise of the measurement uncertainties, if the accelerator did not provide sufficient collision energy. This is why the research progress in particle physics is closely linked to the energy reached during the collisions in the particle accelerator [3]. Hence, as technological developments progressed and theoretical physics predicted the existence of new particles, accelerators enabling higher and higher collision energies were built at **CERN**. During these upgrades, some accelerators were connected, serving as pre-accelerators for the **LHC** as part of an acceleration chain, which is shown in figure 1.1a. The circular accelerator collision energy achievable is influenced by the magnetic flux density within the **LHC**'s main dipole circuits, which is a strong motivation for their ongoing development.

Another characteristic quantity of particle accelerators is their luminosity. In this context, luminosity is a measure proportional to the number of collisions that occur in a certain time interval in an accelerator [4]. The next bigger upgrade to the **LHC** is called **High Luminosity Large Hadron Collider (HL-LHC)**. As shown in the timeline in figure 1.1c, over the course of this upgrade the luminosity is to be increased from currently twice the nominal luminosity (nominal luminosity corresponds to the initial design goals) to five times the nominal luminosity. Among other things, this will make it possible to investigate certain interactions such as e.g., rare decays of the Higgs boson even better in the future [5]. Further, these upgrades also include changes to several superconducting magnets in the **LHC**. As part of this upgrade, the quadrupole magnets of the inner triplet are upgraded, which is associated with design changes

as exemplified in figure 1.1b. With the HL-LHC upgrade, higher magnetic flux densities will ensure better focussing of the beam in the future [6] and thus lead to a higher luminosity of the accelerator. These two examples, collision energy and luminosity, illustrate the strong link between the flux densities achievable in accelerator magnets and the experimental feasibility of certain high-energy physics experiments. With this motivation for stronger magnets now in mind, the following section 1.2 will take a closer look at the special properties of the magnets usually used for accelerators with very high collision energies.

1.2 Superconducting magnet circuits

In the course of this section, superconducting magnets and their circuitry will be discussed. As already motivated in the previous section 1.1, sufficiently strong magnets are of great importance for the construction of powerful particle accelerators. Superconductors exhibit zero resistance, which makes them the predominant choice for this application, where otherwise large losses would occur due to the high operating currents. In the following course, section 1.2.1 will first briefly discuss the physical principles of superconductivity. This is followed by section 1.2.2, which highlights the advantages of superconducting magnets and motivates the most common design choices made for their technological utilization. Finally, in the sections 1.2.3 and 1.2.4, the special challenges in the construction of superconducting magnets are discussed.

1.2.1 Foundations of superconductivity

As introduced in the book by Mangin and Kahn, the research on superconductivity has a long history [9]. Superconductivity is a state in which a material loses its electrical resistance. This was first discovered by Onnes, who found that the electrical resistance of distilled mercury drops to exactly zero at very low temperatures [10, 11]. In this state, which is reached when the temperature T falls below the so-called critical temperature T_c , electrons can move unhindered through the material, evading the interaction with lattice vibrations (phonons), which is the cause of electrical resistance in normal conducting materials. In the further course of research on superconductivity, the Meissner-Ochsenfeld effect in 1933 was another significant discovery. This effect described the behavior of superconductors in a magnetic field for the first time, as the magnetic field is displaced from inside the conductor. What this effect signifies for the physical properties of the material can be illustrated considering the flux density equation of an isotropic material

$$B = \mu_0 * (H + M) \tag{1.1}$$

with the magnetic flux density B and the magnetic field strength H , which are linked through the magnetization M and the vacuum magnetic permeability μ_0 . Since the magnetic flux vanishes inside the superconductor, the magnetization M must be of the same magnitude, but opposite sign compared to the magnetic field strength H . As they are correlated with the magnetic susceptibility χ_m by $M = \chi_m * H$, the magnetic susceptibility inside the conductor is given by $\chi_m = -1$. The material therefore exhibits perfect diamagnetic properties [12, 10, 13]. Shortly after the discovery of superconductivity, the search was on for ways to utilize it technically. Experiments with superconducting coils revealed that the temperature is not the only decisive factor on whether a material is superconducting or not. During experiments, which

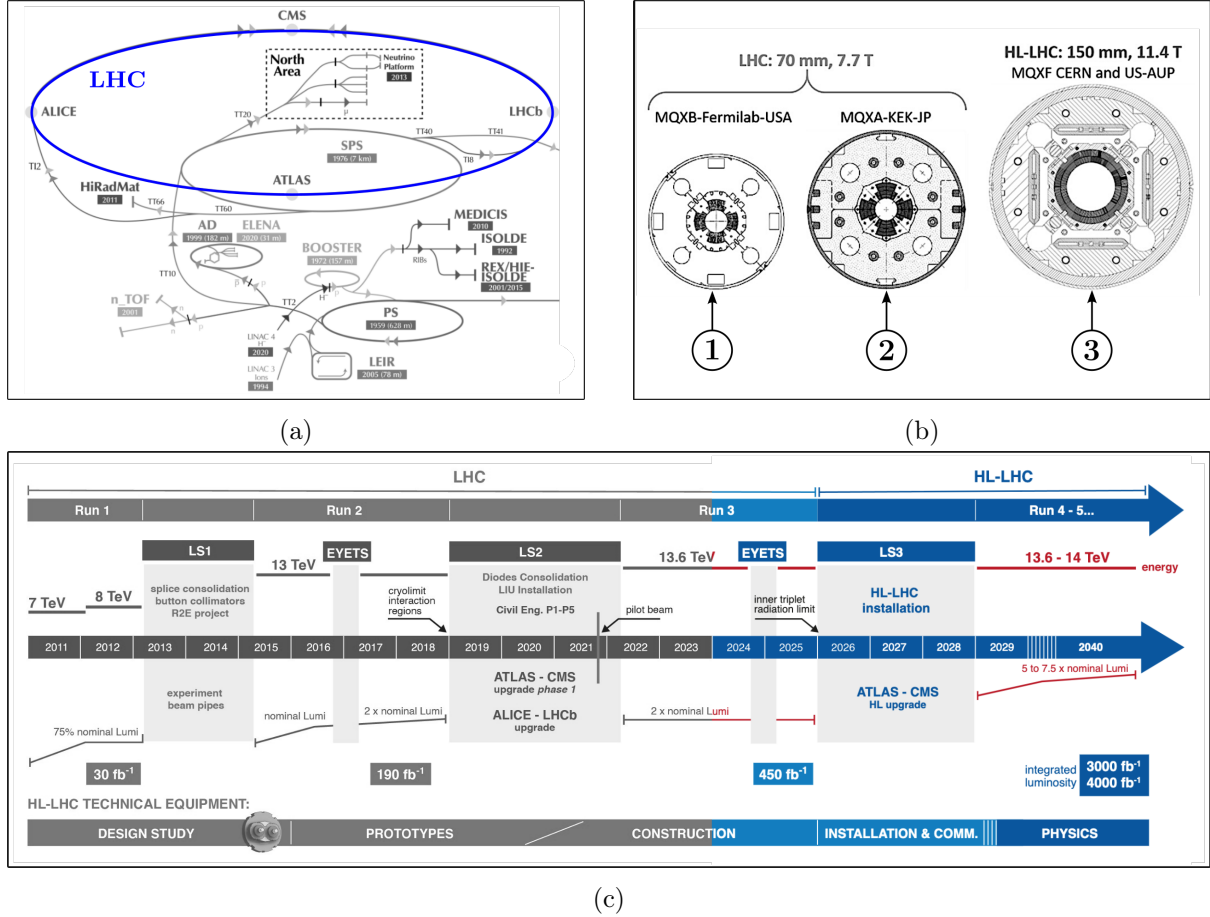


Figure 1.1: Particle accelerators at CERN and their upgrades over time. a) Accelerator complex at CERN (modified image - original from the CERN website [7]). b) Comparison of the inner triplet quadrupoles in the LHC and in the HL-LHC: (1) MQXB magnet as it was installed in the LHC, (2) MQXA magnet as it was installed in LHC, (3) MQXF magnet as it will be installed for the HL-LHC project (modified image - original image from the work of Rossi and Brüning [6]). c) Overview of the various upgrades at the LHC since its initial commissioning (modified image - original from the HL-LHC website [8]). The colored area of the timeline lies in the future, whereas all steps in grey lie in the past.

aimed for investigating the possibility for the construction of very strong, iron-free magnets, it was discovered that the phase transition of superconductors also depends on the strength of the external magnetic field applied to the superconductor [14, 15]. In this context Abrikosov provided the decisive theoretical analysis [9, 16], showing that a distinction is to be made on how conductors lose their superconducting state, when being exposed to a magnetic field. In type I superconductors, the superconducting phase disappears abruptly after a critical field strength H_c is exceeded. However, this abrupt drop does not occur with superconductors of the second type. Instead, a gradual phase transition can be observed, which begins when a low critical field strength H_{c1} is exceeded. The complete disappearance of the superconducting phase in type II superconductors is only completed after a significantly higher critical field strength of H_{c2} is exceeded. This is illustrated in figure 1.2a, in which lead, exhibiting type I behavior, is compared with lead-indium alloys, which exhibit type II behavior. Generally only superconductors made out of pure elements show type I behavior [9]. The critical field strength and the critical temperature are closely linked. Their relationship can sometimes be approximated as parabolic [15, 17, 13, 18]. However, this can only apply to special cases, as the phase transitions usually require more complex descriptions. In 1916, Silsbee showed that the current density influences the superconducting phase in a way that is analogous to the magnetic field strength's influence [15]. There is also a critical current density j_c which, if exceeded, is accompanied by a phase transition. For type I superconductors in particular, there is a close relationship between the critical field and the critical current density. In type II superconductors, on the other hand, the critical current density is determined more by the metallurgical material properties [9]. In summary, these three physical quantities namely the current density, the external magnetic field and the temperature influence the superconducting state and span the so-called critical surface which is shown in figure 1.2b. The critical surface can be interpreted as a boundary in phase space that separates the state vectors leading to a superconducting phase from those associated with a normal conducting phase. If the state vector is within this critical surface, the material is superconducting. Outside the critical surface, superconductivity is lost. Accurately describing the physics behind superconductivity posed an enormous challenge at the beginning of the twentieth century, as the underlying quantum mechanical effects were not yet sufficiently well understood. The brothers Fritz and Heinz London developed a phenomenological theory, London theory, which was able to describe the Meissner-Ochsenfeld effect for the first time [9]. With this theory, for the first time a characteristic length was defined that could describe the depth of penetration of the magnetic field into the conductor [9]. Several other theories were developed after the London Theory. However, it was not until the 1950s that a quantum theory of superconductivity was found with the BCS theory [9, 13]. Even before the development of the BCS theory, there were indications that superconductivity was not solely determined by the interactions between the electrons, but also by the properties of the lattice atoms, as suggested by the isotope effect. The isotope effect states that the critical temperature of a material scales linearly with the reciprocal of the square root of the material's atomic mass [9]. This is exactly where the BCS theory comes in, which explains superconductivity based on the interaction of electrons with lattice movements. When an electron crosses the lattice, it attracts lattice atoms, which then move from their initial position in the direction of the electron. This displacement of the atoms from their initial position results in local accumulations of positive charge, which can in turn have an attractive effect on a second electron. The two electrons involved, the first that distorted the lattice and the second that was attracted to it, therefore interact with each other

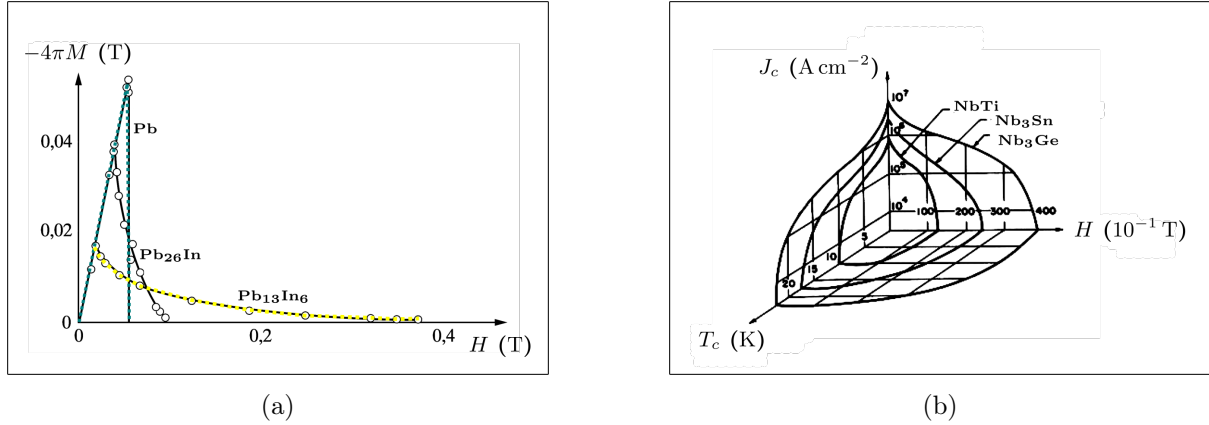


Figure 1.2: Properties of superconducting materials. a) Measurements of magnetization M as a function of magnetic field strength H for type I and type II superconductors (modified image - original from Gerthsen [17], data from Livingston [19]). While lead in its elemental form (green graph) exhibits type I behavior, its alloys with indium (black and yellow graphs) exhibit type II behavior. b) Critical surface of Nb-Ti, Nb₃Sn and Nb₃Ge superconductors (modified image - original from Beasley [20]).

via lattice oscillations (phonons). This pairwise interaction of the electrons via the exchange of virtual lattice vibrations (phonons) is called a Cooper pair [21]. The formation of Cooper pairs is associated with a reduction in the total energy of the electrons, which opens up the so-called superconducting gap. Put simplistically, this is comparable to the band gap in semiconductors: since there are no electronic states feasible within the gap, it becomes impossible for the electrons to change their energy through collisions with the lattice. Their energy is therefore preserved when moving through the lattice, resulting in zero electrical resistance.

1.2.2 Superconducting cables and magnets

Superconducting cables and magnets have a long tradition in the construction of particle accelerators [22, 23]. In order to achieve higher energies in the accelerators, their magnets must provide higher magnetic flux densities for keeping the charged particles on their path, which is most easily achieved with higher coil current densities. However, comparing the quadratic increase of ohmic heat losses with the only linear increase of a magnets flux density, for an increasing coil current, the great advantage of a disappearing ohmic resistance in superconducting coils becomes apparent. Another problem with the simple upscaling of conventional magnets for stronger fields is saturation effects of the magnet's iron yoke. Normal conducting magnets are often equipped with a ferromagnetic core or yoke, which (as can be seen from equation 1.1) can amplify the magnetic field. However, as the ferromagnetic material enters a saturation at high magnetic field strengths, this contribution is less effective when a magnetic field flux of about 2 T or higher is to be obtained [24]. Thus, leveraging the disappearing resistance of superconducting coils for the design of high field accelerator magnets offers decisive advantages. However, in order to fully exploit this potential, a wide range of design objectives must be taken into account, including the selection of a suitable material.

Thanks to intensive research, several superconducting materials are known today. Over the




Figure 1.3: Types of superconducting cables. a) Rutherford cables in various designs: (1) Conductor with about 3000 Nb-Ti filaments, in a copper matrix. (2) Nb₃Sn Conductor with 192 filaments (modified image - original from Larbalestier et al. [26]). b) Cross-section of a wire in channel conductor (modified image - original version from Guo et al. [27]).

course of time, new records have been set for higher critical temperatures, which soon even exceeded the boiling point of liquid nitrogen [18]. In the proceeding of these advances, a distinction was made in the literature between **Low-Temperature Superconductors (LTS)** which have a critical temperature T_c below 20 K and **High-Temperature Superconductors (HTS)**, featuring a critical temperature T_c above 20 K. However, a major drawback of HTS is their brittleness, complicating the production of HTS cables on an industrial scale [18, 15]. In the design phase of the LHC, the alternatives of Nb-Ti at 1.9 K and Nb₃Sn at 4.2 K were discussed, but it was ultimately the uncertainty about the industrial maturity of Nb₃Sn at this time that led to the choice of Nb-Ti [23]. In order to ensure that the normal conducting zone spreads quickly and evenly, the individual superconducting filaments are usually embedded in a thermally and electrically conductive matrix material such as copper or aluminum [25]. In most superconducting cable designs, two distinct regions of the cable are left without filament: one is located in the center of the cable, the other one is ring-shaped at the edge of the cable as can be seen in figure 1.3a. The space in the middle of the cable is intended to allow filament transposition, whereas the free space at the edge of the cable is left out for manufacturing reasons [25]. However, the shapes and designs of the cables are diverse, and offer various advantages and disadvantages. The so-called wire-in-channel cable, for example, is embedded in an additional bulk of matrix material as shown in figure 1.3b. This serves as an additional heat sink in case the superconducting filaments turn back to the normal conducting state (this is called a quench and discussed in more detail in section 1.2.4) [25] but it leads to a higher cross-section, which is a disadvantage for the construction of very compact magnets. Depending on the arrangement of the strands in a cable, a distinction can be made between different cable types. What most designs have in common however, is that the individual strands are fully transposed along the length of the cable. This ensures a more even current density distribution over the cross-section of the cable, and thus reduces the AC losses of the cable [28]. Finally, the most common cable designs are briefly presented below:

- Rutherford cables: The Rutherford cable is the most common design in accelerator appli-

cations. This is due to the high current densities and packing factors achievable as well as its high mechanical stability [28, 22].

- Rope type cables: In this design, the strands are twisted with each other. The resulting thicker conductor strands are then partially twisted again, resulting in a solid strand, like a rope. This cable design leaves cavities inside the conductor which can be used for cooling the cable [28]. Hence, one usual application of this design is cable-in-conduit conductors [28, 29], in which the cable is inside a conduit and in direct contact with its coolant.
- Braid type cables: As the name suggests, the individual strands in this design are braided with their respective nearest neighbor strand. It is therefore very similar to the design of the rutherford cable, but in comparison with poorer mechanical stability [28].
- Nuclotron cables: The design of nuclotron cables is comparable to that of rutherford cables. The nuclotron cable however, features a thin channel that runs through the center of the cable, which is used for cooling [28, 30]:
- Roebel cables: The design of the Roebel cable is realized for flat tape cables, as usually used with high-temperature superconductors. Additionally, just like in other cables, pairs of flat layers of tape are transposed with each other [28, 31].

1.2.3 Losses in superconducting cables

In most technical applications, superconducting cables are exposed to changing magnetic fields. This applies to the LHC as well, where the fields change during the machine ramp up and ramp down, resulting in transient losses in the cables. As these losses influence the performance and efficiency of each superconducting magnet, and can even influence the occurrence of a quench, the most important losses occurring in superconducting cables are briefly presented here:

- Inter-filament coupling losses: When a current flows through the resistive matrix material from one filament to the other (but within the same strand), the ohmic loss associated with this is referred to as an inter-filament coupling loss [32, 33].
- Inter-strand coupling losses: If current flows from one strand of the cable to the other, this is referred to as inter-strand coupling losses. Here, the current flows through a contact resistance, which results in ohmic losses [32].
- Hysteresis losses: If the superconducting material undergoes a change in its internal magnetic field due to the movement of magnetic vortices, this is associated with so-called hysteresis losses [32, 34].
- Eddy current losses: If an eddy current is induced in metallic elements close to the cable, the corresponding energy loss is referred to as an eddy current loss [35].
- Ferromagnetic losses: Every ferromagnetic material is subject to hysteresis when the external magnetic field changes. If the cable has ferromagnetic constituents, the corresponding losses are referred to as ferromagnetic losses [36].
- Mechanical losses: If there is a movement of a cable due to Lorentz forces acting on it, this transfer from energy stored in the field, to kinetic and finally thermal energy is associated with losses [35].

1.2.4 Quench

A superconductor quenches when the phase vector describing the conductors state is leaving the region below the critical surface in the (j, B, T) phase space [31, 37]. Since the specific heat capacity of most materials decreases with falling temperatures, very small energy dispositions in the conductor, caused e.g. by transient losses, can heat the superconductor enough to cause a quench. In this context, the amount of energy needed to quench a superconductor is referred to as the **Minimum Quenching Energy (MQE)** [24]. The energy stored in the superconducting magnets, on the other hand, can be very high, which, together with the fragility of the superconducting phase, bears considerable potential for damage [24]. When an initially small disturbance causes a localized quench in a filament, the change in resistance of this filament causes a redistribution of the currents within the cable. This, drives a part of the current flow through the matrix material, which in this context is referred to as the current sharing regime [31, 38]. Due to the current flow in the resistive matrix material, further energy is dissipated into the conductor, leading to a localized temperature increase, the so-called quench hot-spot [39]. The temperature of this hot-spot is a relevant parameter in the design process of superconducting magnets.

After the initial energy perturbation, the further temperature characteristic of a superconducting cable can develop in two fundamentally different ways: thermal runaway or recovery [31].

In the case of thermal runaway, there is a further propagation of the normal conducting zone. Due to the continuous propagation of the normal conducting zone (which is primarily advancing in the transport current flow direction [25]) and thus the further increasing energy disposition in the magnet, the magnet is increasingly heated [31]. This uneven propagation of the normal conducting zone can lead to damages, as the formation of large temperature gradients can cause high mechanical stresses in the magnet. The speed at which the normal conducting zone propagates is therefore an important factor in the design of magnets, as high propagation speeds reduce these gradients. In the literature, this velocity is called normal-zone propagation velocity, which can be in the range of 5 - 20 ms⁻¹ [24].

If the thermal runaway described above does not occur, the superconductor may recover instead. In this case, the local cooling of the superconductor is strong enough to compensate the energy disposition of the initial energy perturbation and can thus restore the conductor to the superconducting state [31].

1.3 Quench protections systems

General Motivation

As explained in the previous section 1.2.4, quenching is associated with the transfer of magnetic energy into heat. In superconducting magnet circuits, which are often characterized by high inductances and operating currents, this energy can be very high. For instance, the potential magnetic energy stored in the 154 bending magnets of the main dipole circuit, operating at approximately 11.8 kA can be approximated as follows [24, 38]:

$$E_m = \frac{154 \cdot 0.1 \text{ H} \cdot (11.8 \text{ kA})^2}{2} \approx 1.1 \text{ GJ.} \quad (1.2)$$

If this amount of energy were to be dissipated locally in the magnet coils, this could lead to damage, as already explained in section 1.2.4. The aim of quench protection systems is hence to quickly dissipate the magnet energy while preventing the development of localized zones of largely increased temperature. With the objective of a successful quench protection strategy now introduced, the next section will take a closer look at their technological implementation.

Types of quench protection systems

As introduced in the work of Ravaioli [25], quench protection systems can be divided by two different criteria. A distinction is made between systems with internal and external energy dissipation and, depending on the type of mechanism, between active and passive systems. This is exemplified in figure 1.4. In this context, active refers to the existence of electronic

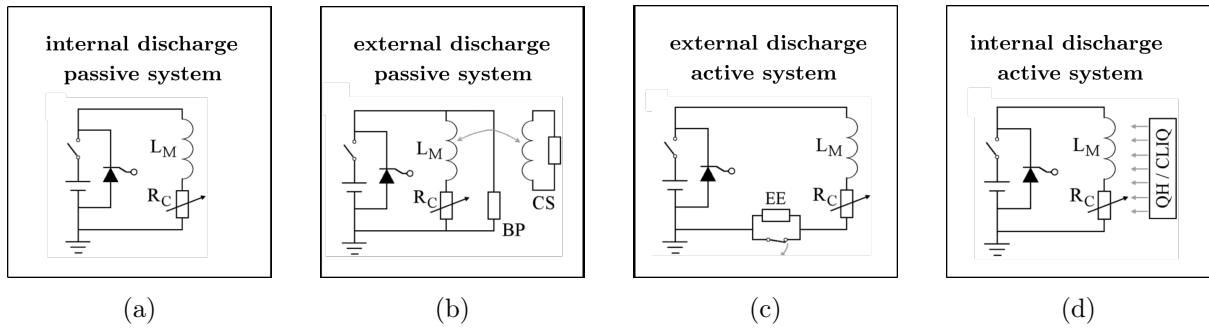


Figure 1.4: Different categories of quench protection systems. a) Passive systems with internal discharge. b) Passive systems with external discharge, c) Active systems with external discharge, d) Active systems with internal discharge (modified image - original from the work of Ravaioli [25])

components that initiate active measures (like activating a switch) in order to protect the circuit after detecting a quench event with a **Quench Detection System (QDS)** [25]. In passive systems, on the other hand, such active components are not implemented. Furthermore, internal energy dissipation refers to the main part of the magnetic energy being dissipated in the magnet coil itself, whereas external energy dissipation refers to the main part of the magnetic energy being dissipated in a component outside the magnet.

Self-Protected magnets

As it does not feature any additional components as shown in figure 1.4a, the simplest case of a quench protection is the internal passive system, which is called self-protected in the literature [40, 25, 39]. This variant is a very cost-effective and, due to its passivity, reliable protection for superconducting magnets [25]. In this system, a large part of the energy disposition takes place directly inside the magnet. However, this approach can only be applied, if the quick formation of a uniformly distributed normal conducting zone can be ensured in case of a quench. Otherwise, localized hot-spots could lead to damage of the magnet. However, even though this type of quench protection has the aforementioned advantages of reliability and ease of implementation, it is unsuitable for high energy-density magnets with high field strengths as the possible current densities with this design are comparatively limited [25].

Systems with by-pass elements

Another variant are passive safety systems with by-pass elements as shown in figure 1.4b. The idea behind this is to offer the transport current an alternative branch in case of a quench in order to by-pass the actual conductor and thus reduce ohmic loss in the magnet coil where the quench started. In most cases, these components are diodes or resistors [25, 41]. However, a design compromise must always be found when using resistors as by-pass elements. On the one hand, the resistor must be small enough to provide an effective alternative path to the current in case of a quench, so that a large proportion of the current flows through the by-pass resistor [25, 41]. On the other hand, low by-pass resistances are associated with current heat losses during normal operation of the magnet. This is due to leakage currents which may occur during current changes in the magnet as described by the following formula [41]:

$$I_{\text{leak}} = \frac{dI(t)}{dt} \cdot \frac{L_{\text{magnet}}}{R_{\text{parallel}}}.$$

As shown by this equation, the leakage current I_{leak} , is dependant on the current through the magnet $I(t)$, the inductance of the magnet L_{magnet} and the resistance R_{parallel} of the parallel resistor. Hence, to reduce the leakage current, either the magnet inductance should be decreased or the parallel resistance should be increased. Another strategy for reducing the leakage current is the use of diodes as by-pass elements. This has the advantage of low leakage currents, because the voltage across the magnet is usually lower than the forward voltage of the diode, which is comparatively high at low temperatures [42, 25].

Energy extraction system

If the magnetic energy is dissipated in an external series resistance, this is referred to as a system with **Energy Extraction (EE)**, which falls under the category of active external systems [43]. The basic electrical scheme of such a system is shown in figure 1.4c. In normal operation, a switch is closed in parallel with the energy extraction resistor, so that the transport current flows only to a small extent through the energy extraction resistance. However, if the QDS is triggered, the (electrically controlled) switch is opened in order to force the current through the energy extraction resistor. This accelerates the discharge of the magnets. Additionally, most of the energy dissipation no longer takes place inside the magnet coil but predominantly outside. This reduces the risk of excessive heat input into the magnet and relieves the thermal load on the cryogenic system. Ideally, the EE-resistance could be set to very high values in order to drive the energy dissipation forward as quickly as possible. However, concerning the design value of the EE-resistance a compromise must be made, as the associated voltages across the EE system could possibly become unacceptably high for safe laboratory operation [25].

Quench Heaters and CLIQ

Active systems with internal energy dissipation are another option for the protection of superconducting magnets. In this context, the **Quench Heater (QH)** system [44] and the **Coupling Loss Induced Quench (CLIQ)** system [25, 45] are particularly worth mentioning. Both systems actively return the part of the superconductor to its normal conducting state in a controlled manner after a quench has been detected.

The QH systems installed in the LHC are approximately 0.025 mm thick strips of stainless steel embedded in a polyamide insulation [46]. They are densely pressed in between the cable guides and the coil windings as depicted in figure 1.5a, in order to transfer the forces acting on the cables as directly as possible to the magnet collar and to ensure a high thermal coupling to the coil windings. Most sections of the quench heater are copper coated so that the strips predominantly emit heat in the vicinity of the bare steel, high-resistance areas, the so-called heating stations.

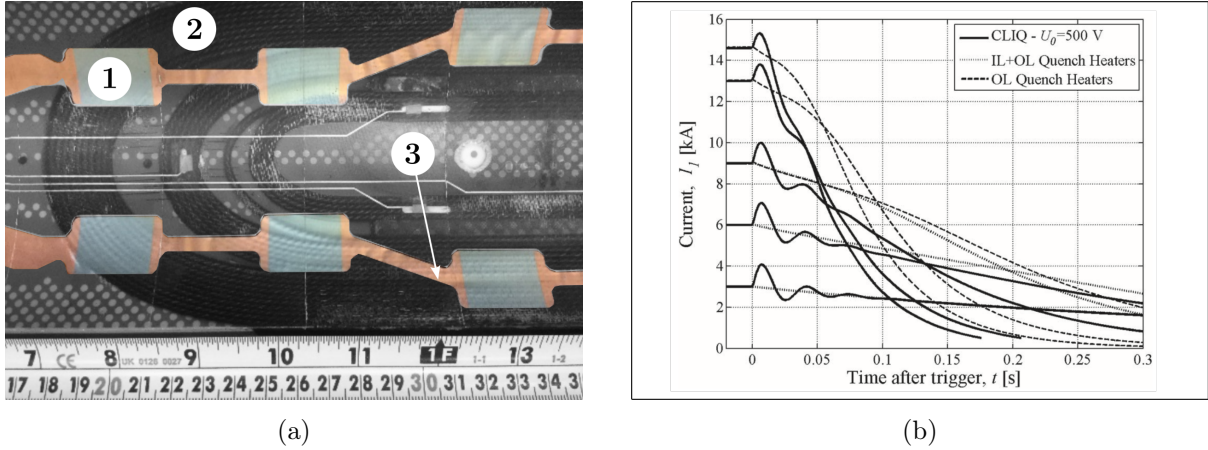


Figure 1.5: Active internal quench protection systems as they are installed in the LHC. a) quench heater in the inner layer of a magnet: (1) uncoated part of the stainless steel quench heater strip, acting as a heating station, (2) coil winding, (3) copper coating (modified image - original from a presentation of Todesco [47]). b) Magnet current over time after trigger for CLIQ and Quenchheater systems (modified image - original from the work of Ravaioli et al. [48]).

The relatively novel CLIQ system, on the other hand, is an alternative to quench heaters that offers a few advantages. With the CLIQ system, inter-filament and inter-strand coupling losses are used to heat the cables. This has the advantage of a more uniform, faster and more direct heat transfer compared to quench heaters, whose heat transfer is based on relatively slow thermal diffusion. These losses are introduced by discharging a bank of capacitors through dedicated current leads directly connected to the magnet coils. This leads to rapidly oscillating currents, as shown in figure 1.5b. In many cases, a faster discharge of the magnet can be achieved using this novel operating principle [49, 48, 50].

1.4 LHC main dipole circuit and FPA events

The LHC consists of more than 2000 individual circuits, which can be subdivided into different circuit families [1]. Since this thesis focusses on the improvement of circuit models of the main dipole circuit, this type of circuit will be introduced in more detail below, based on figure 1.6. The LHC comprises eight individual main dipole circuits (RB), one for each octant of the LHC. Each of these circuits comprises 154 magnets with a magnetic length of approximately 14.3 m [51, 52], which are needed to keep the particle beams within the accelerator on their circular paths. Each dipole magnet is protected by quench protection heaters and a by-pass diode, in order to

provide a low resistance path for the current, once a magnet turns back to normal conducting state due to a quench. The dipole circuits are powered by a high-precision **Power Converter (PC)** (i.e. a power supply) with a maximum output current of 13 kA [51]. In addition, there are two so-called **EE** units, one that is attached at the end of the magnet chain (called EE even) and one that is installed in the middle of the magnet chain (called EE odd) [52, 38]. If disturbances within the circuit, such as e.g. a quench in one of the magnets are detected, a **Fast Power Abort (FPA)** is triggered. This **FPA** can be subdivided into two stages [52]. First, the power

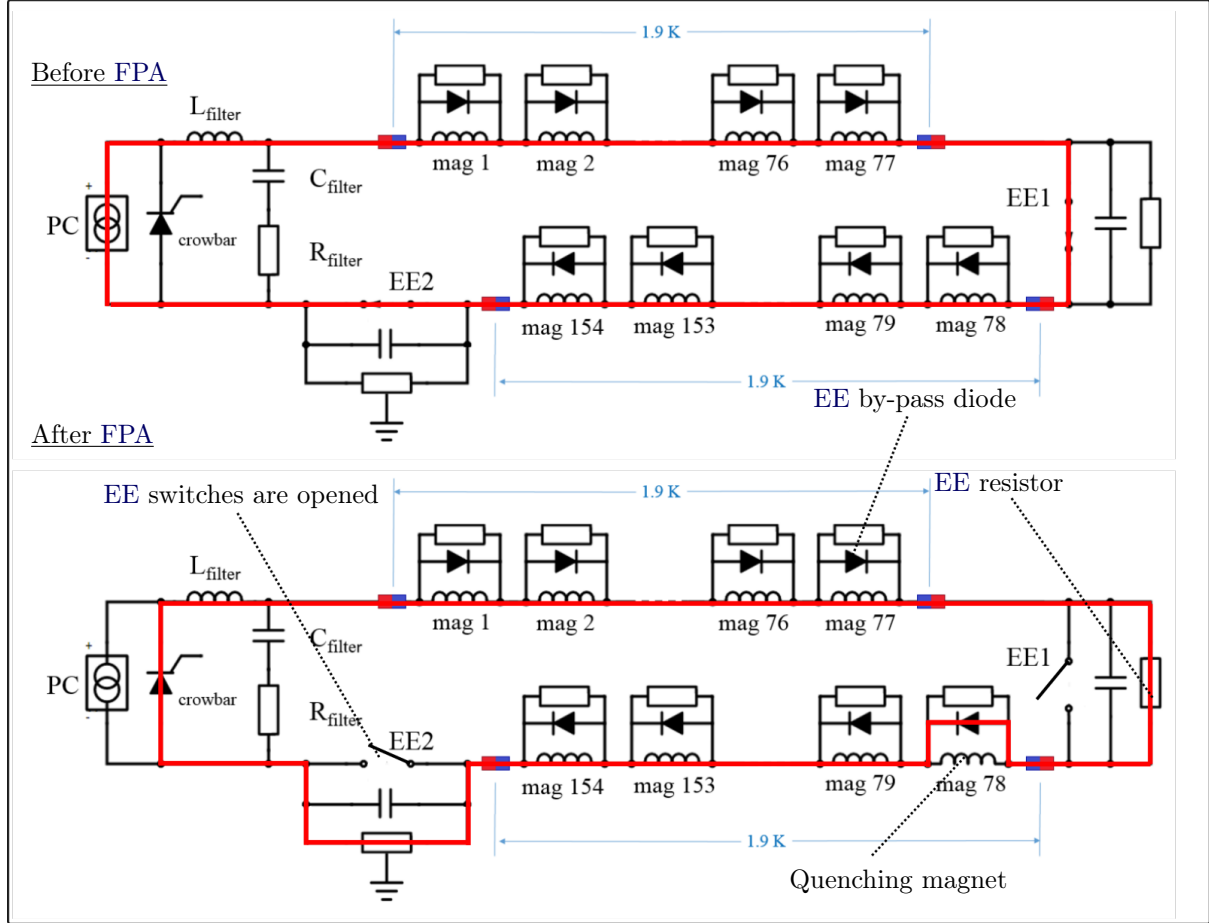


Figure 1.6: Simplified schematic of the main dipole circuit (modified image - original taken from the work of Rowan [51]). Top: Current flow before the **FPA** is triggered. Bottom: Current flow after the **FPA** is triggered, assuming a quench in magnet 78 occurred.

converters are switched off, in order to initiate the discharge of the circuit. However, even after switching off the power supply there is still a current flow which is forced by the high inductance of the magnets. The power converters are hence by-passed by a crowbar, providing a branch for the current flow once the power converters are switched off. Subsequently, the energy extraction is triggered in order to rapidly reduce the circuit current, and thus ensure that the operation of the circuit can quickly be resumed. In this step, the switches of the two energy extraction systems are opened one after the other, causing an increased energy dissipation within the **EE** resistor modules.

Hereinafter, the most important components of the RB dipole magnets shall briefly be presented, based on figure 1.7. The main dipole magnets have two apertures within the same cold




Figure 1.7: Cross-section of a main dipole magnet (modified image - original taken from the CERN document server [53]).

mass, through which two particle beams travel in opposite directions. Inside the beam pipe is what is known as the beam screen. This beam screen is essentially a tube of metal inside the beam pipe. It is used to shield the magnet from the synchrotron radiation and is cooled with liquid helium to increase its lifetime. The coils themselves are located inside the magnet collar, which is made out of a paramagnetic material that absorbs the forces acting on the coils and holds them in place. The coil windings are partially interrupted in the circular cross-section by copper wedges to achieve field homogeneity in the middle of the magnet [51]. The magnet collar is housed in an iron yoke, which plays a key role in shaping and guiding the magnetic field [51]. The iron yoke not only contains the magnet collar, but also a heat exchanger tube, which is used to cool the entire cold mass of the yoke with liquid helium. Since dipole magnets alternate with quadrupole magnets in transverse direction of the LHC, there are two different types of busbars providing different current levels in the magnet cross-section - the main quadrupole busbars and the main dipole busbars. Both types of busbars are located in individual grooves within the iron yoke. In order to prevent the cold mass from heating up due to thermal exchange with the environment, the inner part is covered with so-called superinsulation. To improve the thermal insulation, there are additional shells around the inner structure, the spaces between which are evacuated. Finally, the entire structure is surrounded by the so-called vacuum vessel, which also acts as a hard outer shell for the safe transport of the magnets.

1.5 Simulation of superconducting magnet circuits

Simulations are a cost and time-efficient complement to experiments for researching the technology of superconducting magnet circuits. In contrast to measurements however, simulations enable studying physical processes independently of any measurement related disturbance, which is another reason why they have become an essential analysis tool for this application. The level of detail and accuracy of these simulations is strongly linked to the technological feasibility of

new generations of compact and powerful accelerator magnets. However, systems already in operation also benefit from the advantages of simulations. After a magnet quench, for example, informative analyses of associated currents, voltages and temperatures can be conducted. These physical quantities are not only crucial for the safe operation of the magnet, but also have a significant influence on the magnetic field, which is subject to specific requirements. In dipole magnets for example, a strong magnetic field that is as homogeneous as possible is required, as otherwise, even the smallest differences in position of the particles passing through the magnet could lead to an undesirable broadening of the beam. With regard to temperatures, the hot-spot temperature is of particular interest, as high hot-spot temperatures can potentially damage the magnet as already explained in chapter 1.

Since the physical processes in superconducting magnet circuits are complex multiphysics problems, simulations must take into account a wide range of interdependencies between physical quantities.

For small to medium-scale simulation problems, such intricate codependencies are often solved by using multiphysics software packages such as COMSOL, which rely on **Finite Element Method (FEM)** approaches.

However, since the physical mechanisms determining the behavior of superconducting magnet circuits span different scales, it is not possible to exclusively simulate them using these conventional approaches described above, as the simulations would become computationally too expensive.

Conclusively, the electrothermal simulation of superconducting magnet circuits requires the development of new tools and methods, which are not yet commercially available. Hence, the **Simulation of Transient Effects in Accelerator Magnets (STEAM)** framework was developed, which is further described in the following chapter 2.

Methods and tools for optimizing superconducting magnet circuit models

2.1 STEAM

2.1.1 Objective of STEAM

To efficiently solve the challenges discussed in section 1.5, CERN developed a simulation framework, the so-called **STEAM** framework. By linking different, partly in-house developed software packages and tools within a modular approach, **STEAM** can be used to split complex multi-scale problems into individual domains, solving each domain with a specialized tool. Among other tools, the framework includes tools like SIGMA [54], which can be used for the automatic generation of COMSOL models, FiQuS a **FEM** tool for superconducting magnet quench simulations [55, 56], and ProteCCT, a tool for the simulation of canted cosine-theta magnets [57]. Additionally, **STEAM** offers the possibility to use several tools simultaneously within a **cooperative simulation** (co-simulation) in order to solve the overall problem computational efficiently on different scales [58, 59]. In this way, for example, netlist simulation tools (such as PSPICE or XYCE) can be used for macroscopic simulations of entire circuits, while individual magnets, which are largely determined by microscopic effects, are analyzed using **FEM**-based tools such as **STEAM-Lumped Element Dynamic Electro Thermal (STEAM-LEDET)** [60, 25]. Furthermore, the framework automates large parts of the traditionally time-consuming process of setting up a simulation analysis, through relying only on a minimalistic input file in `.yaml` format [38, 61, 62].

2.1.2 Structure of the STEAM-SDK

The **STEAM-Software Development Kit (STEAM-SDK)** [63, 64] defines the structure of the analyses and provides classes with auxiliary functions. The **STEAM-SDK** consists of multiple modules, each providing different functionality as presented below:

- **Analyses:** This subfolder contains the main class **AnalysisSTEAM**, in which the structure of the analyses is defined. The **AnalysisSTEAM** class can also be initialized by a `.yaml` file in order to set up the analysis defined in it fully automatically.
- **Builders:** The builders are used to decompose model data into the data required for the respective simulation tool used (e.g. COMSOL, XYCE, **STEAM-LEDET**...). In a later step this tool specific model data is often exported in order to create the input file for the simulation tool used.
- **Data:** This module contains data classes that define attributes of other modules. It relies heavily on the use of the pydantic library, which ensures automatic data validation, type checking and documentation of all data objects defined within the data class.

- Drivers: In [STEAM](#), python classes for launching the individual tools are referred to as drivers. In the simplest case, such a class can start a sub-process that calls the tool executable together with the tool's input file.
- Parsers: Within a [STEAM](#) analysis, several software tools can be used, which come with a variety of different file formats to handle. For this, so-called parsers are used to translate and export the model data into different formats. For instance, [STEAM](#) can employ the [ParserPSPICE](#) class for writing a netlist solver input file (`.cir`), based on information stored in the respective data object.
- Parsims: Software modules that are intended for setting up multiple analyses with varying simulation parameters are called parsims (parametric simulations) in [STEAM](#). These modules are used for parameter studies and optimizations, such as the [STEAM-Dakota Interface \(STEAM-DI\)](#), whose development was continued within this thesis.
- Postprocs: This module contains methods for postprocessing the analysis data, enabling the comparison of measurement and simulation data by determining error metrics.
- Viewers: This visualization module can be used to display measurement data and simulation data for comparison, relying on a configuration file, as further discussed in section [2.1.3](#).
- Other utilities: In addition to the modules essential for this thesis, the [STEAM-SDK](#) additionally contains several other modules and utility classes that contain further functionalities for setting up analyses.

2.1.3 File conventions

In the following, the most important files used within the [STEAM](#) framework will be briefly presented, as their functionalities are essential for the methods described in the further course of this chapter.

- Analysis file (`.yaml`): The analysis file serves as an input for simulations with [STEAM](#), containing all user defined input necessary for setting up an `AnalysisSTEAM` object. The minimalist encoding (e.g., no brackets like in `JSON`) as well as the hierarchical structure of this file format not only ease file manipulation but also enable the stored data to be directly parsed into the data classes of the framework.
- Stimulus file (`.stl`): As already explained in section [1.2.1](#), the behavior of superconducting magnets during a quench is determined by nonlinear dynamic interactions of currents, temperatures and fields, which makes them extremely complex systems. Since the complexity involved in modelling them can not be easily captured within netlist simulation tools, the behavior of the magnets is determined by other tools and only the pertinent aspect, essential for circuit simulation - a time-dependent resistance response of the magnet - is later on imported to the netlist solver tool. This time-dependent resistance response is defined in tabular form within a so-called stimulus file. To generate the stimulus file, [co-simulations](#) are carried out that can describe the behavior of the quenching magnet within the complete circuit very precisely [\[38\]](#). As the discharge of the magnet is strongly

dependent on the current level, at which the magnet quenched, [co-simulations](#) were repeated at several current levels. For quenches at current levels that have not already been taken into account in a previous [co-simulation](#) however, the magnet's response is obtained by performing an interpolation between available solutions. Hence, the import of presimulated magnet responses within stimulus files enables to perform very accurate simulations of large circuits, almost as accurately as with [co-simulations](#) - but much faster.

- Configuration file (`.yaml`): The configuration file contains metadata on measurement and simulation data, necessary for postprocessing. This includes the names of various simulation signals required for plots, as well as information about which types of measurement and simulation data represent the same physical quantity. Furthermore, the metadata contained in the config file is queried by various steps, such as the `RunViewer` step or the `CalculateMetrics` step, as described in section 2.2.
- Event file (`.csv`): In case a quench event occurs, a [FPA](#) is triggered and measurement signals of the circuits are stored within two databases - Post Mortem and [NeXt CERN Accelerator Logging Service \(NXCALs\)](#). Before the [LHC](#) can then be put back into operation, the quench cause must be identified from these signals, in order to ensure safe operation. This is done using analysis scripts implemented in jupyter notebooks, which were developed as part of the SIGMON project [65, 66]. In this process, essential characteristics, such as the current level at which the magnet quenched, are exported and saved in `.csv` files. These files, which are called event files in the context of this thesis, contain all the information about the quench event needed for a circuit simulation. Hence event files can serve as an input for the so-called `ParsimEvent` step implemented in [STEAM](#), which will be discussed in section 2.2

2.2 Automated simulation and validation of magnet circuits

2.2.1 Netlist simulations

Within this thesis, simulations are mainly modeled with so-called netlist solvers. In netlist format, circuits are defined by linking individual nodes through the definition of circuit components connected to them. This is done in a list in which the type of component, the connected nodes and the physical characteristics of the component are stored line by line. This clear structure of the circuit, which does not rely on a graphic representation of the circuit's schematic, makes it easy to programmatically read, edit, and run the circuit models, which is why netlist files are a widely used input format for circuits simulations. The system of differential equations defined by the entries of the netlist is solved by a so-called netlist solver. Within this thesis, the PSPICE netlist solver was used for this purpose, as it provides direct options for plotting individual signals and debugging the circuit via a graphical user interface.

2.2.2 Analog behavior modelling

Within this thesis, several circuit models of components have been implemented, whose real-world electrical characteristics are strongly temperature-dependent. Therefore, their temperature must be considered within the simulation by evaluating their thermodynamic equilibrium

Table 2.1: Analogies between thermal and electrical systems used in ABMs.

Thermal domain	Electrical equivalent
Temperature T (K)	Voltage V (V)
Heat P (W)	Current I (A)
Heat Capacitance C_v (J K^{-1}) $(T - T_0) = \frac{1}{C_v} \cdot \int P_{\text{heat}} \cdot dt$	Electrical capacitance C (F) $(V - V_0) = \frac{1}{C} \cdot \int I \cdot dt$
Thermal conductance h (W K^{-1}) $P_{\text{cool}} = h \cdot (T - T_0)$	Electrical conductance G (Ω^{-1}) $I = G \cdot (V - V_0); \quad G = \frac{1}{R}$

(according to first law of thermodynamics). However, since the netlist simulation tools used did not provide functionalities for direct calculation of thermodynamic quantities, thermal modelling had to be carried out using a so-called **Analog Behavior Model (ABM)**. Within these models, analogies between electrical and thermal systems are used to describe the thermodynamic balance of the modelled circuit component [67]. The analogy used in the thesis is based entirely on the following two principles:

- Temperatures are represented by electrical voltages.
- Power (heating/cooling) is represented by electrical currents.

These principles directly lead to further analogies: thermal heat capacities behave like electrical capacitors, and thermal heat conductors behave like electrical conductivities, as illustrated in table 2.1. In practice however, the aforementioned concepts cannot always be implemented. Particularly in the case of temperature-dependent heat capacities and thermal conductivities, the component behavior cannot be modeled using passive components, as these can only be defined statically in the PSPICE netlist solver. In such cases, voltage-controlled current or voltage sources are implemented that allow the component parameters to be defined dynamically.

2.2.3 Structure and sequence of the analysis

As already explained in section 2.1, each **STEAM** analysis is divided into individual steps. During this thesis, the **STEAM** framework was expanded in order to include functionalities that enable both simulation and quantitative validation of the results within the same analysis. Subsequently, the individual steps of such an analysis are presented, following the corresponding analysis file's structure as shown in 2.1b. At first, the **MakeModel** step is executed. In this step, the input file (`.cir`) for the simulation is written to the simulation folder, based on the model data stored in the `steam_models` package. In principle, this file would already be sufficient to simulate the circuit. For validation and optimization however, the simulation's input file must be modified so that the output includes all relevant signals. This is done in two consecutive **ModifyModelMultipleVariables** steps. As part of these steps, individual or multiple features of the original input file can be modified. The initial **ModifyModelMultipleVariables** step ensures the simulation results are saved in the common simulation data format (`.csd`), by changing the input file's probetype key from its default **PROBE** to **CSDF**. Subsequently, the second of the two **ModifyModelMultipleVariables** steps defines which signals should be included in the output file of the simulation. In this step, the list of signals to be considered within the validation

is derived from the entries specified in a configuration file (see section 2.1.3 for details about the configuration file). At this stage of the analysis, the simulation's folder structure has been generated, and it has been ensured that all necessary signals will be included in the output file. However, another `ModifyModelMultipleVariables` step has to be invoked in order to align the simulation parameters with the current parameter vector chosen by the optimization algorithm. Once all parameters have been adjusted, the actual simulation study can be arranged using the `ParsimEvenet` step. In the `ParsimEvenet` step, one or more simulations are prepared based on the information stored in the event file. Initially, a list of quench events to be simulated is generated (parsim sweep file), detailing the parameter configurations needed for their simulation. Based on this list, a comprehensive parameter study, covering all magnet quenches included in the eventfile, is set up. Additionally, other files required for the simulation (such as the stimulus files introduced in section 2.1.3) are written to the simulation folder during this step. Upon completion, an initialization file for the viewer is written, containing all information needed for initializing a viewer object. Only now, within the `RunSimulation` step, the actual simulation is carried out. This involves the call to a tool-specific driver, which launches the respective simulation tool. Once the simulations are completed, the analysis proceeds with the postprocessing of the simulation results, invoking the `RunViewer` step and the `CalculateMetrics` step. In the `RunViewer` step executed first, a `.pdf` report is generated showing comparative plots of measurements and simulation data as specified in the config file. Subsequently to any analysis or optimization, these plots serve as an initial sanity check, helping to identify unexpected behavior or potential flaws in the models. In the concluding `CalculateMetrics` step, a quantitative validation of the simulation results is carried out. As part of this step, error metrics are calculated, quantitatively describing the agreement between simulation data and measurements. Finally, these metrics are saved within an attribute of the `AnalysisSTEAM` object for later processing.

2.2.4 Acquisition of measurement data

Whenever a model is to be validated, experimental data is needed. Since the magnet circuits of the LHC are subject to constant monitoring, a large amount of measurement data is available - stored within CERN's internal databases: `Post Mortem (PM)` [69] and the `NXCALS` [70]. These two databases differ primarily in terms of the time spans and resolutions of the measurement data they provide. Even though `PM` does not provide continuous logging data like `NXCALS`, it profits from a higher resolution in the time domain with respect to `NXCALS`, leaving both databases with distinct advantages. Consequently, using measurement data from both databases is beneficial, especially because the reliability of the validation profits from a larger and more comprehensive reference dataset. Hence, a python script was developed to facilitate the signal acquisition by automatically querying and saving all signals needed for the validation as specified in the configuration file. The script is largely based on the `LHC-Signal-Monitoring-API (LHC-SM-API)` [71] developed at CERN, a python library that provides methods for accessing the API interfaces of the databases. Instead of acquiring the data during the analyses, the measurements were acquired before the analysis and saved in an offline folder in order to mitigate latencies associated with server queries and thus further speed up the validation process. The measurement data were saved in `.feather` file format, which benefits from high read and write speeds for tabular time series data [72]. To avoid accidentally comparing the wrong measurement data to a simulation, all data series were saved with an unique identifier containing the

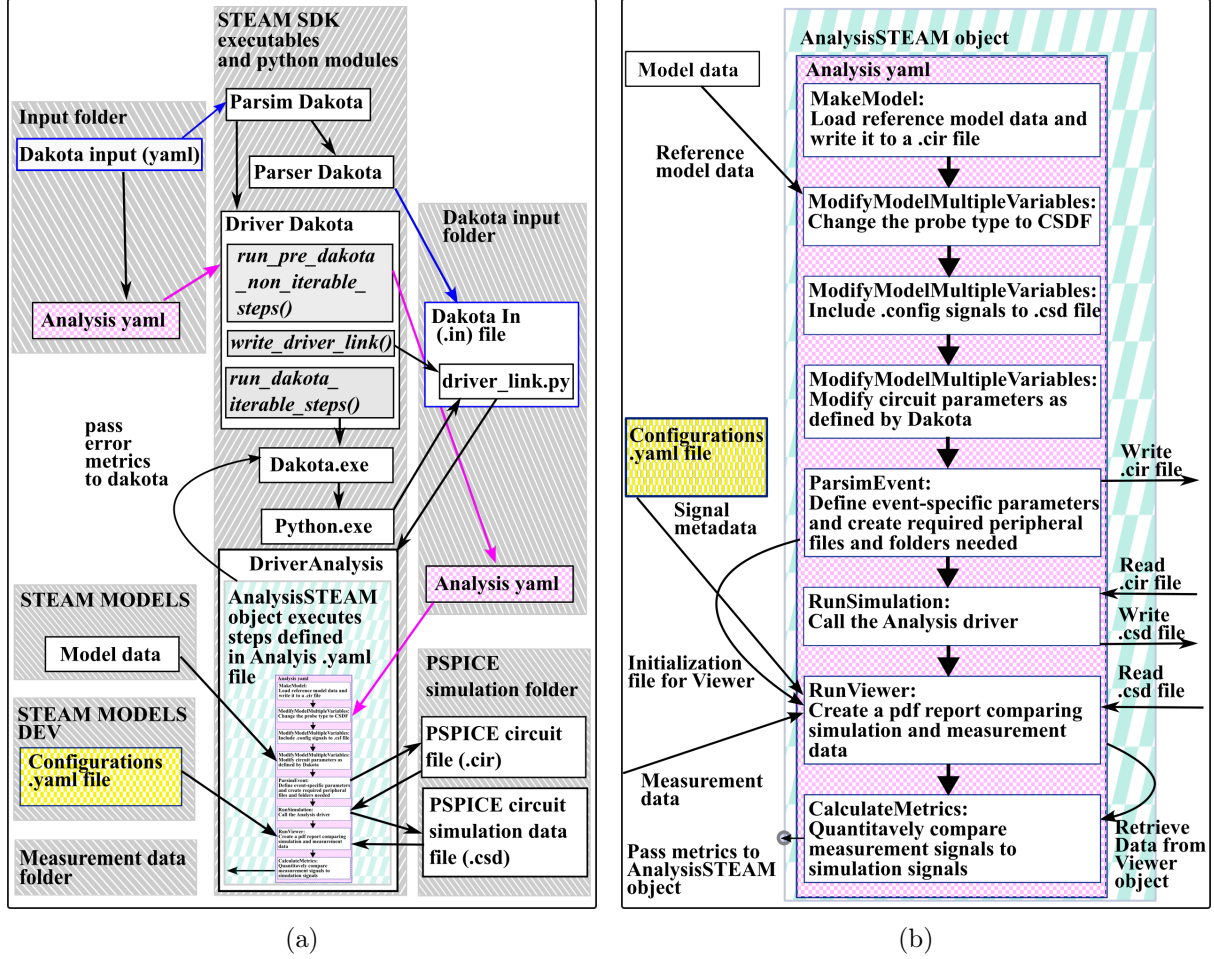


Figure 2.1: Structure of the STEAM-DI with detailed visualization of the processes executed by the AnalysisSTEAM object (image is based on a figure on the steam-sdk github website [68]). a) General structure of the STEAM-DI. For optimization with the STEAM-DI, additional files are required, which are obtained from multiple peripheral folders, like the measurement data folder, or packages such as the `steam_models` package. b) Steps of an analysis for circuit simulation and validation with the STEAM framework. The graphic features a closeup of the green analysis block shown in figure a.

time stamp of the corresponding quench event as well as the name of the circuit.

2.3 Closed-loop multi-objective optimization

The optimization of models relying on multiple parameters can be highly time-consuming, especially if the sensitivity of the model to individual parameters is unknown and the effort for model validation is high [73]. The optimization of superconducting magnet circuits is also affected by this, as the influence of individual components is usually non-linear, and the simulations, which can still take several minutes on a powerful computer, are too resource-intensive to simulate the entire parameter space. For this reason, it is desirable to reduce the number of simulations until finding an optimum parameter configuration. With this objective in mind, this thesis employs **Closed-Loop Optimization (CLO)** techniques for parameter fitting. In these methods, the choice of the next parameter vector to be tested is informed by feedback from the last optimization step [73]. In the context of this thesis, this means that the physical parameters of the next simulation are selected based on the error metrics obtained in the preceding optimization steps. For this purpose, an interface between the parameter optimization toolkit Dakota and **STEAM** was further developed.

2.3.1 STEAM-Dakota interface

Today, numerous algorithms are available for optimizing model parameters, each of them having their unique advantages and drawbacks. However, manually implementing such algorithms directly within the simulation framework can be time-consuming and error-prone, especially with more complicated algorithms. In order to alleviate these drawbacks, third party libraries containing optimization algorithms hold a promising solution. These optimization toolkits not only profit from a wide variety of algorithms available, but also from a wider user base and hence better validation. For the optimizations during this thesis, the Dakota optimization toolkit [74, 75] developed by Sandia National Laboratories was employed, as it is widely used and easily available, being open source. In order to enable the use of Dakota with the netlist solver PSPICE, an already existing interface [76] was further developed within this thesis, which in the following is referred to as the **STEAM-DI**.

Hereinafter, the processes within an optimization using the **STEAM-DI** will be introduced by following its process diagram shown in figure 2.1a. When performing an optimization with Dakota, the system to be optimized is treated as a black box, meaning that the optimization does not consider details within the system but solely its input and output parameters. Hence, the user needs to define interfaces inside an input file (`.in` file), which are used for both changing the black-box parameters as well as for retrieving feedback from it (value of an objective function). The **STEAM-DI** follows a principle of minimal user input for being set up, solely relying on a Dakota input yaml file containing only the most essential information for the optimization. This also includes which of the analysis steps defined in the analysis file must actually be repeated during the optimization iterations. In this way, e.g. the **MakeModel** step described in the example analysis in figure 2.1b, which would otherwise repeatedly generate a similar folder structure, can be excluded from the iteratively repeated analysis steps, thus saving computational effort. This distinction between steps that only need to be executed once at the start of the optimization and those that need to be repeated and possibly modified in each iteration is made by defining

two lists `initial_steps_list` and `iterable_steps_list` within the `ParserDakota` inputfile (`.yaml`). The script that is started initially is the `ParsimDakota` script. Since the `STEAM-DI` iteratively changes the parameters of `STEAM` analyses, it is located in the `parsims` subfolder of the `STEAM-SDK`. In the very beginning of the optimization's set up phase, the `ParserDakota` inputfile (`.yaml`) is translated into a Dakota infile (`.in` - the input file format required by Dakota), using the `ParserDakota` class. Later on, the `DriverDakota` is called for starting the optimization. Moving forward, the driver consecutively launches three distinct functions. First, the steps from the `initial_steps_list` are run by calling the `run_pre_dakota_non_iterable_steps()` method. Second, the `write_driver_link()` function writes a python script `driver_link.py` into the Dakota input folder, which acts as the actual interface between the system to be optimized and Dakota. Finally, the driver invokes the function `run_dakota_iterable_steps()`, calling `Dakota.exe` with the Dakota infile (`.in`), whereupon Dakota starts the optimization process. The optimization process, can be subdivided into three iteratively executed phases:

- First Phase - Definition of a parameter configuration: Based on the metrics from the preceding iteration steps, the parameter configuration to be tested in the subsequent step is determined and passed to the black-box model. In the first iteration, where no feedback is yet available, Dakota passes initial values which are defined by the user.
- Second Phase - Execution of the blackbox code: In the second phase, a python script `driver_link.py` is executed with the parameters designated in the first phase. This script is a wrapper utilized to enhance the structure of the code. Its essential function is to launch the `DriverAnalysis` passing along some parameters previously specified in the Dakota infile. Within the analysis driver, a new analysis file is prepared in order to align the simulation to the parameters specified by Dakota. For this, the parameters are retrieved from Dakota and incorporated into a `ModifyModelMultipleVariables` step within the newly generated analysis file. In addition, the novel analysis file's simulation number key is set to the current optimization iteration number, so that instead of overwriting simulation results, a new simulation folder is made for each iteration. Finally, all steps defined within the analysis file are executed by initializing an `AnalysisSTEAM` object based on the modified analysis file. The metrics calculated in the last part of the `CalculateMetrics` step are saved in an attribute (`AnalysisSteam.summary`) of the analysis steam object and then transferred to Dakota by the analysis driver.
- Third Phase - Retrieving the objective function: In the final phase, the value of the objective function (error metric) is retrieved and evaluated. If a termination criterion defined by the user is met, the optimization ends. If not, the three phases are executed again.

During the optimization, Dakota writes tabular data of the tested parameter configurations and their corresponding objective functions to a file, enhancing the processes' traceability.

2.3.2 Algorithms

The algorithms implemented in Dakota capable of being employed into a closed-loop multi-objective optimization can essentially be divided into three subgroups, which are briefly presented below [74]:

2.3.2.1 Gradient-descent methods

In gradient descent methods, gradients of the objective function within parameter space are used to determine the impact and sensitivity of individual parameters. For this, Dakota can retrieve gradient information either analytically or numerically. In case of analytical gradients, the gradient information directly stems from the system optimized, being passed along with its other outputs to Dakota. However, if such information is not available, because the system does not provide it, Dakota can obtain the gradients numerically. In such cases, Dakota evaluates the objective function at different parameter vectors, and subsequently applies finite differencing to obtain the gradient information. Even though an additional simulation step per parameter is needed to calculate a gradient in such cases, employing a gradient descent method can strongly reduce the amount of iterations needed until an optimum is reached compared to traditional parameter studies. However, this approach requires reliable gradients and thus a smooth objective function, as otherwise the optimization may not converge to an optimum. Furthermore, gradient-based methods search for the optimum closest to the starting point, leaving large parts of the parameter space untested. Even though this usually is a computational advantage, they are thus unreliable for searching global optima.

2.3.2.2 Derivative-free local methods

In derivative-free approaches, the parameter space is explored based on designated sampling patterns. As implied by the name, the parameter vector's shift is not determined by evaluating a gradient, but rather by considering the objective function at different parameter vectors, making it more robust in case of unreliable or non-smooth gradients. Since the positions of the parameter vectors within the same pattern do not rely on each other, this type of algorithm can be easily parallelized by testing several parameter vectors within a pattern congruently. However, this comparatively uninformed selection of parameter vectors often comes at the cost of a slower convergence rate [74].

2.3.2.3 Derivative-free global methods

The search for global optima is always associated with an extensive exploration of the parameter space. Even though global optimization algorithms such as evolutionary algorithms aim to reduce the extent of sampling needed, they still are computationally more expensive than local optimization methods.

2.3.2.4 Algorithm used

For the optimizations within this thesis, a derivative-free local method (colony pattern search) was chosen for the following reasons:

- Global method versus local method: The use of global methods in this application is of little interest, as the approximate position of the optimum parameter vector is already known in advance. A global optimum far from the assumed position would rather be an unphysical artifact of overfitting than the result of a valid deviation from the circuit parameters' expected designvalues, which are subject to regular monitoring.

- Derivative-free method versus gradient descent method: Due to the unstable gradient information (some simulations might freeze and have to be aborted due to unfavorable parameter configurations), a derivative-free method has to be used. This method benefits from increased robustness compared to local gradient-based methods.
- Unconstrained versus constrained method: There are no physical boundary conditions for the optimization that are not already implemented in the simulation code itself. Hence, the selected optimization algorithm does not necessarily have to allow the definition of further constraints.

Model of the LHC main dipole circuit's energy extraction resistors

As part of this thesis, a thermoelectric model of the main dipole circuit's (RB) energy extraction resistors was implemented. This model considers the resistors' increase in resistance due to heating effects associated with the current flows through them. Section 3.1 first provides a general motivation and discusses the modelling details. Subsequently, the optimization process and its preparation are discussed in more detail in section 3.2. The results of the simulations and model optimizations are then presented in section 3.3. Finally, the results are summarized and discussed in section 3.4.

3.1 Model

3.1.1 Thermodynamic equivalence model

Each of the eight LHC main dipole circuits comprises two energy extraction units. The units mainly consist of three EE resistor modules, which are connected in parallel to a switch branch as illustrated in figure 3.1. If the switch branch is opened, shortly after a FPA is triggered, the magnet current is forced to flow through the resistors, each having a design resistance of 225 mΩ. Before the Long Shutdown 1 (LS1) (2013 - 2015), two modules connected in series were installed with another module in parallel, resulting in an overall resistance of $(\frac{1}{2 \cdot 225} + \frac{1}{225})^{-1}$ mΩ = 150 mΩ. However, as the LHC was upgraded and its magnets were powered to higher current, a smaller extraction resistance was needed in order to keep the voltages across the EE at the design level. Hence, after 2015 the modules were connected in parallel, resulting in a total resistance of approximately $\frac{225}{3}$ mΩ = 75 mΩ. The resistance of the modules is mainly determined by 84 stainless steel (Type 304L) plates as shown in the middle of figure 3.1, which are connected in series. Furthermore, one of the three modules is connected to the LHC's earthing system after 42 plates, as illustrated in the left of figure 3.1. Since the current flow through the resistive plates during a discharge is associated with ohmic heat losses, the resistance of the modules increases significantly over the course of a discharge. In order to prevent the modules from overheating, the resistive plates, being placed with a distance of 4 mm to each other, are cooled by a forced air flow generated by four fans placed in the front and back of the module as shown in the right-hand side of figure 3.1. The hot air is then lead through two heat exchangers, transferring the heat to a water reservoir, with a water exchange rate of about 30 L min⁻¹ [43].

In the further course, a PSPICE model of the EE resistance is described which is taking into account changes in the resistance due to aforementioned heating and cooling effects. The model previously implemented at CERN, on which the model implemented in this thesis is based, neglected the effects of heating and cooling, which was associated with increased errors in both the voltage across the energy extraction V_{EE} and the current through the circuit I_{MEAS} , especially at high FPA current levels I_0 . In order to model the heating and cooling behavior,




Figure 3.1: Structure of the RB EE System (images based on technical drawings from CERN). Left: Simplified electrical schematic showing the three energy extraction modules and the way they are connected with respect to the circuit and its earthing system. Middle: Detailed view showing the stainless steel plates which are connected to the two power input pins and the earthing system. Right: Structure of a single EE module and its cooling system.

a thermodynamic state equation of the system is derived by considering an EE module's thermodynamic equilibrium, after a FPA is triggered. This is done starting from the first law of thermodynamics:

$$\sum_{l=1}^m \dot{Q}_l + \sum_{r=1}^e \dot{W}_r + \sum_{i=1}^k \dot{m}_i \cdot \left(h + \frac{v^2}{2} + gz \right)_i = \frac{d}{d\tau} \left(\sum_{j=1}^n m_j \cdot (\bar{u} + \bar{e}_{\text{kin}} + \bar{e}_{\text{pot}})_j \right). \quad (3.1)$$

In this equation, \dot{Q}_l are the heat flows transferred across the system boundaries, \dot{W}_r are the powers introduced or extracted from the system and \dot{m}_i are the mass flows across the system's border having an enthalpy h_i , a velocity v_i and a potential energy gz_i . This equals to the system's rate of change in energy, which is given by the right-hand side of the equation. This rate of change on the right-hand side considers \dot{m}_j as the mass of the j^{th} particle or body of the system which has mass specific internal energy \bar{u} , kinetic energy \bar{e}_{kin} and potential energy \bar{e}_{pot} . In the model derived, the resistive plates are considered as a homogeneous, coherent body with a uniform resistance, heat capacity and heat transfer across the system boundaries. The thermodynamic control volume is hence defined as shown in figure 3.2. As the control volume is closed and the heated mass is not subject to changes in kinetic or potential energy, equation 3.1 can be simplified as follows:

$$\sum_{l=1}^m \dot{Q}_l + \sum_{r=1}^e \dot{W}_r + \sum_{i=1}^k \dot{m}_i \cdot \left(h + \frac{v^2}{2} + gz \right)_i = \frac{d}{d\tau} \left(\sum_{j=1}^n m_j \cdot \left(\bar{u} + \overset{0}{\bar{e}_{\text{kin}}} + \overset{0}{\bar{e}_{\text{pot}}} \right)_j \right). \quad (3.2)$$

Furthermore, the heat transfer is considered to be limited exclusively to the cooling water, which is in the following denoted with $\dot{Q}_{\text{exchange}} = P_{\text{exchange}}$. This is justified by design constraints during the construction phase of the EE resistor modules, not allowing increased heat dissipation




Figure 3.2: Modeling of a EE module (left image based on a technical drawing from CERN and an image from Dahlerup-Petersen et al. [43]). a) Cross section of a single 225 mΩ EE resistor module. b) Sketch illustrating the thermodynamic model derivation of the resistor module.

to the surrounding air as the modules are installed in underground areas [43]. Additionally, the work introduced to the system is dominated by the ohmic heat inputs into the resistor plates, subsequently denoted with $\dot{W}_{\text{ohmic}} = P_{\text{ohmic}}$. Hence, the thermodynamic equilibrium of the EE resistor control volume is given by:

$$\underbrace{\dot{Q}_{\text{exchange}}}_{=P_{\text{exchange}}} + \underbrace{\dot{W}_{\text{ohmic}}}_{=P_{\text{ohmic}}} = \underbrace{\frac{d}{d\tau} m_{\text{plates}} \cdot \bar{u}}_{=P_{\text{Cv}}}, \quad (3.3)$$

with P_{Cv} being the power absorbed or released by the resistor's heat capacity. With this equation in mind, it becomes apparent that the system's energy equilibrium further results in a power equilibrium, relating the power absorbed or released by the resistor's heat capacity P_{Cv} to its heating and cooling powers:

$$\underbrace{P_{\text{ohmic}}}_{=I_{\text{res}}^2 \cdot R(T_{\text{res}})} + \underbrace{P_{\text{exchange}}}_{=-h_{\text{th_EE}} \cdot (T_{\text{res}} - T_{\text{amb}})} - \underbrace{P_{\text{Cv}}}_{=\frac{d}{dt} C_v(T_{\text{res}}) \cdot T_{\text{res}}} = 0. \quad (3.4)$$

(I)
(II)
(III)

In this formula, the ohmic heating power equates to the module's temperature-dependent resistance $R(T_{\text{res}})$ multiplied by its current I_{res} squared. Furthermore, the cooling power P_{exchange} is modeled with Fourier's law, describing the heat transfer to be proportional to the difference between the ambient temperature T_{amb} and the resistor's temperature T_{res} with the respective proportionality constant being a thermal conductance $h_{\text{th_EE}}$. The environment is considered as an infinite heat sink with a constant temperature of $T_{\text{amb}} = 293.15$ K. This is represented by a voltage source E_{amb} , which keeps the voltage node T_{amb} at a constant voltage of $V(T_{\text{amb}}) \equiv 293.15$ V. The neglect of radiative cooling is motivated by the fact that the maximum temperature of the resistive plates (usually smaller than 500 K) is comparatively low. Additionally, the calculation of higher order polynomial functions can lead to numerical instability of the circuit solver,

which is another reason why the Stefan Boltzmann law is not applied in the simplified model introduced here. Furthermore, convective cooling is neglected for complexity reasons, as the plates are cooled with a turbulent, time-dependent air flow which is cooled by a heat exchanger, causing the complexity of the convective cooling mechanisms involved to exceed the limits of what can reasonably be modelled with a 0-dimensional lumped-element model. Hence, there are three terms that contribute to the thermodynamic equilibrium: one term that describes the ohmic heating (I), one that describes the cooling (II) and one term describing the influence of the system's thermal capacity (III). Since each of the contributions corresponds to a power, they can be modelled by a current source using the analogies introduced in section 2.2.2. Hence, an equivalent electrical circuit diagram as shown in figure 3.3 can be derived, as the control volume's thermodynamic state equation now corresponds to a Kirchhoff node rule. In this circuit analogy




Figure 3.3: ABM for simulating an EE resistors's thermal equilibrium. The voltage source E_{amb} has a constant voltage of 293.15 V, corresponding to ambient temperature.

model, the heating and cooling powers are modelled using two PSPICE current sources. Since having multiple PSPICE current sources in parallel can cause numerical instability, the third power contribution P_{Cv} is modelled with a voltage source instead of a current source. Hence, contribution (III) is solved for $(T_{\text{res}} - T_{\text{amb}})$ in order to derive a formulation for the respective PSPICE voltage source $V_{(III)}$, which defines a voltage proportional to difference between the voltage nodes T_{amb} and T_{res} shown in the electric equivalent circuit:

$$P_{Cv}(t) = \frac{d}{dt} C_v(T_{\text{res}}(t)) \cdot T_{\text{res}}(t) \quad (3.5)$$

$$\frac{1}{C_v(T_{\text{res}})} \cdot \int_0^t P_{Cv}(\hat{t}) d\hat{t} \approx T_{\text{res}}(t) - \underbrace{T_{\text{res}}(t=0)}_{=T_{\text{amb}}}$$

For this transformation, the change in C_v over time was considered to be small, which as an approximation was justified by a small rate of change of $C_v(T)$ at ambient temperatures (see figure 3.4). With this assumption, C_v can be excluded from the integration, which fixed previous convergence problems of the simulation. In a future model version, this simplification could be avoided by simulating P_{Cv} with a current source, which would make the transformation shown in equation 3.5 obsolete. However, it should be noted that the heat capacity $C_v(T_{\text{res}})$ as well as the resistance $R(T_{\text{res}})$ are temperature-dependent variables. Accordingly, there is a mutual dependency between the temperature and these two quantities, so that the equations introduced

in this section and those described in the next section for calculating the material behavior are both needed to fully describe the module's thermo-electric behavior. In the next section, the still missing part of defining temperature-dependent material properties in PSPICE is introduced.

3.1.2 Temperature-dependent material properties

As introduced in the previous section, the resistivity and heat capacity of the modules' stainless steel plates are temperature-dependent. In the following it is explained how these temperature dependencies were considered in PSPICE, using figure 3.4. For each temperature-dependent




Figure 3.4: Implementation of temperature-dependent material behavior in PSPICE: Left) Plot of the look-up table data used to describe the temperature dependence of the resistivity ρ of stainless steel. The table's values stem from the [STEAM](#) material library [77]. Middle) Plot of the look-up table data used to describe the temperature dependence of the volumetric heat capacity per volume c_v of stainless steel. The values of this table were taken from the work of Bentz and Prasad [78] rather from the STEAM material library, for their increased accuracy in the high temperature range ($T > 300$ K). Right) Schematic corresponding to the PSPICE netlists that were implemented to model temperature-dependent material properties in PSPICE.

parameter, a separate node (T_{res1} and T_{res2}) is defined with a voltage corresponding to T_{res} . This is realized by the two voltage sources $E_{\text{Input_cv}}$ and $E_{\text{Input_}\rho}$. These nodes now act as inputs for two distinct subcircuits, cv_table and ρ_table . Within these subcircuits, another voltage source is defined which assigns a voltage across the subcomponent to the voltage at the first pin (having a value of T_{res} as previously defined) based on a look-up table. The actual material properties are stored in these look-up tables, whose values are shown in the graphs to the left of figure 3.4. To improve the readability of the model, and ease the access to ρ and c_v , the voltage across each subcomponent is mirrored with two distinct nodes (c_v and ρ), using the voltage-controlled voltage sources (PSPICE E-components) E_ρ and E_{cv} :

$$\begin{aligned} V(\text{cv}) &= V(d1) - V(T_{\text{res1}}) \\ V(\rho) &= V(d2) - V(T_{\text{res2}}) \end{aligned}$$

The voltage of the nodes (c_v and ρ) now corresponds to the temperature-dependent material properties as defined in the look-up tables and can be used just like any node voltage to define temperature-dependent material behavior.

3.1.3 Full equivalence model of an energy extraction unit

In the following, it is described how the formerly introduced subsystems are integrated into a complete model of an energy extraction unit. All the equations discussed in this section correspond to the definition of a PSPICE component or parameter, which can be found in the full system's equivalent schematic shown in figure 3.5. The overall system is made up of three EE modules connected in parallel. As PSPICE resistors (R components) do not allow dynamic changes in their resistance value, this is realized via voltage-controlled current sources (PSPICE G-components), whose internal resistance is defined by a variable voltage. However, since the parallel connection of three voltage-controlled current sources, one for each of the modules, would lead to numerical instabilities, another implementation has to be followed.

Modelling simultaneously the first two energy extraction modules that are not connected to the earthing system within a single subcircuit, showed to improve numerical stability and simulation time significantly. This implementation is discussed in the following.

In the schematic shown in figure 3.5, the nodes 1_pIn and 1_pOut are connected in parallel to the switch branch. Furthermore, two so-called "virtual" voltages sources ($E_{monitor_in}$ and $E_{monitor_out}$) are implemented, which can be used to measure the current in and out of the circuit for debugging purposes. The term "virtual" in this context means that the sources do not have any effect instead of a current measurement on the circuit, as the voltage across them is zero. Going from left-hand side to right-hand side, the monitoring sources are followed by parasitic capacitances and inductances $L_{parasitic}$ and $C_{parasitic}$, which have been inherited from the previous model and furthermore help the model's numerical stability. In the next branch, which contains the subcircuit, the actual calculation of the temperature-dependent resistance of the first two parallel modules takes place. Within the subcircuit, there is a single voltage-controlled current source whose internal resistance corresponds to that of two parallel modules ($R_{mod1||mod2}$), i.e. half a module resistance:

$$I(G_{ABM1}) = \frac{V(G_{ABM1})}{V(R_{mod1||mod2})}. \quad (3.6)$$

In order to ensure that the initial resistance at $T = T_{amb}$ is equal to an initial measured resistance R_{meas} (across the entire EE unit), the module's resistance is scaled with a factor f_{res} , which can be obtained from measurement data and is defined as:

$$f_{res} = \frac{3 \cdot R_{meas} \cdot A_{res}}{l_{res} \cdot \rho(T_{amb})}.$$

Hence, the internal resistance of G_{ABM1} is defined as:

$$R_{mod1||mod2} = \frac{1}{2} \cdot f_{res} \cdot \frac{l_{res}}{A_{res}} * V(\rho),$$

with $l_{res} = 84.3.12$ m and $A_{res} = 0.22 \text{ m} \cdot 0.004 \text{ m} = 88e^{-5} \text{ m}^2$ being the resistor equivalent length and cross-section and $V(\rho)$ being the temperature-dependent resistivity obtained as described in section 3.1.2.

Since two modules in parallel are modeled, and the current through them is hence twice as high as for a single module, the system equations of the thermodynamic submodel have to

be adapted accordingly to match the mass and cooling power of two parallel energy extraction modules:

$$I(G_{\text{ohmic}}) = I(G_{\text{ABM1}})^2 \cdot V(R_{\text{mod1} \parallel \text{mod2}})$$

$$I(G_{\text{exchange}}) = -2 \cdot h_{\text{th_EE}} \cdot (T_{\text{res}} - T_{\text{amb}})$$

$$V(E_{C_v}) = -\frac{2 \cdot f_{C_v_EE}}{V(C_v)} \cdot \int I(E_{C_v}) dt$$

with $f_{C_v_EE}$ being a scaling factor of the thermal capacitance, accounting for imperfections in the heat capacitance's look-up data and $V(C_v) = V(c_v) \cdot A_{\text{res}} \cdot l_{\text{res}}$ being the voltage corresponding to a single module's heat capacity at constant volume. It can be noted that despite modelling two parallel EE modules within the subcircuit, the freely selectable fitting parameters $h_{\text{th_EE}}$ and $f_{C_v_EE}$ still correspond to the thermal conductance and heat capacity scaling factor of a single module, which eases their interpretation. Additionally, the way in which the temperature-dependent properties ρ and c_v are retrieved remains unchanged as introduced in section 3.1.2. Since the subcircuit models only two out of three energy extraction modules, the third one remains to be considered. The third module can not be modelled within the subcircuit as well, as it is connected to the LHC's earthing system. Hence, the third module is modelled by two voltage-controlled voltage sources (PSPICE E-components) which are connected in series ($E_{\text{mod3_1st_half}}$ and $E_{\text{mod3_2nd_half}}$). The voltage sources each have an internal resistance of $R_{\text{mod1} \parallel \text{mod2}}$, corresponding to the resistance of two parallel modules which is equivalent to the resistance of half a module:

$$\begin{aligned} V(E_{\text{mod3_1st_half}}) &= R_{\text{mod1} \parallel \text{mod2}} \cdot I(E_{\text{mod3_1st_half}}) \\ V(E_{\text{mod3_2nd_half}}) &= R_{\text{mod1} \parallel \text{mod2}} \cdot I(E_{\text{mod3_2nd_half}}) \end{aligned}$$

Hence, their total series resistance corresponds to that of a single module. However, as parent circuits cannot access the nodes of their sub-circuits in PSPICE, the value of $R_{\text{mod1} \parallel \text{mod2}}$ is retrieved the following way:

$$R_{\text{mod1} \parallel \text{mod2}} = \frac{V(1_pOut, 1_pIn)}{I(E_{I_mod1 \parallel \text{mod2}})},$$

with $V(1_pOut, 1_pIn)$ being the voltage across the subcircuit and $I(E_{I_mod1 \parallel \text{mod2}})$ being the current through the virtual voltage source $E_{I_mod1 \parallel \text{mod2}}$ which is inline with the subcircuit. The model's optional connection to the earthing system, which is only implemented in case of one of the two energy extraction units, is considered by a node in between the two PSPICE E-components ($E_{\text{mod3_1st_half}}$ and $E_{\text{mod3_2nd_half}}$) that either floats or is connected to the earthing system.

3.2 Setup of the optimization

3.2.1 Selection of suitable events

In order to fit the new EE resistance model to measurement data, suitable FPA events had to be selected. Hence, three different FPA events, occurring at different current levels were selected,

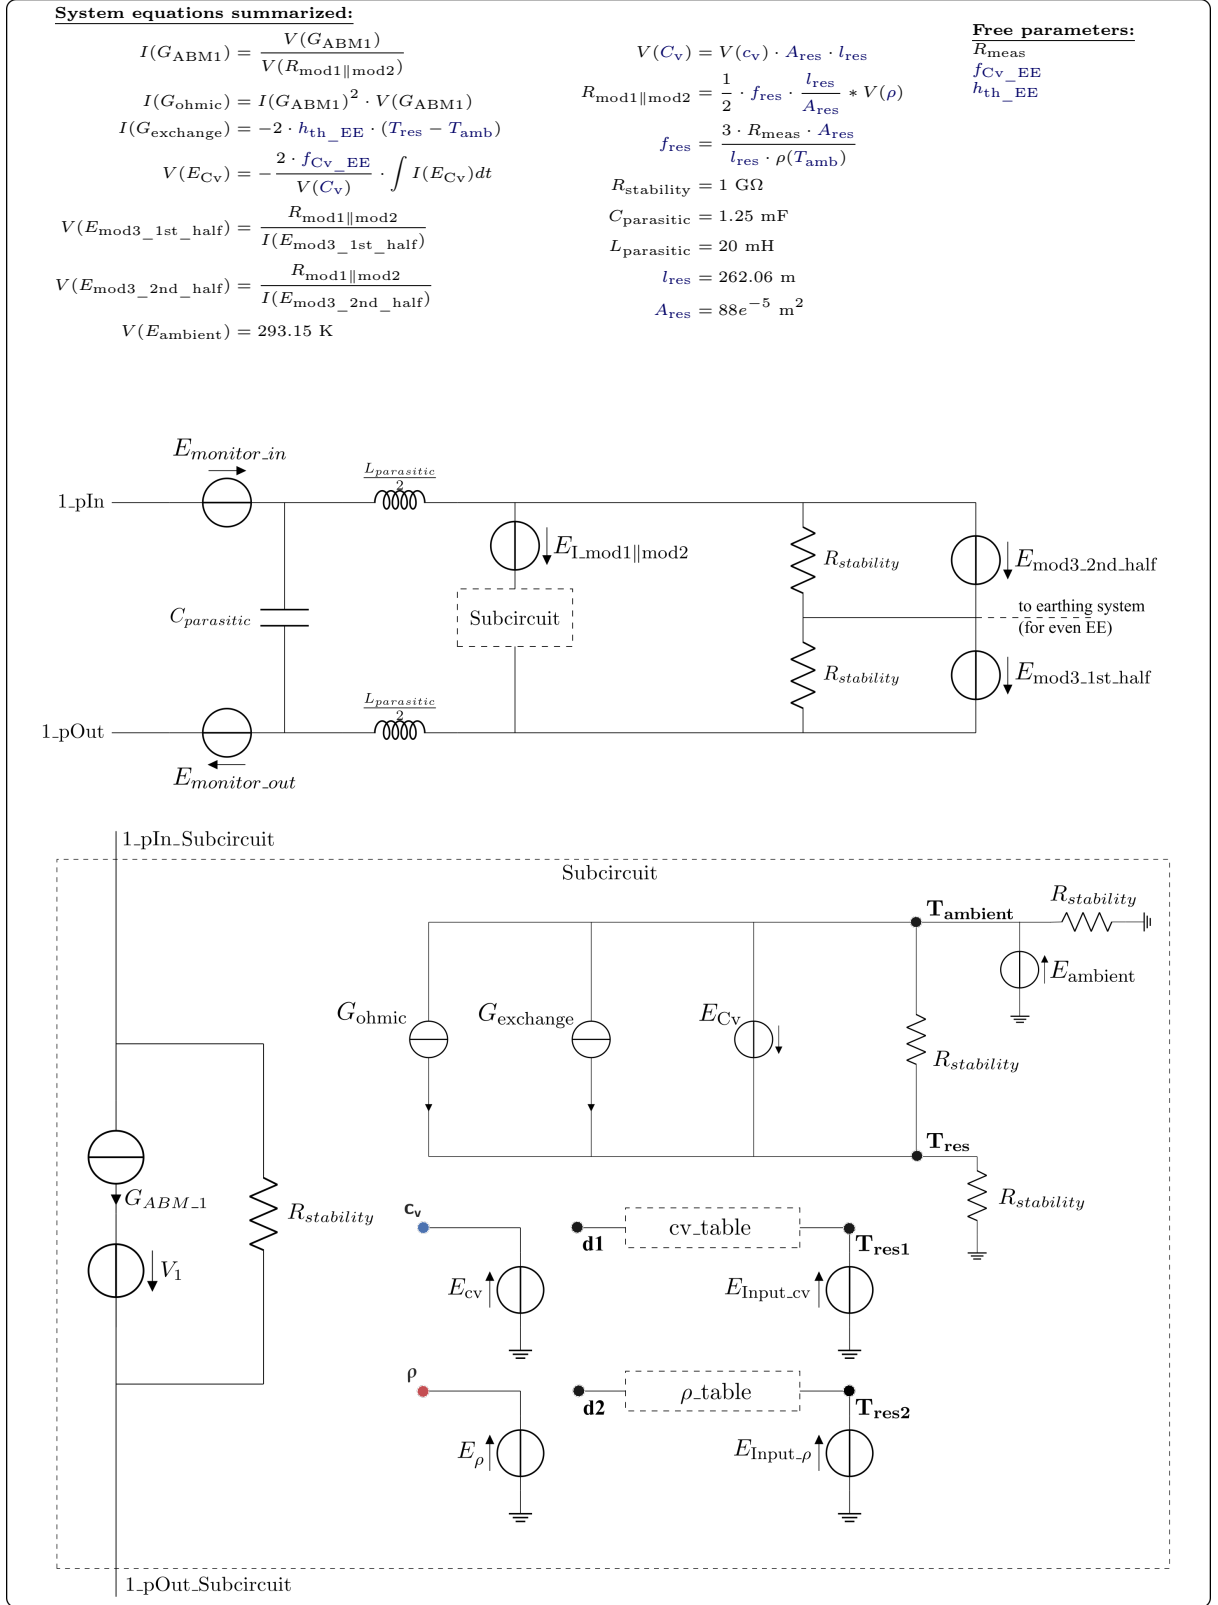


Figure 3.5: Equivalent schematic of a full RB EE unit's PSPICE model without its switch branch.

which had to fulfill the following criteria:

- No quench of the magnets: By excluding events with a quench, the influence of incorrect modelling of the resistance of a quenching magnet as the cause of deviations in the simulated circuit current can be excluded.
- Timestamp of the event must be after 2015: As explained at the beginning of the chapter, the energy extraction featured a modified module configuration before 2015, which does not match that of the model implemented. Accordingly, the [FPA](#) events selected for the optimization must have occurred after 2015.
- Consistency of the eventfile: Over the course of the project, it became clear that some eventfiles had incomplete or inconsistent headers. However, as this header specifies the boundary conditions of the simulation, it had to be ensured that all the necessary keys were assigned and that they were within a valid value range.
- Current level: The current level at the start of the discharge should be distributed as evenly as possible across the selected event files so that the validity of the model parameters for different [FPA](#) current levels is ensured.
- Availability of measurement data: Since the optimization process is based on iteratively comparing simulation data to measurement data, the measurement data of the selected events had to be available in [PM](#) and [NXCALIS](#).

Since checking the aforementioned criteria by hand would be impractical, a python script filtering the eventfiles was implemented. This script was used to search thousands of events, provided by the signal monitoring team (SIGMON) at CERN [65]. In the end three different events without a magnet quench were found which occurred at [FPA](#) currents ranging from $I_0 = 3499$ A to $I_0 = 11500$ A.

3.2.2 Selection of a suitable optimization algorithm

Initially, gradient descent methods were considered to be the fastest way to reach convergence to the optimum point. However, the process of prototyping physical models posed an additional challenge of handling parameter configurations which prevent convergence of the simulation. Hence, the simulation driver was complemented with a timeout parameter which would abort any simulation process not converging within 900 s. In order to mitigate these cases of failed or frozen simulations, and to provide the optimization algorithm with feedback about the failure of the analysis, the analysis driver was further supplemented to return a high penalty value for the objective function in such cases. Since this strategy nevertheless causes wrong numerical gradients, which would prevent a gradient descent method from its correct function, the ultimate choice was a derivative free pattern search algorithm, which has already been briefly introduced in section 2.3.2.2. Prior to optimization, the variables were automatically normalized with respect to their maximum and minimum permitted values in order to compensate for differences in sensitivity. In the [STEAM-DI](#) this can be done by setting the keyword `scale_type = "auto"`. In addition, the initial step size is an important parameter, significantly influencing the convergence of the optimization. Here, an initial step size of $\pm 10\%$ of the start value (keyword: `initial_delta = 0.1`), proved to be effective.

3.3 Results

In the following two sections 3.3.1 and 3.3.2, the **Old energy extraction Model (OM)** is compared to the model developed within this thesis, which is in the following referred to as the **New energy extraction Model (NM)**. As introduced in the previous sections, the **NM** takes into account the heating and cooling of the resistor. Later on within the following sections 3.3.3 and 3.3.4, two different fits for the free parameters introduced in section 3.1.3 are presented. These two parameter configurations fitted, one fitted with constant magnet inductance $L_{\text{magnet}} = \text{const.}$ and one with current-dependent magnet inductance $L_{\text{magnet}} = f(I)$ are presented in order to highlight the significance of the improved inductance description implemented. The models after fitting are in the following referred to as **New energy extraction Model after Fit (NMF)** and **New energy extraction Model after Fit considering variable Inductance (NMFI)**. In order to draw a comparison between the different models and optimizations performed, a set of signals derived with different models was compared to measurement data. These signals include the voltage across the energy extraction V_{EE} , the current through the circuit I_{MEAS} , as well as the resistance of the **EE**. Since the resistance of the **EE** units is not directly measured, it was calculated from **PM** signals as follows: $R_{\text{EE}} = \frac{V_{\text{EE}}}{I_{\text{MEAS}}}$. Furthermore, the increase in temperature of the **EE** modules was compared to measurement data. To do so, measurement data from both databases had to be considered in order to profit from both the high resolution in the time domain provided by **PM** and the increased time span of the logging data provided by **NXCALS**. Within both databases, the temperatures of the individual modules are saved individually. In order to make these signals comparable to the simulated temperature within the **NM**, in which all **EE** modules are assumed to have the same resistance and temperature, the data had to be preprocessed. Within this preprocessing step, the mean value of the three individual energy extraction module temperatures was calculated and saved within a single dataframe enabling a more direct comparison. Furthermore, only the increase in temperature ΔT , but not the absolute temperature of the **EE** modules was compared, as the initial module temperatures may vary due to prior **FPA** events, which is not yet considered within starting values of the simulation. Another characteristic temperature which was considered as part of the comparison is the hereinafter so-called experimentally-derived temperature ΔT_{calc} . This quantity is the temperature increase which would be necessary to explain the measured increase in resistance, under the assumption of uniform temperature distribution in the **EE** resistor. It is derived by mapping the increase in resistance to an increase in resistivity, which can then be translated into an increase in temperature through the model's material look-up data. For each different model, three different **FPA** currents I_0 : 11500 A, 5923 A and 3499 A were tested.

In most cases only the results corresponding to 11500 A are presented in this chapter, as the conclusions and characteristics across the various current levels are very similar for all I_0 , but most often only differ in how much they are pronounced. To simplify the comparison of the models, the normalized error ε between the measured values and the simulation is presented. This error is calculated as follows:

$$\varepsilon = \frac{Q_{\text{sim}} - Q_{\text{meas}}}{\max(\text{abs}(Q_{\text{meas}}))} \quad (3.7)$$

with Q_{sim} being the simulated quantity, Q_{meas} being the measured quantity and $\max(\text{abs}(Q_{\text{meas}}))$ being the maximum absolute value of the measured quantity. In order to provide a clearer

picture across all models, different fits, current levels and signals, the mean value of this error is summarized within table 3.1 of the summary. At last, a meta analysis of the fitting processes, showing the strengths and weaknesses of the applied algorithm and parameters is presented in section 3.3.5.

3.3.1 Old energy extraction model (OM)

3.3.1.1 Characteristic electrical signals

In the following, the OM's characteristic signals and their respective errors are presented, based on figure 3.6. Both the simulated and the measured signal of V_{EE} , shown in figure 3.6a, exhibit a sharp drop immediately after the FPA, which is observed for both EE units. This voltage drop is associated with the opening of the switch, forcing the circuit current I_{MEAS} to flow through the energy extraction. Despite the qualitatively similar course of both simulated and measured voltage signals, a considerable deviation between the signals can be noticed. This deviation, which is illustrated by the normalized error ε , steadily decreases after the FPA, until reaching its minimum after at approximately 45 s. The maximum absolute value of this error is above the 5% mark. Later on, however, the simulated and measured voltage curves equalize again, which is also reflected in the decreasing absolute of the error ε . This deviation is significantly influenced by the model's underlying assumption of a constant EE resistance. Since the resistor heats up over the course of a discharge, its resistance increases, which is not considered within the model. Hence, the model underestimates the voltage drop across the EE resistance, which is the cause of the negative errors observed. The observed positive error after approximately 150 s however can not be attributed to an underestimated resistance, but may be caused by deviations between the modelled and real magnet inductance. A similar conclusion can be drawn for the course of the current in the circuit I_{MEAS} , which is shown in figure 3.6b. Both the simulated and the measured current follow the expected semi-exponential drop. While they are still in good agreement immediately after the FPA, an increasing deviation at later stages of the discharge is observed. Considering the errors, this deviation increases steadily until approximately 140 second after the FPA, and subsequently stagnates. The maximum value of the error is at approximately 4%. The overestimated current can be attributed to an underestimated EE resistance. Finally, the course of the simulated and measured resistance is shown in the figure 3.6c. A clear increase in the measured electrical resistance can be observed. Furthermore, after approximately 150 s the odd energy extraction's resistance exhibits a small decrease, whereas the resistance of the even energy extraction further increases. This decrease in the measured resistance, observed for the odd EE in later stages of the discharge (after approximately 200 s), may be attributed to the module's cooling. The simulated resistance however is treated as constant within the OM and hence significantly deviates from the measured value.

3.3.2 New energy extraction model (NM)

In the following section, the simulation results of the characteristic signals and temperatures as calculated with the NM are presented. In order to ease interpretation and validation, the simulations presented within this section were conducted considering the module to be adiabatic $h_{th_EE} = 0$ and unchanged with respect to its heat capacity $f_{Cv_EE} = 1$.

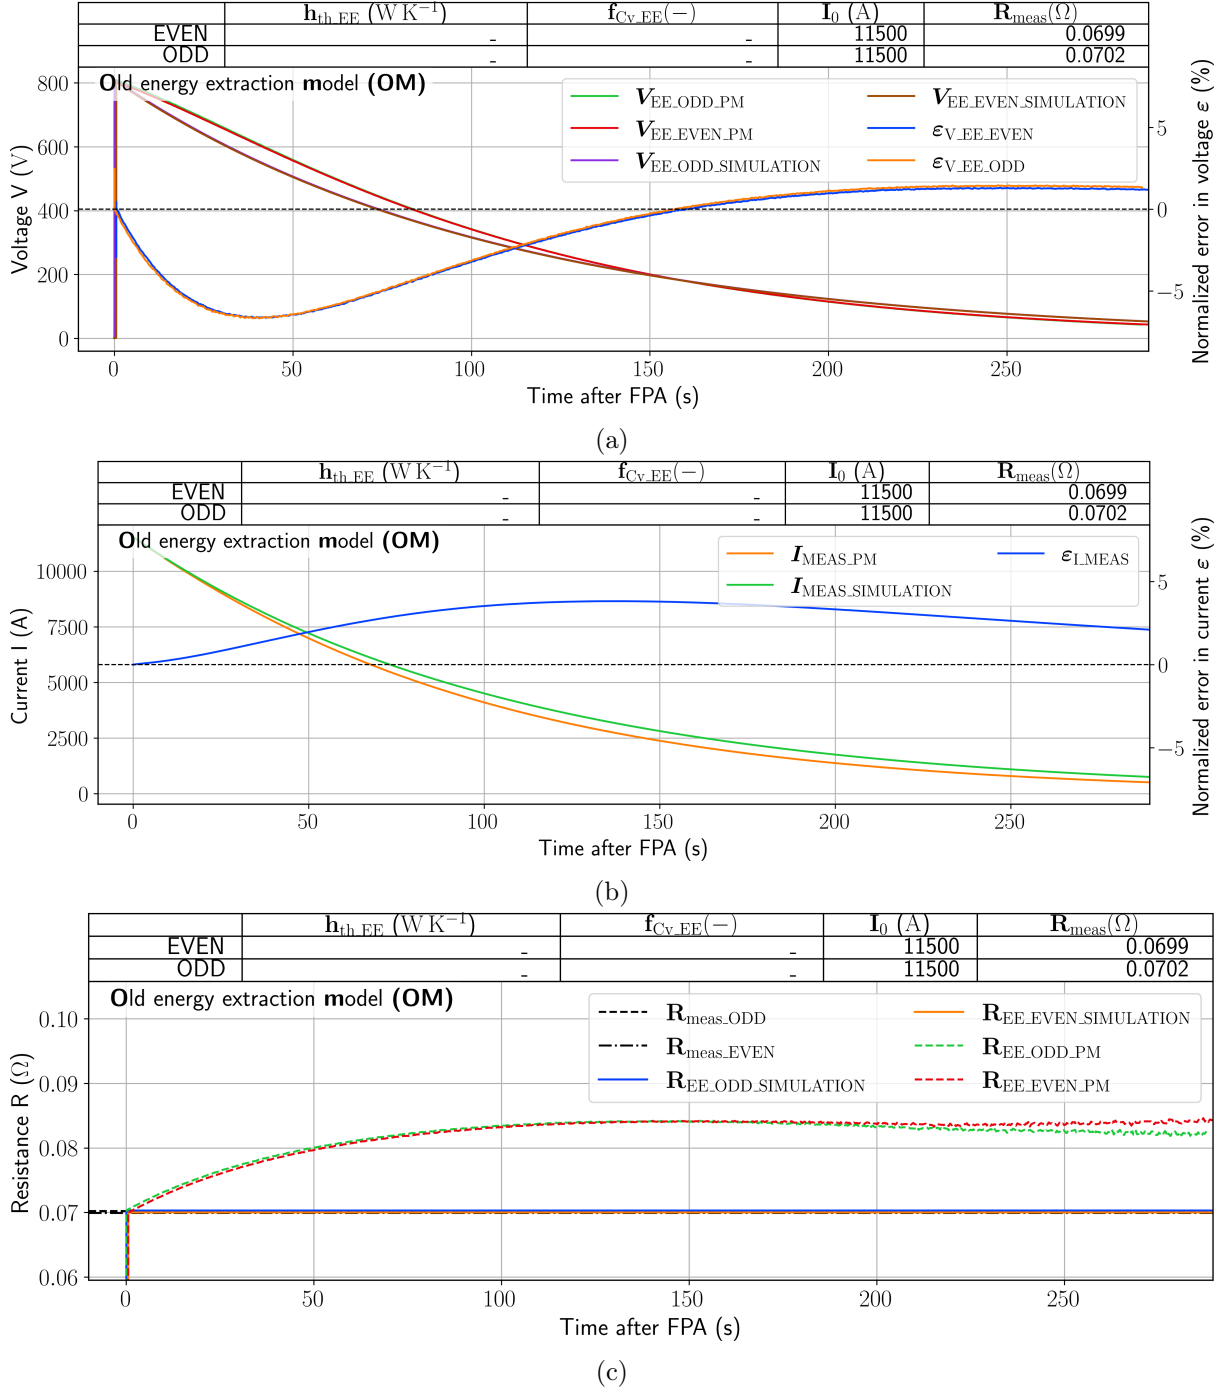


Figure 3.6: Characteristic electrical signals obtained with the OM for a discharge (without quench) at $I_0 = 11500$ A. a) Voltage across the energy extraction V_{EE} , b) Current through the circuit I_{MEAS} , c) electrical resistance of the energy extraction R_{EE} . The OM does not feature a thermodynamic description, hence it does not employ the free parameters h_{th_EE} and f_{Cv_EE} which is denoted by a "-" within the table above the respective plots.

3.3.2.1 Characteristic electrical signals

In the following, the NM's characteristic signals shown in figure 3.7 are presented. For the voltage across the energy extraction V_{EE} , shown in figure 3.7a, only a slight deviation between the measured and simulated signal can be noticed. This deviation steadily decreases over the course of the discharge, until reaching a minimum after approximately 50 s. The maximum absolute error in V_{EE} is below 2%. Just like for the OM, the simulated and measured voltage curves equalize in the further course, which is also reflected in a decreasing error, which almost drops to zero towards the end. This improvement in the match between measured and simulated signals of V_{EE} can be attributed to the model considering the temperature related increase in resistance of the EE. For the NM, the simulated and measured current almost coincide as shown in figure 3.7b. With the NM only a slight error in I_{MEAS} below 0.2% can be observed, having its maximum at approximately 25 s after the FPA. However, after this maximum is reached, the curves steadily equalize. Conclusively, also I_{MEAS} exhibits an improved accuracy in comparison to the OM. Finally, the course of the resistance as simulated with the NM is shown in figure 3.7c. Both the modelled resistance, as well as the measured resistance, increase steadily over the course of the discharge, proving a qualitatively consistent implementation of the model. Furthermore, the characteristic feature of the even energy extraction exhibiting a lower resistance than the odd energy extraction is correctly represented within the NM. However, there is still a clear deviation between the two resistances, as the measured resistance is larger than the simulated resistance over the entire course of the discharge. At last, the courses of measurement and simulation equalize again towards later stages of the discharge.

3.3.2.2 Simulated and measured temperatures

In the following section, the module temperature as calculated through the NM is compared to measurement data for all current levels tested. Figure 3.8 shows the simulated and the measured increase in temperature. At first, the temperature increase for the lowest tested FPA current $I_0 = 3499$ A, which is illustrated in figure 3.8a will be discussed. Considering the course of ΔT , it is noticeable that for both EE units the simulated rise in temperature is significantly faster than the measured temperature increase. Both curves however reach a similar steady state temperature increase of about 17 K. The experimentally-derived temperature increase ΔT_{calc} however does not exhibit a steady state. On the contrary, it exhibits a steeper rise in the first 100 s than both the simulation and the measurement. In the course of the rise, both curves approach a local maximum of about $\Delta T_{calc} = 25$ K, which is reached after approximately 50 s. Subsequently, a temporary decline in the curve can be observed for both energy extraction units, lasting until about 130 s. Afterwards, both curves of ΔT_{calc} exhibit a second, much steeper rise in temperature, which increases up to values several times higher than what is observed in the measurements. This strong rise in ΔT_{calc} at late stages of the transient is an artifact, the cause of which will be discussed further at the end of this section.

Comparable observations can be made for the temperature increase analyzed at $I_0 = 5923$ A, which is illustrated in figure 3.8b. Here too, the initial rise in the simulated temperature increase is larger than that of the measurement. In addition, just as in the previous case, an overshoot of the experimentally-derived temperature ΔT_{calc} increase can be observed, which, in contrast to the previous current level discussed, seems to be approaching a steady state. Remarkably, the final temperature increases of measurement, simulation and ΔT_{calc} seem to approach each

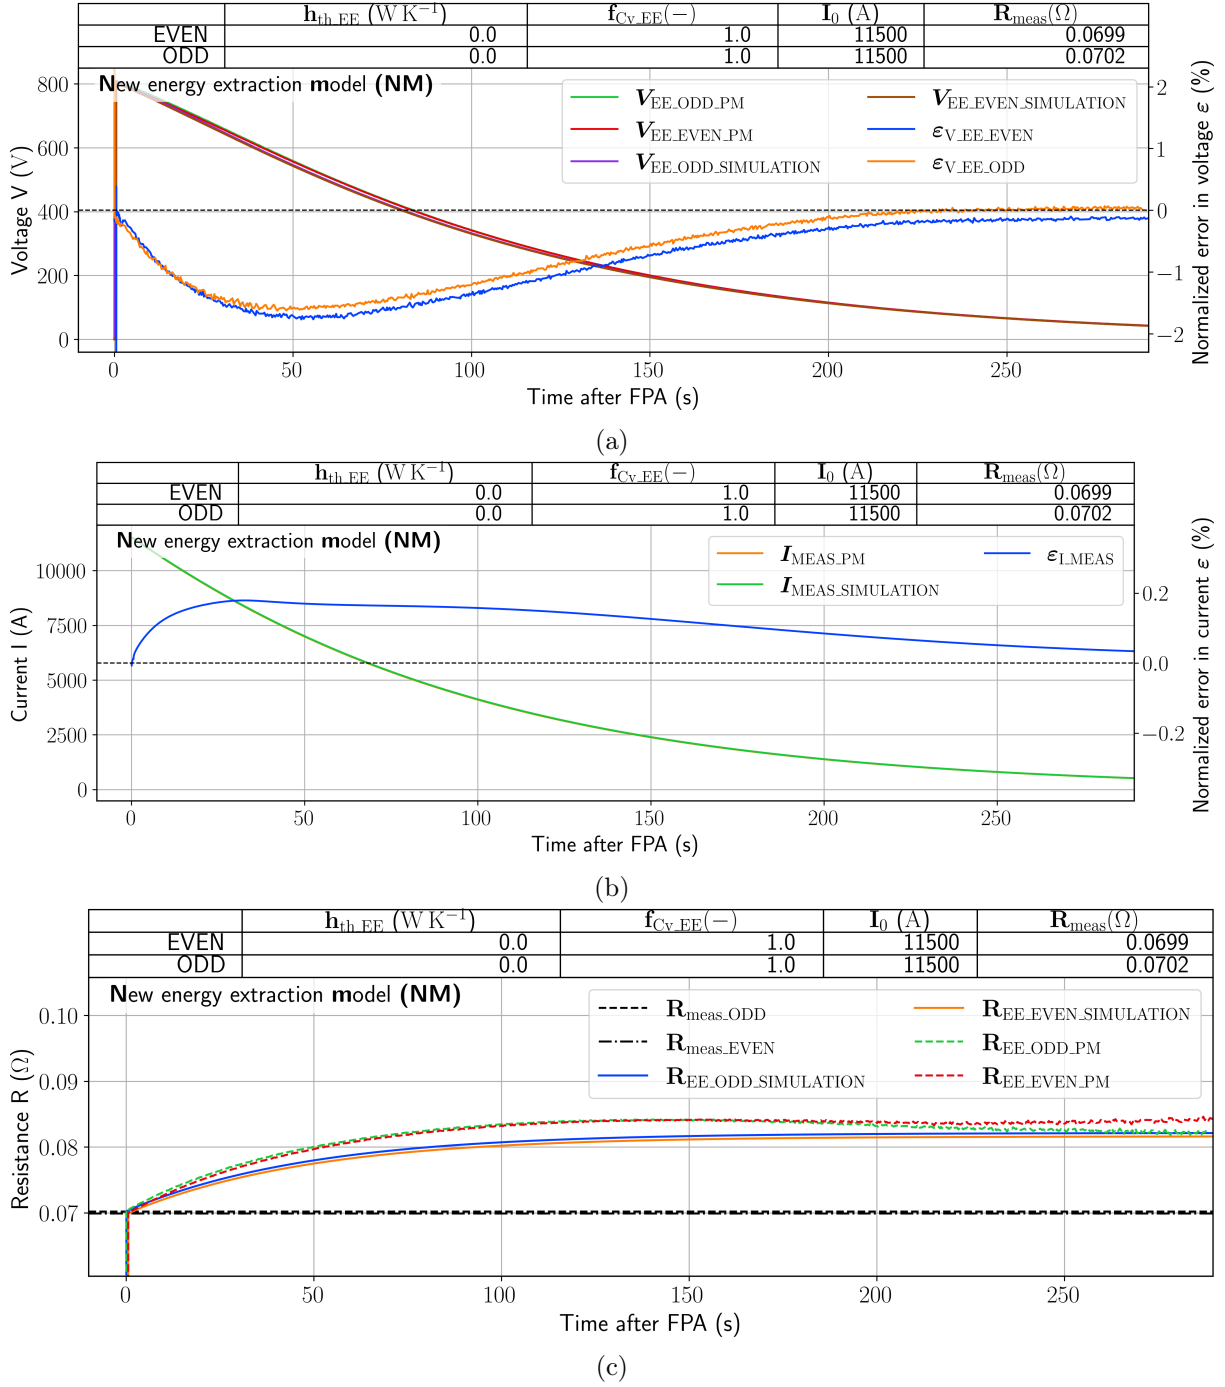


Figure 3.7: Characteristic electrical signals obtained with the NM for a discharge (without quench) at $I_0 = 11500$ A. a) Voltage across the energy extraction V_{EE} , b) Current through the circuit I_{MEAS} , c) Electrical resistance of the energy extraction R_{EE} . In order to ease the comparison, cooling was turned of by setting $h_{th_EE} = 0$. Further characteristic model parameters are given in the tables above the respective plots.

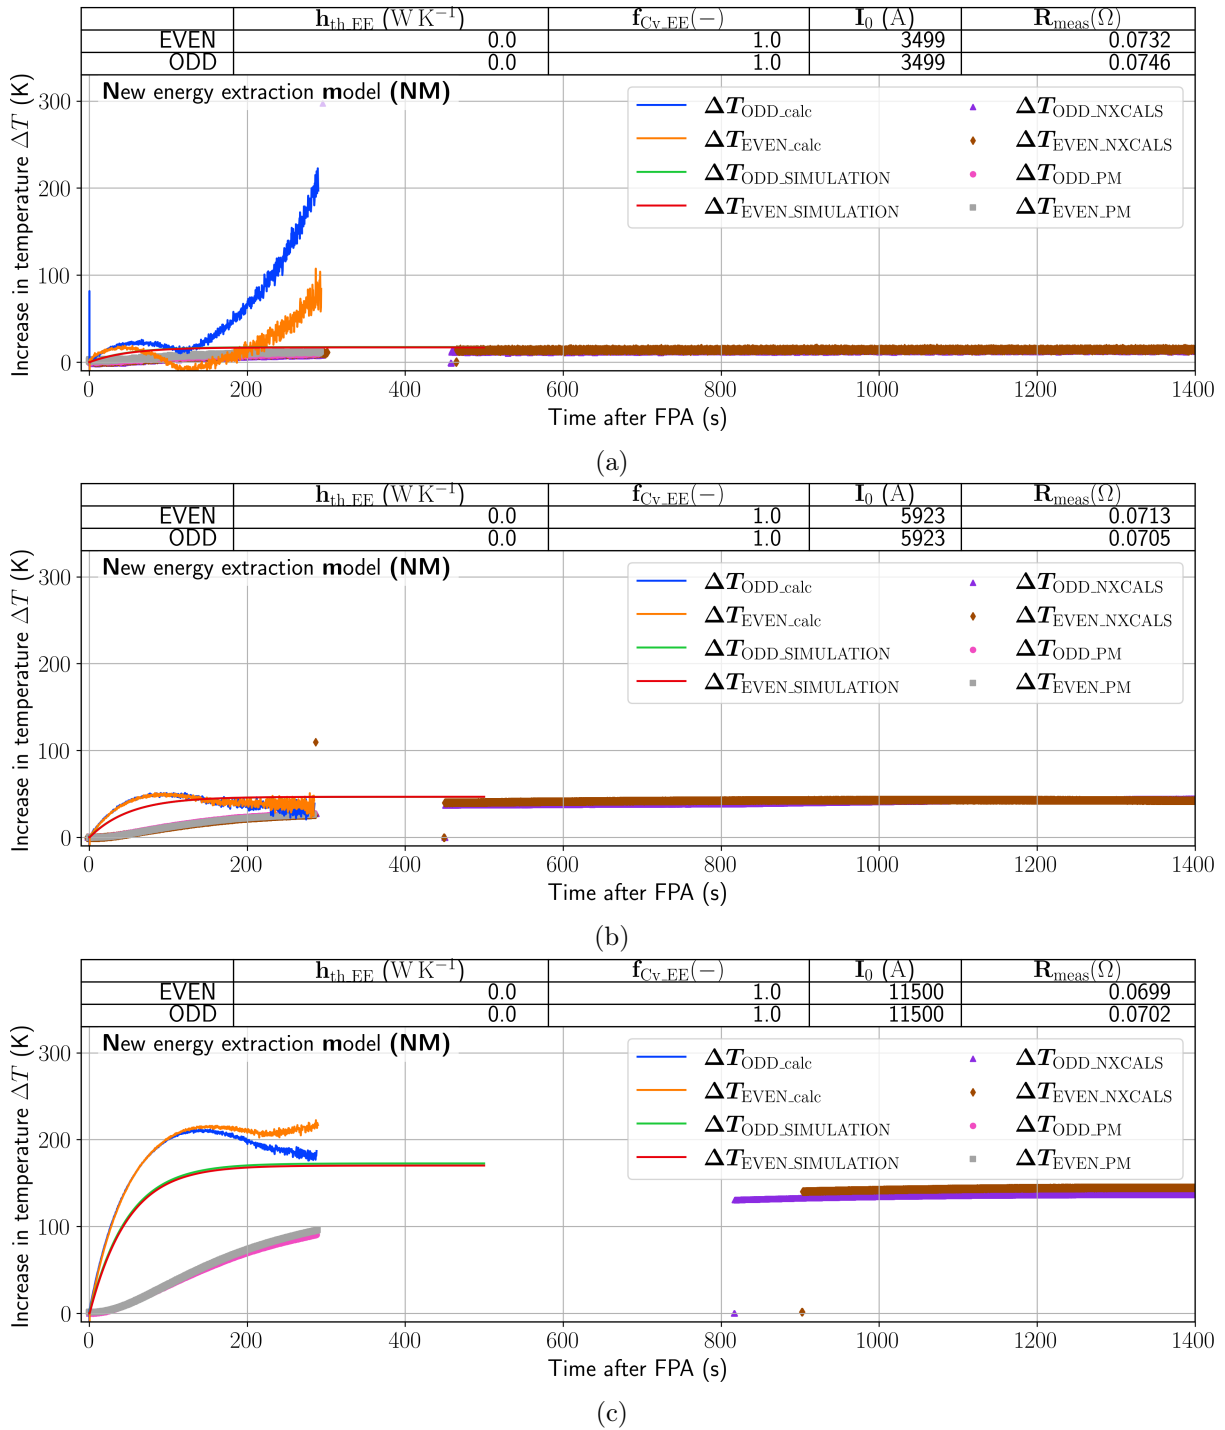


Figure 3.8: Increase in resistor temperature as simulated with the NM in comparison to measurement data and the temperature increase which would be necessary to explain the measured increase in resistance ΔT_{calc} (assuming homogeneous resistivity, heat capacity and temperature across the resistor). The plots correspond to different FPA events, which occurred at different current levels I_0 .

other at approximately 45 ± 5 K. This steady state temperature increase is higher than that observed for $I_0 = 3499$ A.

Finally, the temperature increase at $I_0 = 11500$ A as shown in figure 3.8c will be discussed. Considering the course of ΔT , it is noticeable that just like for the other two current levels, the simulated rise in temperature is faster than the measured rise in temperature. Furthermore, the simulated maximum of ΔT , is about 20 K above the measured maximum of ΔT . While the simulated temperature increase reaches a steady state of approximately 165 K, the measured temperature increases no higher than about 150 K. Again, the experimentally-derived temperature increase exhibits a steeper rise than both simulation and measurement. In the course of its rise, the ΔT_{calc} curves of both systems approach a local maximum of approximately $\Delta T_{\text{calc}} = 220$ K, which is reached after about 150 s. Subsequently, a temporary decline in the ΔT_{calc} curve of both systems can be observed, which for the even energy extraction, however is followed by a second rise after approximately 200 s.

In contrast to lower current levels, the final increase in temperature is not matched for $I_0 = 11500$ A. The cause of this observation could be rooted in the modules' cooling control system. The cooling fans of the EE modules are only triggered, once a certain threshold temperature (approximately 50 degree celsius [43]) is surpassed. For all temperatures below this threshold, the simulation's underlying assumption of $h_{\text{th_EE}} = 0$ has more validity since also in reality, the cooling is not triggered, but turned off, resulting in a better match of the steady state temperature increases of measurement and simulation. For higher temperatures however (caused by higher I_0), the cooling system is triggered which may be the cause for which the simulated steady state temperature increase of the adiabatic system exceeds the measured increase in temperature for this case. The fact that the simulated and measured temperatures approach each other for lower currents however signifies that the thermodynamic equilibrium considered within the NM enables an accurate approximation of the module temperatures. Furthermore, an overshooting of ΔT_{calc} could be observed for all tested current levels. This phenomenon of overshooting may be caused by non-uniform heating within the resistor (for example at the welding points of the stainless steel plates). As areas of increased temperature would primarily occur at points with a high current density, they could act like additional serial resistances having a non-negligible influence on the overall resistance of the module. After enough time however, thermal diffusion would cause the temperature gradients within the steel plates to reduce, causing a reduction of the module resistance. This condition where all temperature gradients are evened out by thermal diffusion would be reached when ΔT_{calc} realigns with the simulated increase in temperature, as observed in figure 3.8b. Since such effects are not covered within the model, which assumes the resistor to be a homogeneous mass of uniform resistivity, temperature and heat capacity, the simulated increase in temperature deviates from ΔT_{calc} , especially within early stages of the simulation. The strong rise of the increase of ΔT_{calc} at late stages of the transient is assumed to be an artifact of measurement noise. ΔT_{calc} is approximately proportional to the measured resistance, which is proportional to the inverse of I_{MEAS} . Hence, a vanishing circuit current, together with noise in the measurement within late stages of the discharge could be a possible cause for the unrealistically large values of ΔT_{calc} .

This explanation is further supported by the fact that the large increase in ΔT_{calc} is most pronounced for the event with the lowest considered FPA current of 3499 A. However, since a similar behavior is also observed for $I_0 = 11500$ A, the cause of the second rise observed in ΔT_{calc} can not conclusively be explained.

Another open question is why the measured increase in temperature seems to lag behind the simulated increase in temperature. Since the steady state temperature increase for lower current levels is matched, an underestimated influence of cooling can be excluded as a cause of the observed delay between both signals. One possible reason could be the position of the temperature sensor, as it is not in direct contact with the plates, but located inside a capsule attached to a plate in the middle of the module. Hence, an increased shell thickness of the capsule, and the associated slower thermal diffusion, could explain the delayed reaction of the sensor.

3.3.3 New energy extraction model after fit (NMF)

In the following section, the simulation results of the characteristic signals and temperatures as calculated with the **NMF** are presented¹.

3.3.3.1 Characteristic electrical signals

In the following, the same set of signals already presented in the previous sections but calculated with the **NMF** is presented. After fitting the model with Dakota, a clear improvement in the simulation accuracy of V_{EE} is observable, as can be seen from figure 3.9a. In contrast to the **NM**, the optimized model's error in V_{EE} is less than one percent. However, considering the course of I_{MEAS} in figure 3.9b, it can be seen that the maximum absolute error in current is increased in comparison to the **NM** before the fit. While the **NM** without fitting exhibited errors below 0.2% magnitude in I_{MEAS} , the **NMF**'s error in I_{MEAS} almost approaches 1% in magnitude. This maximum absolute error is reached about 100 - 150 s after the **FPA**. Considering the simulated and measured resistances in figure 3.9c, it is noticeable that the simulated signals almost coincide with the measured values, especially in the first 70 s. On the contrary, after approximately 70 s, the simulated increase in resistance is slightly higher than the measured increase in resistance, resulting in a slight deviation between the two curves. Towards the end of the simulation, the simulated resistance remains at a constant level, while the measured resistance appears to slightly reduce. To conclude, it shall be highlighted that the fit has resulted in an improved match for voltage and resistance signals, but also in a slight deterioration of the current signal in comparison to the **NM** before fitting. Compared to the **OM** however, all simulated signals (also the current) still exhibit a higher accuracy when being simulated with the **NMF**.

Considering the optimum values of the fitting parameters, it is noticeable that the simulated influence of cooling is negligibly small, as signified by a very low thermal conductance $h_{th_EE} = 1.0113 \text{ W K}^{-1}$. Furthermore, it is worth noting that the calculated heat capacity must be scaled down to approximately 76%, in order to reproduce the measured signals.

3.3.4 New energy extraction model after fit considering variable inductance (NMFI)

As stated in the previous section, fitting the model's free parameters f_{Cv_EE} and h_{th_EE} improved the accuracy of the simulation in terms of voltage and resistance. On the contrary, the

¹In terms of the conclusions that can be drawn, the simulated temperature increase of the **NMF** is similar to the **NMFI**'s progression of ΔT which is presented in the following section. Hence, the **NMF**'s temperature signals are only presented in figure A.1 of the appendix, but are not discussed in more detail.

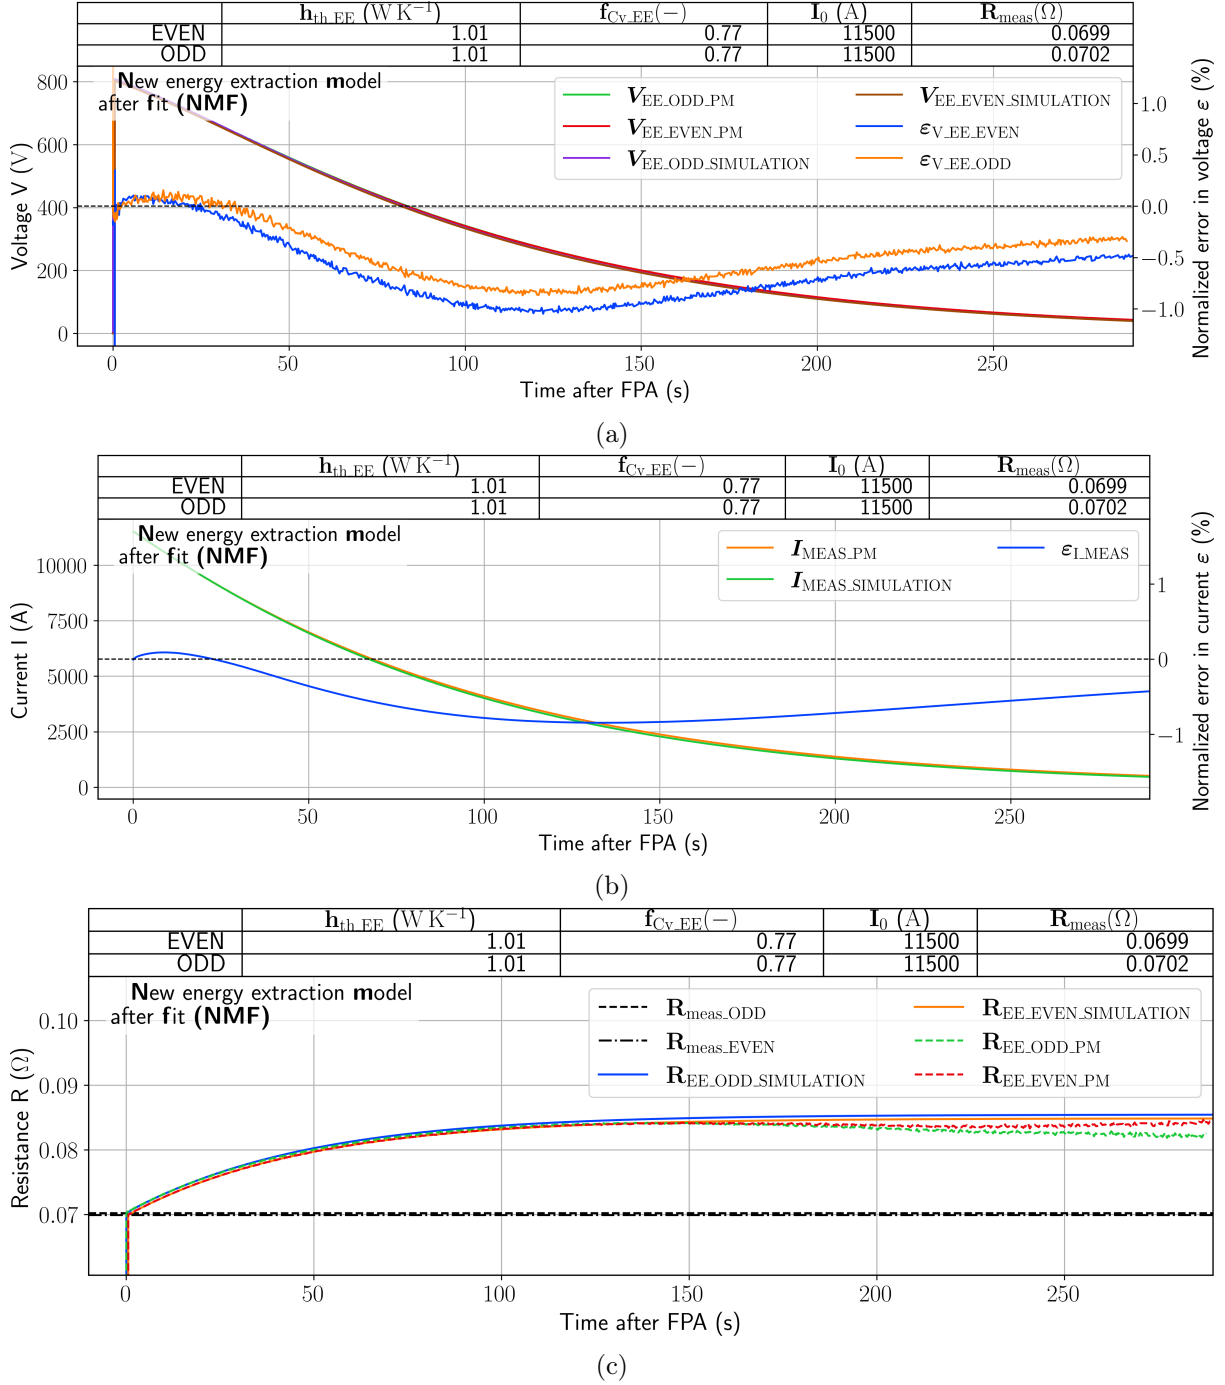


Figure 3.9: Characteristic electrical signals obtained with the NMF for a discharge (without quench) at $I_0 = 11500$ A. a) Voltage across the energy extraction V_{EE} , b) Current through the circuit I_{MEAS} , c) Electrical resistance of the energy extraction R_{EE} . Further characteristic model parameters are given in the tables above the respective plots.

simulation's accuracy in predicting the circuit current I_{MEAS} reduced in the fitting process.

The circuit current during the discharge is largely determined by the total inductance and the total resistance of the circuit. The latter is mainly dominated by the EE resistance. Ultimately, an improved accuracy of the resistance should result in an improved accuracy of the EE current. However, since the error in I_{MEAS} increased due to the improvement of V_{EE} and R_{EE} , a major deviation between the simulated and the actual inductance of the circuit was the most likely cause. Hence, an improved magnet model was implemented, which is introduced within the next section.

3.3.4.1 Improvements to the MB magnet model

In the previous model of the main dipole circuit's magnets (MB), the inductance was assumed to have a constant value of 0.098 H, which is a good approximation. However, the inductance actually changes slightly depending on the current, which caused a deviation in simulated and measured inductance, as shown in figure 3.10a. This can be attributed to the saturation of the iron yoke surrounding the magnet coils. In order to consider this change in inductance during the discharge, the MB magnet model was supplemented with a current-dependent inductance formulation. This was done in the same way as described in section 3.1.2 by defining a current-dependent inductance look-up table. As the magnet differential inductance during a FPA is not directly measured, it had to be calculated through the voltage across the magnet U_{DIODE} (corresponding to the diode voltage signal saved in the PM database) and the current through the magnet (which is equal to I_{MEAS} as long as the magnet does not quench) like exemplified in the following equation:

$$L_{\text{magnet}} = \frac{U_{\text{magnet}}}{\frac{d}{dt} I_{\text{magnet}}} = \frac{U_{\text{DIODE}}}{\frac{d}{dt} I_{\text{MEAS}}}. \quad (3.8)$$

Based on this data, a look-up table of the magnet differential inductance was defined that increases with decreasing current². Within the new magnet model, each magnet inductance component was replaced by a current-dependent look-up table based voltage-controlled voltage source (PSPICE E-component). This resulted in a more accurate match of the measured and the simulated inductance across a widespread range of different current levels, as shown in figure 3.10b.

3.3.4.2 Characteristic electrical signals

After the simulation accuracy of the circuit current during the discharge was improved by supplementing the magnet model as described above, the model of the energy extraction was fitted once more. As this model is the finally implemented and deployed version of the EE , the characteristic signals for all tested FPA current levels are presented to illustrate the improved accuracy across a wide range of different events.

Voltage across the energy extraction V_{EE}

In the following, the voltages across the energy extraction as obtained with the final version of the model (NMFI) are discussed, based on figure 3.11. Just like for the previous fit of the model

²In order to obtain improved numerical stability, the look-up table was defined to be monotonically increasing with decreasing current.

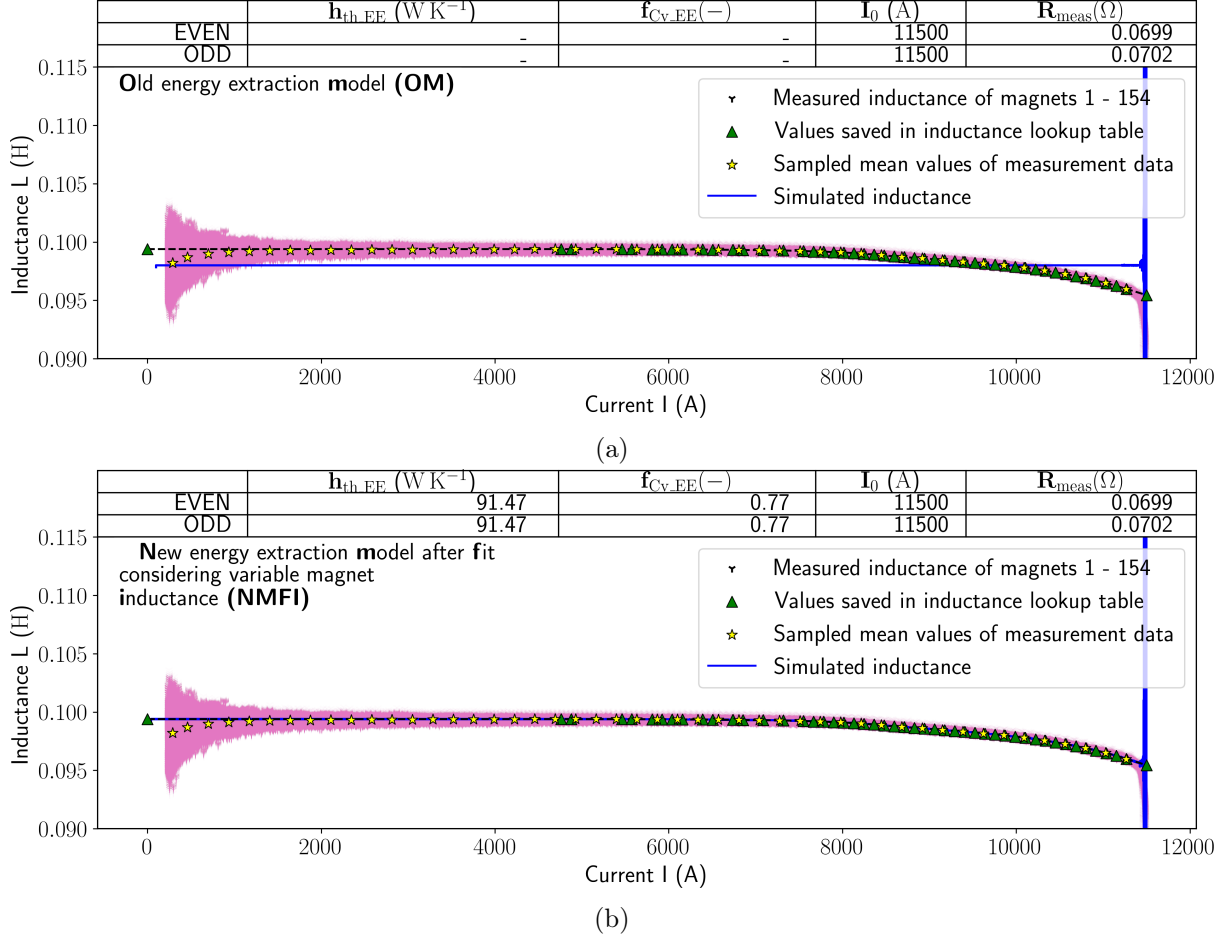


Figure 3.10: Comparison of the simulated and the measured inductance before and after current-dependent inductance was considered in the MB magnet model. The magenta markers illustrate the measured inductance as obtained by equation 3.8. The blue plot corresponds to the simulated inductance. The green triangles connected by the dashed line mark the monotonic subset of values saved within the model's inductance look-up table. The yellow stars correspond to the mean value of the measured inductance. a) Mismatch between measured inductance and the constant inductance assumed in the old MB magnet model. b) Improved accuracy achieved by implementing current-dependent inductance within the MB magnet model.

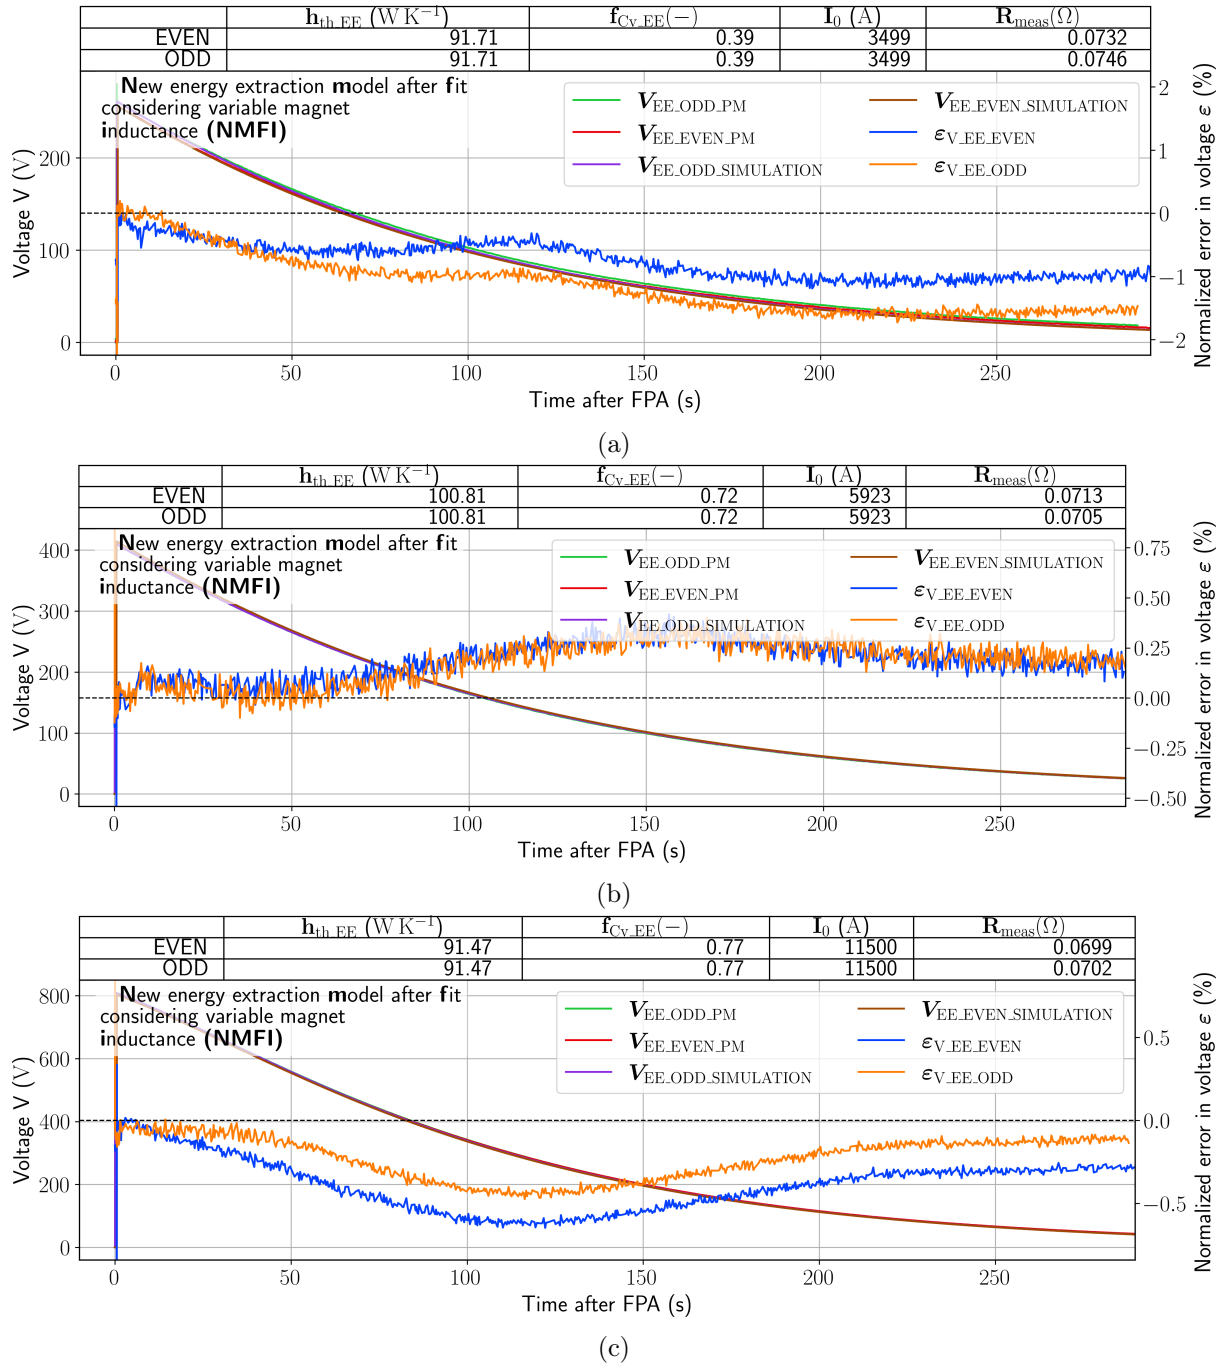


Figure 3.11: Voltage across the energy extraction V_{EE} obtained with the NMF for FPA current levels of: a) $I_0 = 3499$ A b) $I_0 = 5923$ A c) $I_0 = 11500$ A. Further characteristic model parameters are given in the tables above the respective plots.

with constant inductance, it is noticeable that the curves of the measured and simulated voltages of V_{EE} almost coincide. Furthermore, the magnitude of the maximum errors is further reduced across all current levels, in comparison to the previous fit of the model with constant inductance. For all current levels, an initially small error is observed, which however increases over the course of the discharge until reaching a local maximum in magnitude. Hence, the simulation in V_{EE} is most accurate in the early stages after the FPA, where the influence of resistor heating is still small but also towards the end of the discharge, where the currents within the circuits (and hence also the signals and their errors) become smaller.

Circuit current I_{MEAS}

In the following, the currents through the energy extraction as obtained with the final version of the model (NMFI) are discussed, based on figure 3.12. Comparing the figures 3.12c and 3.9b, it can be seen that the error in I_{MEAS} is further reduced in comparison to the NMF. This qualitative observation holds true across all current levels. Furthermore, an almost constantly negative error in I_{MEAS} can be observed, indicating that for most of the decay, the simulated current is slightly underestimated. In all three cases however, there is no consistently rising or falling trend in the absolute error, but rather only a temporary increase in magnitude. The maximum absolute error occurs for all tested FPA current levels between approximately 100-150 s.

Resistance of the energy extraction R_{EE}

In subsequent, the resistance of the energy extraction as obtained with the final version of the model (NMFI) is discussed, based on figure 3.13. Considering the simulated R_{EE} , conclusions similar to those for the NMF can be drawn. Especially at high current levels, the fit has adapted very well to the measured resistance curve, as can be seen from figure 3.13c. Furthermore, by comparing figures 3.13c and figure 3.9c, it is observed that the simulated resistance derived for the NMFI shows better agreement with the measurement data than the resistance simulated with the NMF. Across all current levels, there is a strong agreement between simulated and measured resistance within the first few seconds. However, at later stages of the simulation, the curves of simulated resistance and measured resistance stop to coincide after a certain time point. This time point of separation happens to be increasingly earlier as the FPA current I_0 decreases, which might be attributed to the insufficient physical representativeness of the measured resistances at low currents.

3.3.4.3 Simulated and measured temperatures

In the following section, the module temperature as calculated within the NMFI is compared to measurement data for all current levels tested. At first, the temperature increase for the lowest tested FPA current $I_0 = 3499$ A, as illustrated in figure 3.14a will be discussed. Considering the course of the temperatures, it is noticeable that for both EE units, the simulated rise in temperature happens much faster than the measured temperature increase. Furthermore, the steady state temperature increase of the simulation deviates by more than a factor of two from the measured temperature increase. Additionally, even after fitting, a strong deviation between ΔT_{calc} and the simulated module temperature can be observed. However, within the first 10 s after the FPA, the curves align.

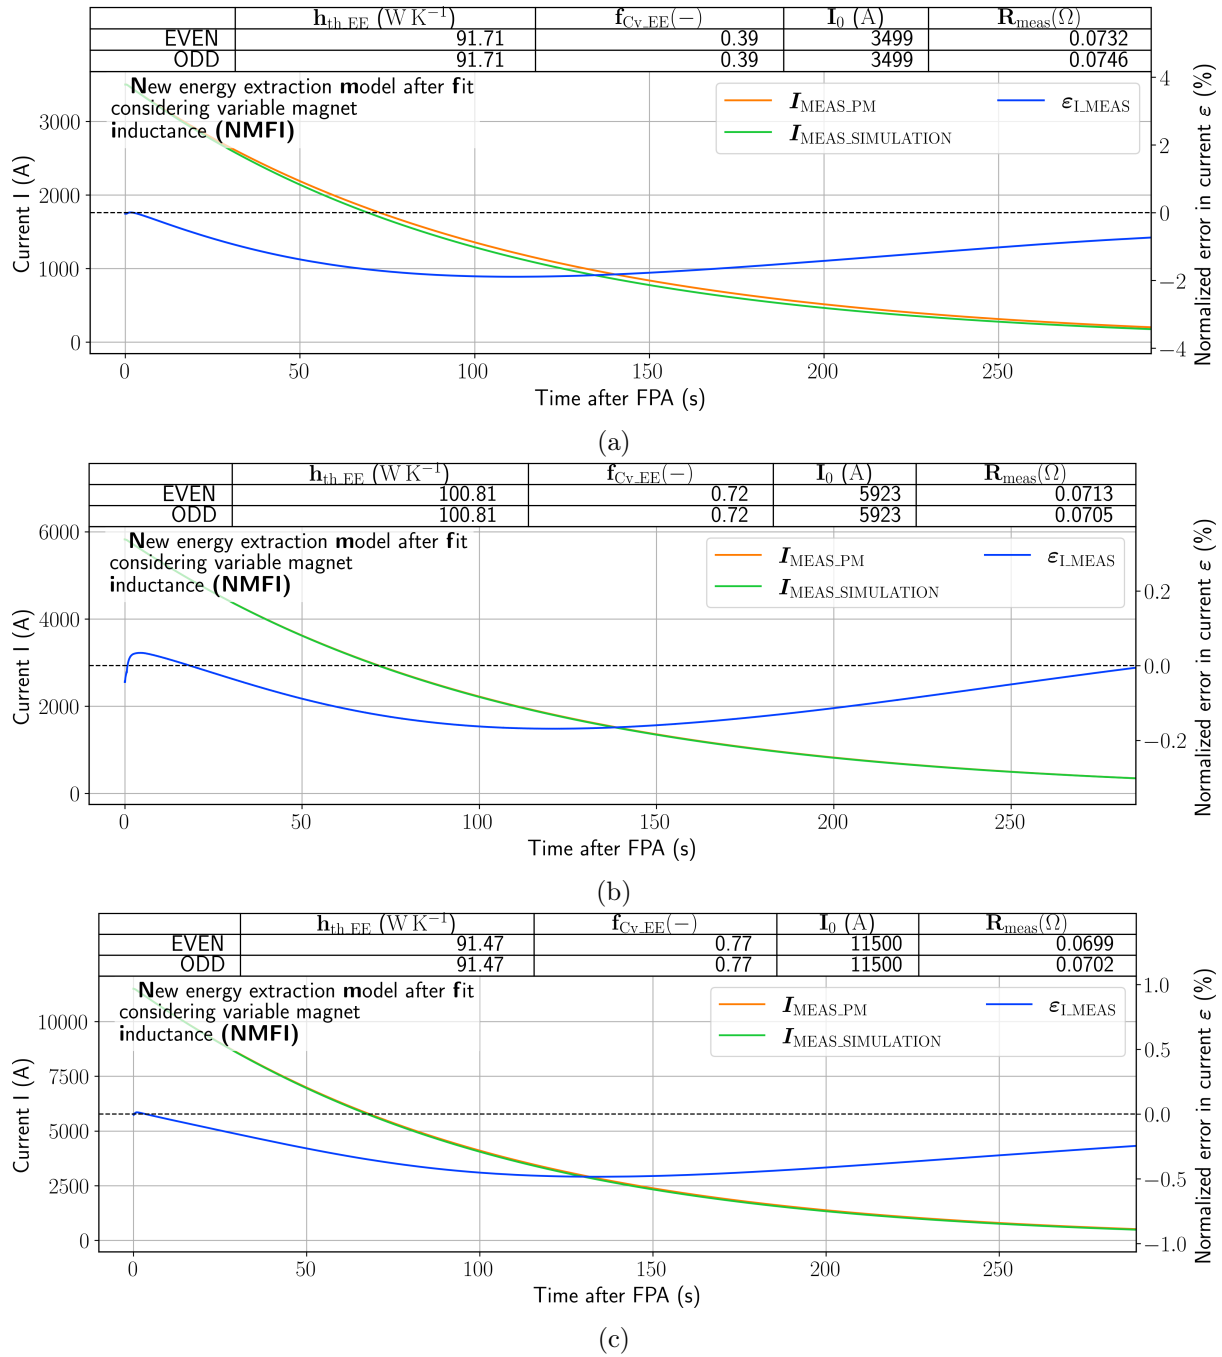


Figure 3.12: Current through the circuit I_{MEAS} obtained with the NMFI for FPA current levels of: a) $I_0 = 3499$ A b) $I_0 = 5923$ A c) $I_0 = 11500$ A. Further characteristic model parameters are given in the tables above the respective plots.

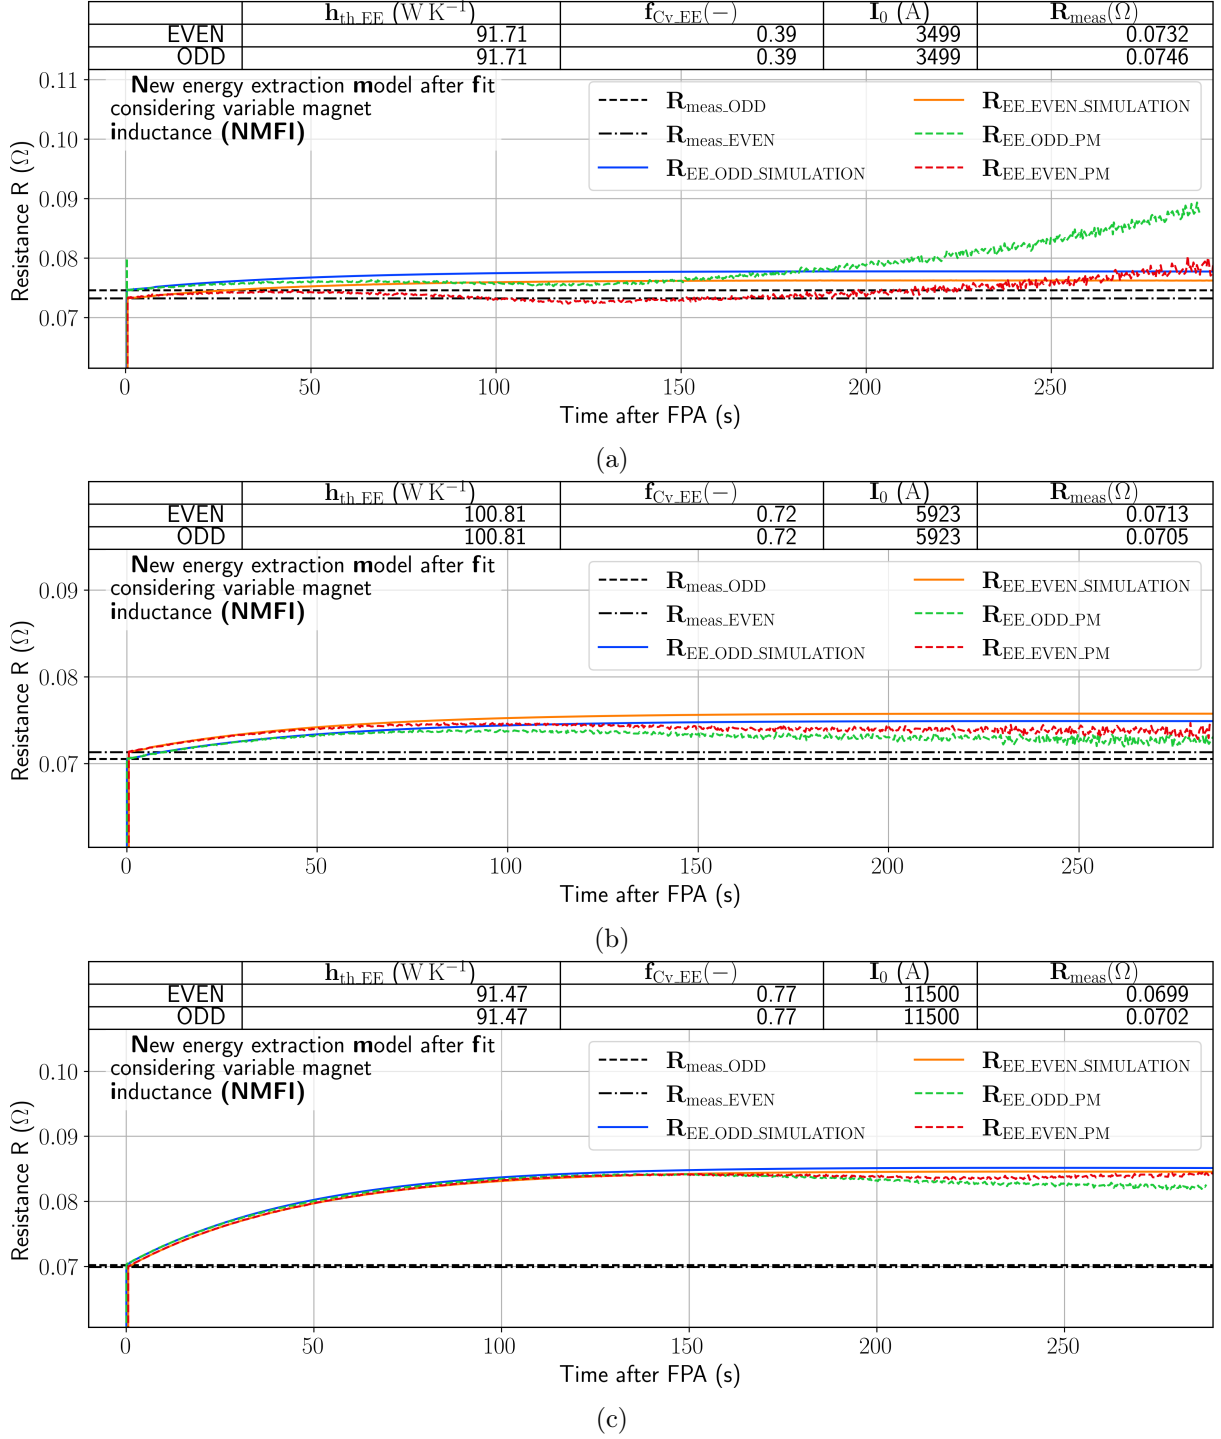


Figure 3.13: Simulated resistance of the energy extraction R_{EE} obtained with the NMFI for FPA current levels of: a) $I_0 = 3499$ A b) $I_0 = 5923$ A c) $I_0 = 11500$ A. Further characteristic model parameters are given in the tables above the respective plots.

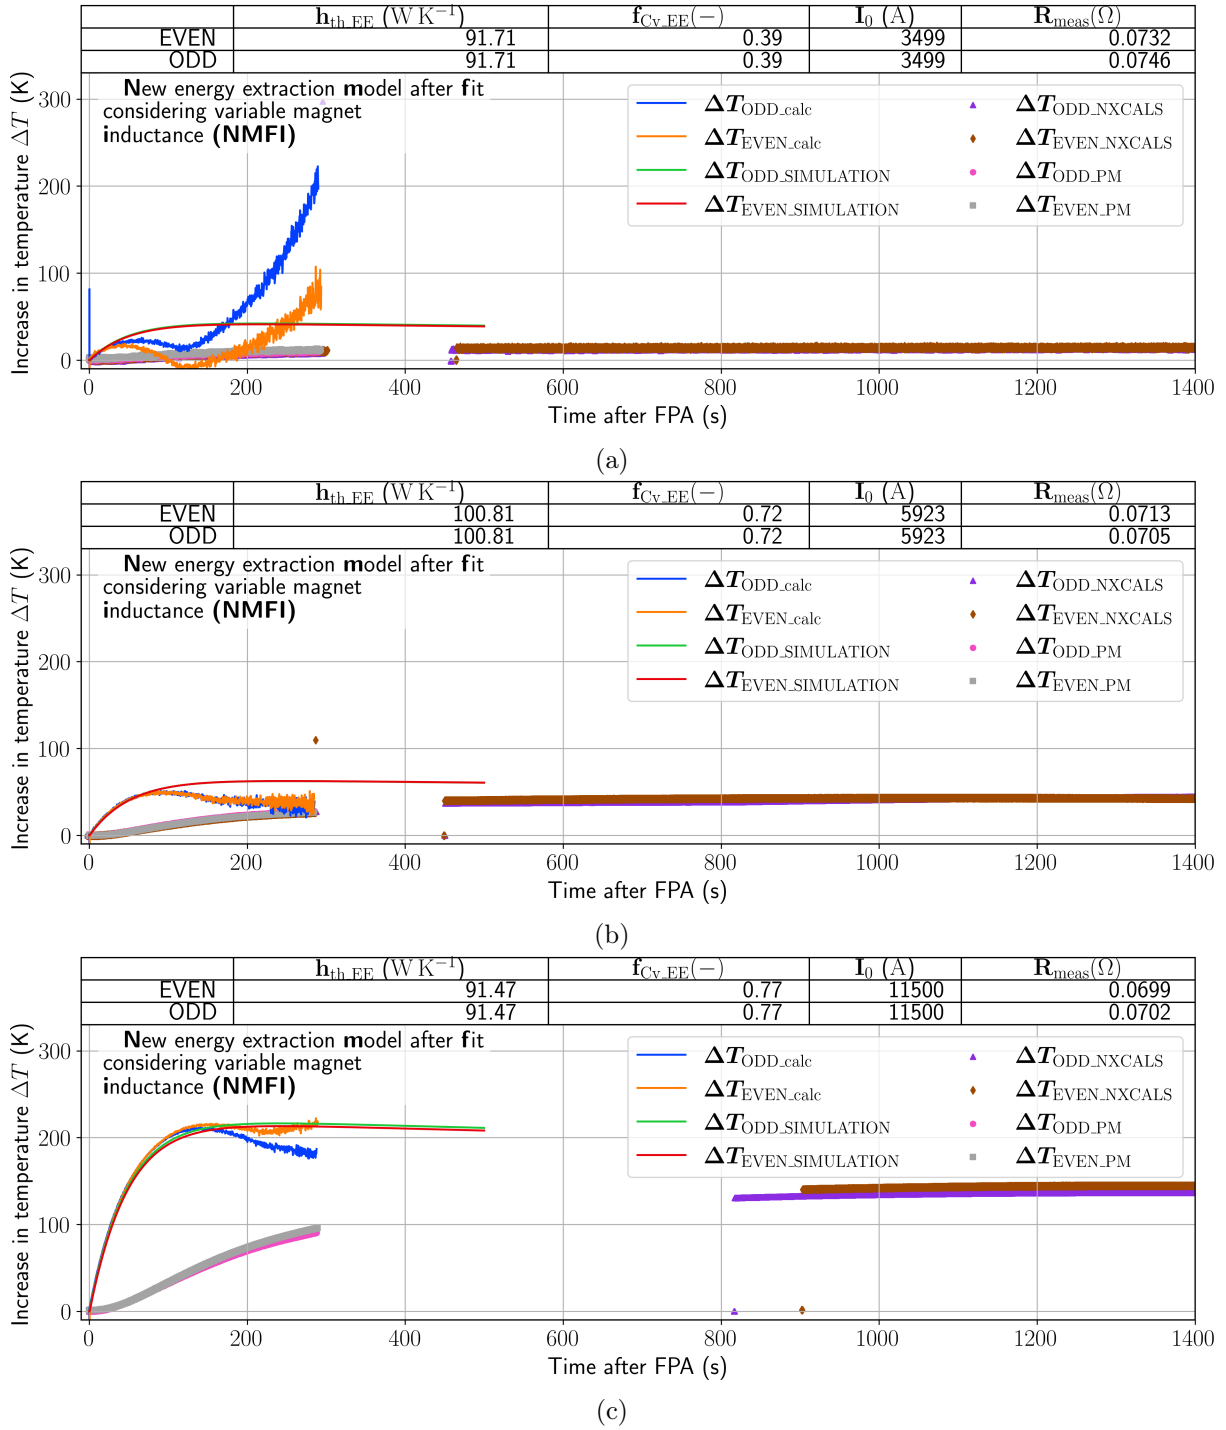


Figure 3.14: Increase in resistor temperature as simulated with the NMFI in comparison to measurement data and the temperature increase which would be necessary to explain the measured increase in resistance ΔT_{calc} (assuming homogeneous resistivity, heat capacity and temperature across the resistor). The plots correspond to different FPA events, which occurred at different current levels I_0 .

Similar conclusions can be drawn for the temperature increase analyzed at $I_0 = 5923$ A, which is illustrated in figure 3.14b. Also in this case, the initial rise in the simulated temperature increase is greater than that of the measurement. Furthermore, just like observed for lower I_0 , the experimentally-derived temperature ΔT_{calc} and the simulated temperature align within the first 100 s. The steady state of the simulated increase in temperature of about 60 K however does not align with the measured increase in temperature, which remains at about 40 K.

Lastly, the temperature increase at $I_0 = 11500$ A as shown in figure 3.14c will be discussed. The curves of the simulated temperature increase and ΔT_{calc} coincide at the beginning of the transient. Compared to the lower current levels however, this overlap is significantly longer and lasts for almost 180 s. Here too, the final temperature increase obtained through simulation does not coincide with the final temperature increase of the **NXCALS** measurement signal.

The observations made for the simulated temperature can be interpreted as follows. During the fitting process, the objective function, evaluating the alignment between the measured and simulated signal V_{EE} is minimized. Hence, the simulated temperature is more and more aligning with the experimentally-derived temperature ΔT_{calc} in order to match the measured resistance. However, as already discussed in section 3.3.2.2, the measured resistance (and hence also ΔT_{calc}) is subject to perturbations, most likely due to elevated influences of signal noise when measuring small current signals. Therefore, the models free parameters are fitted to erroneous resistances, resulting in comparatively big errors within the simulation signals. The reason for which the steady state temperature increase is not matched anymore after fitting can be explained by the changes in $f_{\text{Cv_EE}}$. The fitting algorithm can only match the initial overshoot of ΔT_{calc} , if $f_{\text{Cv_EE}}$ is reduced, as this causes a rise of temperature and resistance. However, changing $f_{\text{Cv_EE}}$ further causes the full ohmic heating energy to be deposited within a downscaled resistor heat capacity, causing the simulated steady state temperature increase to rise, and hence deviate from the measured steady state temperature increase. In a future work, the **EE** resistor model could therefore be further refined in order to model temperature gradients in the resistor, by subdividing the resistor module's volume into multiple parts. This would allow a better match of simulated and measured temperatures than it is possible with the model derived in this thesis, which primarily focussed on the accurate simulation of electrical signals.

3.3.5 Meta analysis of the parameter fits

Since the further development of the **STEAM-DI** was an essential part of this thesis, the following section will take a closer look at the fitting process. The optimization aims to minimize a so-called objective function, which is defined as:

$$\text{objective function} = \frac{\bar{\varepsilon}_{V_EE_ODD} + \bar{\varepsilon}_{V_EE_EVEN}}{2} \quad (3.9)$$

with $\bar{\varepsilon}_{V_EE_ODD}$ and $\bar{\varepsilon}_{V_EE_EVEN}$ being the mean simulation errors (definition given in equation 3.7) for V_{EE} of the two energy extraction units. The courses of the optimization processes leading to the optimum parameter configurations finally selected for deployment of the model are now discussed in more detail, based on figure 3.15. In these optimizations, the initial point in parameter space was selected to be ($f_{\text{Cv_EE}} = 1.0$; $h_{\text{th_EE}} = 100 \text{ W K}^{-1}$). This point, whose actual value is far from the point at which the optimum was expected was chosen in order to better investigate the optimization capabilities of the algorithm.

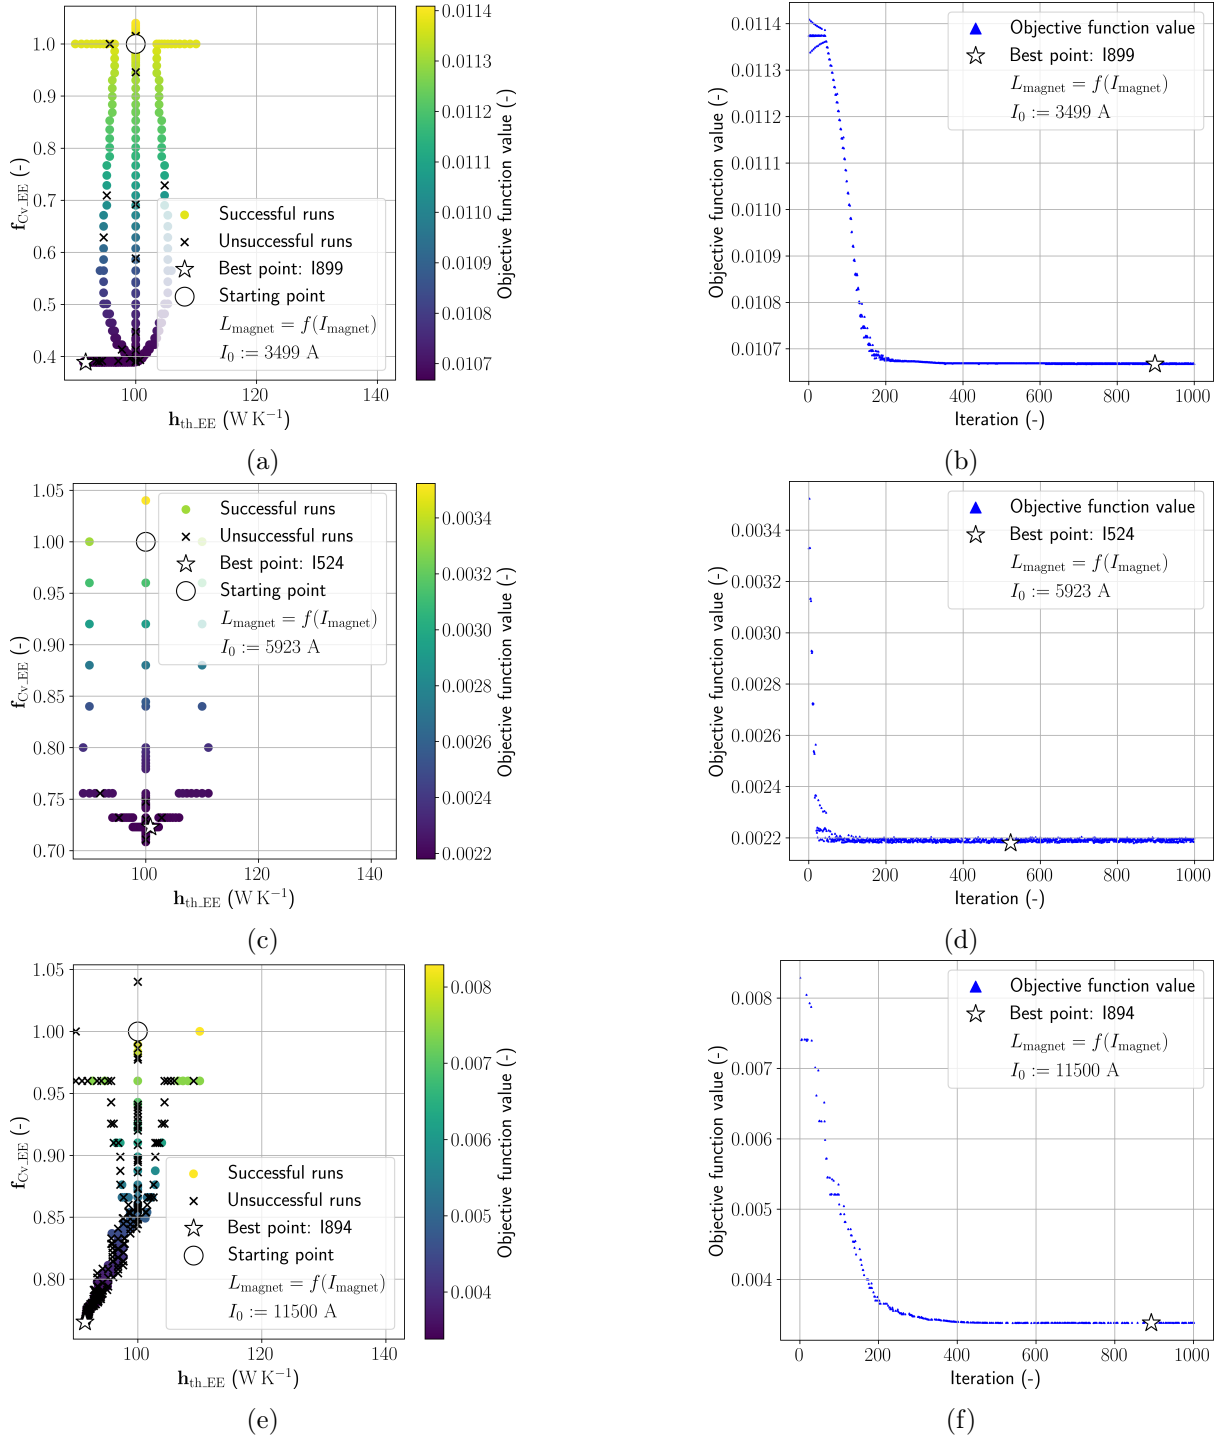


Figure 3.15: Parameter optimization of the newly developed EE model using a pattern search algorithm. a), c), e) Objective function landscape showing the initial point and the best point of the optimization as well as the movement of the tested points within parameter space for the different FPA current levels tested. b), d), f) Value of the objective function over the course of the optimization for the different FPA current levels tested. Each row of the figure corresponds to a different FPA event, which occurred at a different current level I_0 .

In all three optimizations, the algorithm first reduces the heat capacity by scaling down $f_{\text{Cv_EE}}$ in order to increase the simulated temperature and hence the simulated resistance. The final values corresponding to the events at 5923 A and 11500 A are $f_{\text{Cv_EE}} = 0.72$ and $f_{\text{Cv_EE}} = 0.77$.

Considering the parameter landscape of the the optimization corresponding to the FPA event that occurred $I_0 = 3499$ A (shown in the figures 3.15a), it is noticeable that with this optimization, the optimum point of $f_{\text{Cv_EE}}$ is much lower. As explained in section 3.3.4.3, the measured resistance for this event appears to increase to unrealistically high values towards the end of the discharge, which may be attributed to noise within the signals. The optimization algorithm is however aiming to mimic this behavior with the simulation and thus reduces the heat capacity in order to increase the simulated resistance. The optimum values found for this optimization could therefore deviate from the actual physically representative values.

Furthermore, the final optimum points found at all current levels feature a thermal conductance close to $h_{\text{th_EE}} = 100 \text{ W K}^{-1}$. However, comparing these with the values for $h_{\text{th_EE}}$ found in a previous optimization study (in which the magnet differential inductance was constant, see A.2), it is noticeable that the optimum values are widely scattered, and are in some cases very close to the starting point of the optimization.

Hence, the sensitivity of the objective function to $h_{\text{th_EE}}$ is assumed to be small. On the contrary, as the value $f_{\text{Cv_EE}} = 0.77$ found at $I_0 = 11500$ A is consistent across different starting points of the optimization (see figures 3.15e and A.2c for comparison), it can be assumed to be physically representative. However, it should be noted that the optimizations carried out within this thesis are local optimizations, which provide "good" optima, but not necessarily "the best".

In the early stages of all optimizations, the objective function can steadily be reduced. For all FPA currents tested, convergence of the optimization is achieved after about 400 iterations at the latest, which is noticeable through the sudden flattening of the objective function's curves shown in figures 3.15b, 3.15d and 3.15f. This sudden flattening indicates convergence of the optimization. After the objective function reaches a plateau due to the convergence of the optimization algorithm, any further iteration causes only minimal improvements in the objective function's magnitude.

For the optimization corresponding to the event that occurred at $I_0 = 11500$ A (shown in figure 3.15e), a lot of unsuccessful simulations can be noticed. However, the algorithm nevertheless converges to a minimum, as can be seen from the course of the objective function's magnitude shown in figure 3.15f. This indicates the robustness of the method developed, especially in the process of model prototyping, where stable initial starting values of the model parameters are unknown before the optimization³.

3.4 Summary

This chapter described the implementation and optimization of a new model of the LHC main dipole circuit's energy extraction resistors. This new model considers heating and cooling effects

³The multitude of failing optimizations observed in figure 3.15e is not a cause of a poorly chosen initial starting point, as other simulations with parameter vectors close to those that caused a failure in some cases ran successfully. The observed behavior may rather be attributed to a too short process timeout chosen (900 s), after which frozen simulations are aborted.

Table 3.1: Summary of the mean simulation errors in predicting current, voltage, and resistance of the EE for all model configurations and current levels tested. The parameters found for the optimizations of FPA events with $I_0 = 3499$ A are set in braces, as they might be physically less representative than the values found for other current levels, as explained in section 3.3.5.

model	fitted? (yes/no)	L_{magnet} (constant/ variable)	I_0 (A)	$f_{\text{Cv_EE}}$ (-)	$h_{\text{th_EE}}$ (W K ⁻¹)	$\bar{\epsilon}_{\text{I_MEAS}}$ (%)	$\bar{\epsilon}_{\text{V_EE_odd}}$ (%)	$\bar{\epsilon}_{\text{V_EE_even}}$ (%)	$\bar{\epsilon}_{\text{R_EE_odd}}$ (%)	$\bar{\epsilon}_{\text{R_EE_even}}$ (%)
OM	n	c	3499	-	-	0.92	1.71	1.32	3.97	1.55
OM	n	c	5923	-	-	0.53	0.83	0.8	3.27	3.32
OM	n	c	11500	-	-	2.54	2.53	2.47	13.66	14.27
NM	n	c	3499	1	0	1.16	1.63	1.25	2.87	1.34
NM	n	c	5923	1	0	0.17	0.33	0.31	0.8	0.71
NM	n	c	11500	1	0	0.1	0.67	0.8	1.93	2.81
(NMF)	(y)	(c)	(3499)	(0.32)	(1.09)	(1.68)	(1.48)	(1.11)	(2.64)	(2.72)
NMF	y	c	5923	0.6	1728.96	0.39	0.29	0.28	0.87	0.83
NMF	y	c	11500	0.77	1.01	0.53	0.49	0.64	1.44	0.5
(NMFI)	(y)	(v)	(3499)	(0.39)	(91.71)	(1.2)	(1.15)	(0.76)	(2.56)	(2.2)
NMFI	y	v	5923	0.72	100.81	0.09	0.18	0.19	1.58	1.54
NMFI	y	v	11500	0.77	91.47	0.31	0.23	0.39	1.26	0.38

by assuming homogeneous temperature, heat capacity and resistance across the resistor modules. For the resistors' thermodynamic modelling, two distinct free parameters were introduced: $h_{\text{th_EE}}$, which is proportional to the resistors' cooling and $f_{\text{Cv_EE}}$, used to scale the resistors' initially assumed heat capacity up or down. In the further course, different model configurations were compared across different FPA current levels, based on characteristic electrical signals. These included the voltage across the energy extraction V_{EE} , the resistance of the energy extraction R_{EE} and the circuit current I_{MEAS} . In the following, the mean errors of these electrical signals are summarized, based on table 3.1.

As can be seen, the implementation of the new model (NM) caused a significant improvement in simulation accuracy across all tested FPA current levels. By fitting the model's free parameters (NMF) with Dakota, the simulation accuracy could be improved even further, especially at high FPA currents. In order to better match the simulated circuit current I_{MEAS} , the model of the main dipole circuit's magnets was supplemented by considering current-dependent magnet differential inductance. The current dependence of the magnet differential inductance is attributed to the saturation of the iron-yoke around the magnet coils and was determined based on PM measurement data. After this implementation, the model was fitted once more (NMFI) which caused the desired improvement in accuracy of the simulated circuit current. Comparing the OM and the NMFI for the highest FPA current tested, the mean error in the simulation of I_{MEAS} could be reduced from 2.54% to 0.31%. Furthermore, the errors of the simulated signal of V_{EE} and R_{EE} could be reduced by about 2% and 14% respectively.

In all optimizations, the scaling factor $f_{\text{Cv_EE}}$ was reduced by the algorithm, which is attributed to the effect of uneven heating of the resistor. The optimum thermal conductance $h_{\text{th_EE}}$ notably scattered for different optimizations, but often converged to values corresponding to virtually zero cooling. On the contrary, the optimum value of $f_{\text{Cv_EE}} = 0.77$, determined at $I_0 = 11500$ A is consistent across different starting points of the optimization.

Table 3.2: Comparison of the mean simulation errors in predicting current, voltage, and resistance of the energy extraction. The table compares the finally deployed model (DM) to the Old Model (OM).

model	I_0 (A)	f_{Cv_EE} (-)	h_{th_EE} (W K ⁻¹)	$\bar{\epsilon}_{I_MEAS}$ (%)	$\bar{\epsilon}_{V_EE_odd}$ (%)	$\bar{\epsilon}_{V_EE_even}$ (%)	$\bar{\epsilon}_{R_EE_odd}$ (%)	$\bar{\epsilon}_{R_EE_even}$ (%)
OM	3499	-	-	0.92	1.71	1.32	3.97	1.55
OM	5923	-	-	0.53	0.83	0.8	3.27	3.32
OM	11500	-	-	2.54	2.53	2.47	13.66	14.27
DM	3499	0.77	91.47	1.87	2.15	1.82	2.7	1.41
DM	5923	0.77	91.47	0.68	0.53	0.52	1.27	1.21
DM	11500	0.77	91.47	0.31	0.23	0.39	1.26	0.38

With respect to the simulated temperatures, it could be seen that the fitted models allowed to reproduce the experimentally-derived temperature increase ΔT_{calc} , which is the temperature increase needed to account for the measured resistance increase of the energy extraction resistors (assuming homogeneous resistivity, heat capacity and temperature across the resistor). It was furthermore shown that the NM (in the adiabatic case $h_{th_EE} = 0$) could reproduce the measured maximum temperature increases for low current FPA events to a good approximation (see section 3.3.2.2). The measured resistor temperature can not be reproduced which is due to the model's underlying assumption of a homogeneous temperature distribution across the resistor. This is however beyond the scope of this thesis, since the resistor temperature does not directly affect the transients in the superconducting magnets.

Since the effects of resistor heating were observed to be most pronounced at high FPA currents, the model was deployed with the optimum parameters found for the fit with the highest I_0 , which had values of $h_{th_EE} = 91.47$ W K⁻¹ and $f_{Cv_EE} = 0.77$. To ensure that these parameters also bring the desired improvement for other FPA current levels, all three events were simulated once more with the final parameters $h_{th_EE} = 91.47$ W K⁻¹ and $f_{Cv_EE} = 0.77$. The mean errors of these simulations using the final model parameters are given in table 3.2. It can be seen that in comparison to the OM the finally deployed model (DM) reproduces the measured resistances more accurately. This is the case across all FPA current levels tested. Furthermore, the simulated voltage across the energy extraction as simulated with the DM shows a clear improvement in accuracy in comparison to the OM, as an improvement in mean errors is observed for two out of the three current levels analyzed.

For lower FPA currents it is observed that the improvements in simulation accuracy of R_{EE} do not lead to an improvement in simulation accuracy of I_{MEAS} . The reason for this could lie in inaccuracies in other circuit components, causing the circuit current to be underestimated by the simulation, even before the new energy extraction model was implemented. In such a case, a further reduction in the circuit current which comes with the implementation of the new model causes a reduction in the accuracy of I_{MEAS} despite the EE model being more accurate.

Overall, the newly developed model of the LHC main dipole circuit's energy extraction resistors provides an advanced prediction of various characteristic circuit signals. It will therefore contribute to more precise analyses of FPA events, especially in the case of FPA events that

occurred at high circuit currents.

Model of the LHC main dipole magnet's by-pass diodes

In addition to the model of the main dipole circuit's energy extraction resistors, a thermoelectric model of the main dipole circuit's by-pass diodes was implemented in this thesis. This model is presented in the following chapter. The model considers the diodes' heating and cooling effects, as temperature has a decisive influence on the diodes' I-V characteristic. First, section 4.1 provides a general introduction to the LHC main dipole magnet's by-pass diodes and their modelling in this thesis. In section 4.2, the results of the simulations are presented and compared to measurement data. Finally, the results are summarized and discussed in section 4.3.

4.1 Model

If a quench in one of the LHC's main dipole circuits is detected, its quench protection heaters are triggered in order to safely transfer parts of the magnet coil back to normal conducting state. In order to provide a low resistance path for the current in such a case, each magnet of the main dipole circuits is protected by a parallel-connected by-pass diode [42, 79, 80]. As the diode is installed in the magnet cold mass, it must operate at cryogenic temperatures. As shown in figure 4.1, the diodes are enclosed within an insulating capsule, the so-called press-pack, which serves as a protection against contamination and mechanical forces. Inside the capsule, the silicon wafer, which acts as the actual diode, is attached to two molybdenum plates, which are in contact to two copper pole pieces. Furthermore, a silver foil is placed between the top of the wafer and the upper molybdenum plate, which further improves the thermal and electrical coupling of the wafer to the plates. The diode capsule itself is installed within the so-called diode stack, which is placed in a helium container inside of the cryostat of the magnet as shown in figure 4.2 [80]. Within the diode stack, the press-pack is clamped between two nickel-coated copper heat sinks, pressed together with a force of 40 kN [42]. These high pressure forces, which are applied to the stack by a Cu-Be plate spring [83], are necessary in order to ensure a strong electrical and thermal coupling of the heat sinks to the press-pack. After a magnet's quench protection heaters are triggered and parts of the magnet coil become resistive, the circuit's current primarily flows through the diode. This current flow causes the wafer to heat up, which has a strong influence on the I-V characteristics of the diode. The following section presents how this temperature dependence of the I-V characteristic was modelled in this thesis.




Figure 4.1: Cold by-pass diode of the main dipole circuit. Left) Diode press-pack of the by-pass diode (image taken from the work of Will [80]). Right) Cross-section of the diode press-pack (image taken from a datasheet [82] and modified.)




Figure 4.2: Structure of the diode stack, showing how the diode stack is placed and how it is connected to the cryostat. Left) Diode stack (image taken from [83] and modified). Middle) Technical drawing of the diode stack with the heat sinks being colorized in blue and the diode being colorized in yellow (image based on a modified technical drawing from CERN). Right) image of an open main dipole magnet showing the position of its helium container in which the diode stack is installed (image taken from [84] and modified).

4.1.1 Modelling of the diodes' temperature-dependent I-V characteristic

The current-voltage characteristic of diodes can be modelled with the Shockley equation [85]:

$$I_D = I_S \cdot \left(\exp \left(\frac{V_D}{n \cdot V_T} \right) - 1 \right) \quad (4.1)$$

with the reverse-bias saturation current I_S , the ideality factor n and the thermal voltage V_T of the diode. The I-V characteristic of diodes is strongly temperature-dependent. This temperature dependency is reflected in the equation's parameters V_T , I_S and n , which depend on the temperature of the diode's wafer. The thermal voltage is linearly dependent on the temperature:

$$V_T = \frac{k_B \cdot T_D}{q}, \quad (4.2)$$

with T_D being the wafer's temperature, k_B being the Boltzmann constant and q being the elementary charge of an electron. Another temperature dependence is found in the reverse-bias saturation current I_S . Its temperature dependence is in general non-linear and relates to characteristic parameters of the p-n junction which also have a temperature dependency. The reverse-bias saturation current can be approximated by considering the following equation [86]:

$$I_S = n_i^2 \cdot A_j \cdot q \cdot \underbrace{\left(\frac{D_h}{N_D \cdot L_h} + \frac{D_e}{N_A \cdot L_e} \right)}_{\text{weakly temperature dependent}} \quad (4.3)$$

with A_j being the junction area, n_i being the intrinsic carrier concentration, D_h and D_e being the charge carrier diffusion constants of the holes and electrons, L_h and L_e being the charge carrier diffusion lengths of the holes and electrons and N_A and N_D being the acceptor and donor concentrations on the p and n side of the junction. Within this equation, the most dominant temperature dependency is given with the intrinsic carrier concentration n_i [86], which can be approximated with the following formula:

$$n_i^2 = K \cdot T^3 \cdot \exp \frac{-\theta_0}{k_B \cdot T}, \quad (4.4)$$

with K being a constant and θ_0 being the semiconductor's band gap extrapolated to 0 K [86]. Considering equation 4.4 at temperatures close to 0 K, it becomes apparent that the carrier concentration and hence also the reverse-bias saturation current become vanishingly small. This phenomenon is often referred to as carrier freeze-out [80].

The ideality factor n , which is intended to represent non-ideal diode behavior such as carrier recombination, can also be temperature-dependent. Hence, there is no standardized physical law that describes how n is affected by temperature. In order to analyze the trend of n and I_S , the Shockley equation was fitted to measurement data recorded at different temperatures, using a least-square approach. This survey included the following data sources [80, 81, 79, 87]:

- Measurement data from the CERN diode testing facility (building 272), which was conducted at 298 K and 77 K (denoted with B272-298K and B272-77K).
- Reference data measured at 4.2 K (denoted with D8_Vto_Ref).




Figure 4.3: Measurement data of the diode's I-V characteristic. The measurement data exhibits clear deviations from ideal diode behavior. Therefore, the limits indicated by the dashed lines were applied to the data, in order to fit the Shockley equation only to ranges in which it can describe the diode's behavior sufficiently well.

- Measurements that were conducted at CERN to have a comparison to the manufacturer's (DYNEX) specifications at 298 K and 77 K (denoted with CERN_DYNEX-298K and CERN_DYNEX-77K).
- Forward voltage characteristics from a similar diode's datasheet [82]. Since no other high temperature data was available, the datasheet of a diode comparable to the ones installed at LHC was considered. Within this datasheet, the diode's forward characteristic at 398 K is specified in the form of an equation. This equation was used to synthesize measurement data at 398 K.

However, the Shockley equation as formulated in equation 4.1 could not represent the measurement data very well, as strong deviations between the fit and the measured data were observed. The reason for this could be traced back to the measurement data not following an exact exponential course (like predicted by the Shockley equation), as shown in figure 4.3. This observation is particularly pronounced with the 4.2 K measurement, where against expectations, the forward voltage seems to decrease with increasing diode currents after 400 A. The reason for the observed deviation from ideal diode behavior lies in the measuring process and can be attributed to two effects which are not considered in equation 4.1:

- Increase in temperature of the diode during the measurement: If a diode is measured at very low temperatures, its initial heat capacity is very low. Therefore, the diode can be significantly heated during the measuring process, leading to distorted measurements. This effect appears as a shift of the measurement data towards lower forward voltages as seen for currents higher than 400 A in the 4.2 K measurement shown in figure 4.3.
- Additional series resistances: The measurements were carried out with currents reaching up to several kA. Since voltage and current scale linearly according to ohm's law, the




Figure 4.4: Fits of the Shockley equation (without series resistance) applied to the cropped measurement data.

effects of parasitic series resistances of the diode can cause a pronounced shift of the I-V characteristic towards higher voltages, especially at high currents.

In a further step, a series resistance R_S was added to the Shockley equation as follows:

$$I_D = I_S \cdot \left(\exp \left(\frac{V_D - I_D \cdot R_S}{n \cdot V_T} \right) - 1 \right). \quad (4.5)$$

This equation was fitted to the measurement data, leading to a much better representation of the measurements. The fitting parameters found with this approach however did not expose a clear trend over temperature, but seemed to change arbitrarily with temperature (see figure A.3a and figure A.3b in the appendix). This made reliable and physically justifiable modelling of the parameters impossible. Hence, another approach was chosen. In a first step, the series resistance was neglected by fitting equation 4.1 only to low current measurement data, where the influence of the diode's series resistance is less pronounced. For temperatures from 77 K on, only datapoints of currents below 5000 A were considered, in order to cut out regions where the data does not follow an exponential law anymore. For the data taken at 4.2 K, only datapoints below 400 A were considered in order to cut out the influence of heating effects during the measurement. Furthermore, the data points below 3.8 V were cut off in the 4.2 K measurements, as they do not follow an exponential trend but can rather be attributed to leakage currents of the closed diode. The measurement data trimmed in this way, as well as the corresponding fits to the Shockley equation (without series resistance as given with equation 4.1) are shown in figure 4.4. As can be seen, the Shockley equation without series resistance as a fitting parameter provides a good representation of the trimmed measurement data.

In order to extend this equation to include the influence of temperature, the temperature dependency of n and I_S had to be modelled. The respective values of the Shockley equation's fitting parameters n and I_S , found at different temperatures, are shown in figure 4.5. It is noticeable that the values found for I_S are not monotonic, which makes it impossible to model them using equation 4.4. Hence, a more practical approach had to be chosen, which describes




Figure 4.5: Mathematical modelling of the temperature dependence of the parameters n and I_S . The $\text{clip}(\cdot, l1, l2)$ operator limits the functions to the lower and upper boundaries $l1$ and $l2$, in order to prevent numerical instabilities caused by negative or excessive values.

the course of I_S using an exponential polynomial function. This type of function was chosen to ensure both good adaptability to the data points (polynomial) and consistently positive values (exponential function). The selected fitting function is hence given by:

$$I_S(T_D) = 1A \cdot \exp(c_0 + c_1 \cdot T_D + c_2 \cdot T_D^2 + c_3 \cdot T_D^3), \quad I_S(T_D) \in (1e^{-13}, 0.1) \text{ A}, \quad (4.6)$$

with c_0, c_1, c_2 and c_3 being fitting parameters, whose values are given in the legend of figure 4.5. In order to prevent numerical instabilities, the reverse-bias saturation current was limited to a range of values between $1e^{-13}$ A and 0.1 A. For the ideality factor n , a strong increase with decreasing temperatures is observed. For Schottky diodes this effect of strongly increasing ideality factors for low temperatures is sometimes referred to as the "T0 anomaly" [88, 89], which may be described with the following fit:

$$n(T_D) = a + \frac{b}{T_D}, \quad (4.7)$$

with a and b being fitting parameters [89]. Despite the LHC diodes not being Schottky diodes, a similar higher order formulation showed a good representation of the observed temperature dependence of the ideality factor, which is given by:

$$n(T_D) = a + \frac{b}{T_D} + \frac{c}{T_D^2}, \quad n(T) \in (1.9, 1000), \quad (4.8)$$

with a, b and c being fitting parameters, whose values are given in the legend of figure 4.5. In order to prevent numerical instabilities, the ideality factor was limited to a range of values between 1.9, the lowest value of n found in the fits, and 1000 which was set as an upper limit. These fits found for the temperature dependence of n and I_S were inserted into the Shockley equation in order to compare the modelled I-V characteristic to measurement data.

During the construction of the LHC, the I-V characteristics of all diodes were measured at 298 and 77 K at the ENEA Research Center (Frascati, Italy). Based on these measurements,

Gharib et al. [42] published an interval in which the I-V curves of the LHC diodes may vary. These measurements were overlaid with the modelled temperature dependency of the Shockley equation found within this thesis, as shown in figure 4.6a. The modelled temperature dependency of the I-V characteristic proves to be a good approximation of the diode behavior at 298 K, but not of the measurement data recorded at 77 K, as the modelled curve lies outside the intervals specified by Gharib et al. [42]. This is due to the fact that no series resistance has been considered so far.

In order to remain within the intervals of possible diode characteristics, a constant series resistance of $R_S = 9.677e^{-6} \Omega$ was added to the model. With this series resistance, and the correlations found for $I_S(T)$ and $n(T)$, the modelled I-V characteristic lies within the measurement intervals across all temperatures, as can be seen in figure 4.6b. Figure 4.6c shows the finally modelled temperature dependence of the I-V characteristic over a temperature range of 1 to 400 K. The modelled correlations found provide a good approximation of the diode's I-V characteristic, across its entire temperature range of operation.

To model the diodes' I-V characteristics in PSPICE, an approach analogous to that described in chapter 3 was used. The diode was implemented accordingly by a voltage-controlled current source ($G_{\text{Diode_Main}}$) with the system equation:

$$I(G_{\text{Diode_Main}}) = \begin{cases} V(I_S) \cdot \left(\exp\left(\frac{V(G_{\text{Diode_Main}}) - I(v_{\text{input}}) \cdot R_S}{V(n) \cdot V(V_T)}\right) - 1 \right) & , V(G_{\text{Diode_Main}}) > 0 \\ \frac{V(G_{\text{Diode_Main}})}{\frac{U_{BV}}{I_{BV}}} & , V(G_{\text{Diode_Main}}) \leq 0 \end{cases} \quad (4.9)$$

In this equation $V(G_{\text{Diode_Main}})$ is the voltage across the voltage-controlled current source $G_{\text{Diode_Main}}$. The factors $U_{BV} = 2000 \text{ V}$ and $I_{BV} = 0.01 \text{ A}$ were inherited from the old diode model. These values were chosen, in order to model the diode as a resistor with a value of $R_D = \frac{U_{BV}}{I_{BV}} = 200 \text{ k}\Omega$ in case a negative voltage is applied. The virtual voltage source v_{input} is connected in series with $G_{\text{Diode_Main}}$ and used for its current measurement. The temperature-dependent variables n , I_S and V_T were implemented as voltage nodes. Similar to the implementations introduced in chapter 3, the values of these voltage nodes are defined using voltage-controlled voltage sources (E-components in PSPICE), whose output voltage is given by the equations 4.8, 4.6 and 4.2, respectively.

In this section, the modelling of the diodes' temperature-dependent I-V characteristic was presented. In order to determine the diodes' wafer temperatures during the transient, the diodes must additionally be modelled thermodynamically. This is presented in the next section.

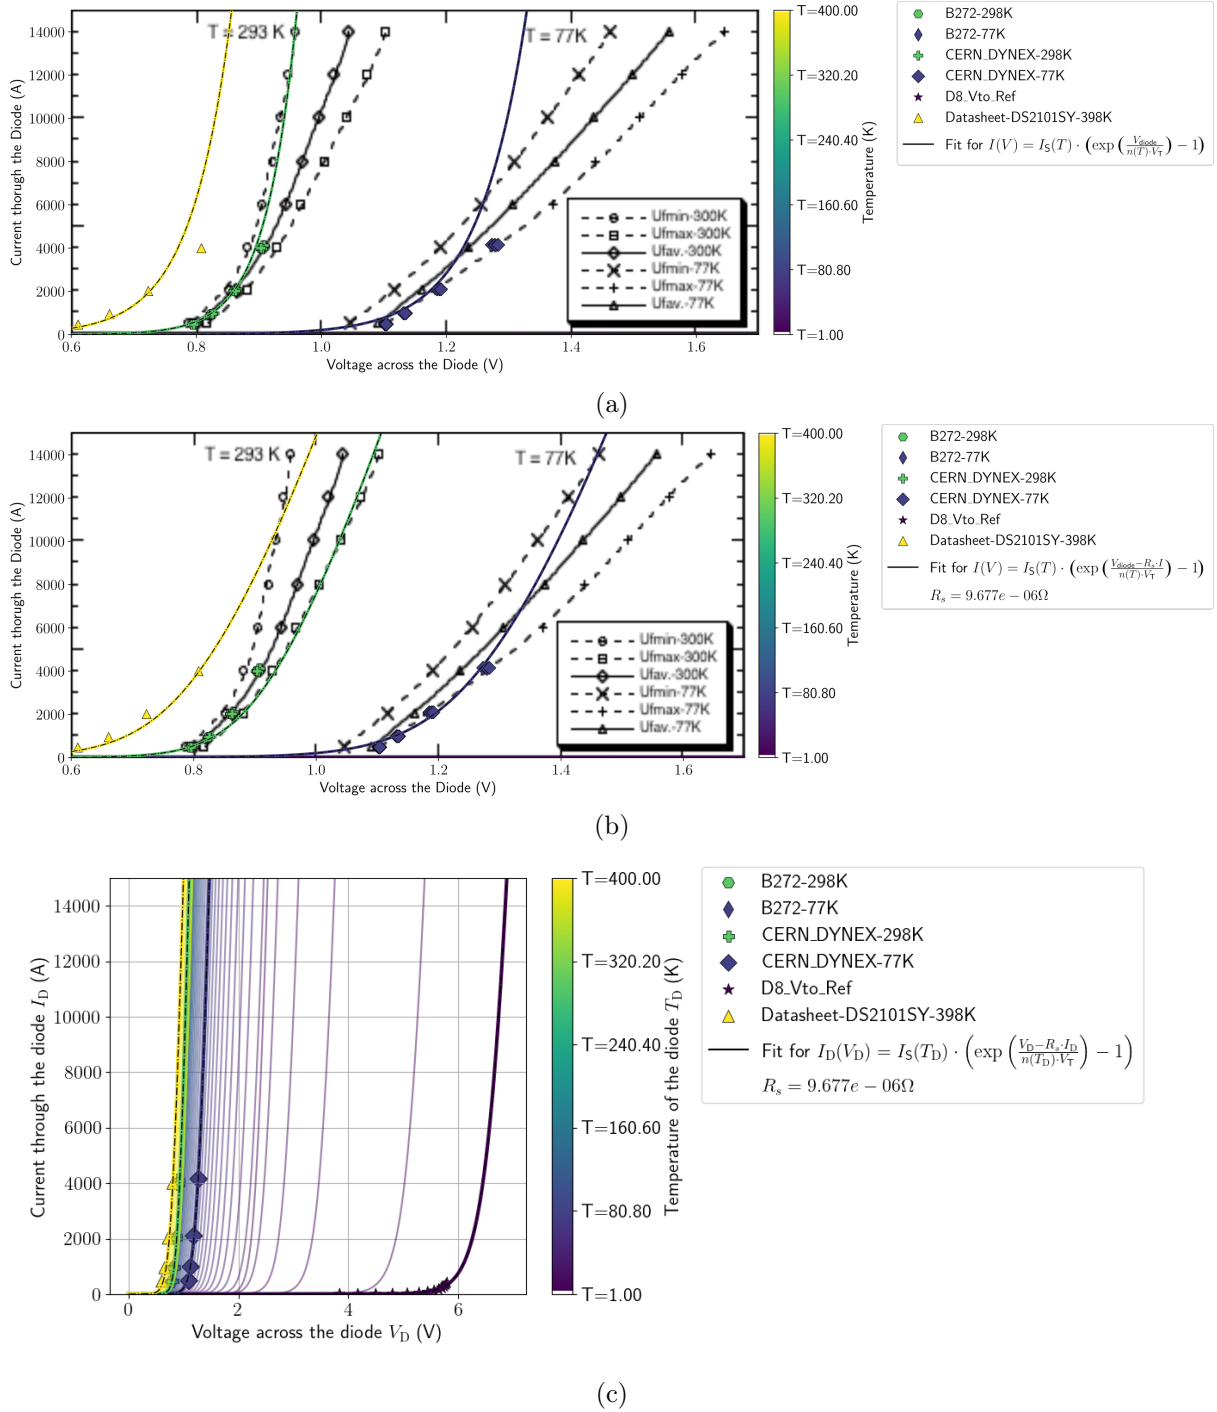


Figure 4.6: a) Comparison of the modelled I-V characteristic derived in this thesis and the measurements of the LHC diodes conducted at ENEA (figures a and b are overlaid with plots from Gharib et al. [42]). a) Comparison to the modelled I-V characteristic without a series resistance. b) Comparison to the modelled I-V characteristic with a series resistance of $R_s = 9.677e^{-6} \Omega$. c) Predictions of the final model of the LHC diode's I-V characteristic for a temperature range of 1 to 400 K.

4.1.2 Thermodynamic modelling of the diodes

The thermal equilibrium of the wafer is highly complex, as it is not only significantly influenced by the heat transfer from the wafer to the diode heat sinks, but also by the cooling effects of the superfluid helium in the cryostat. Furthermore, the temperature and state of the helium in the cryostat are influenced by the energy disposition in the quenching magnets, posing another challenge for accurately modelling the diode's temperature. All models for the thermoelectric simulation of the diodes described in this chapter can therefore only be regarded as highly simplified approximations of the real diode behavior. In this thesis, two different thermodynamic models of the diode body were tested for their suitability to accurately describe the heating and cooling behavior of the silicon wafer during a transient. They are presented in the following sections.

4.1.2.1 Model 1 - constant heat sink temperature

In a first step, an attempt was made to model the thermodynamic equilibrium of the wafer by considering only a single body which is subject to ohmic heating and heat exchange through a thermal coupling to a heat sink of constant temperature. This very simplified model was tested to show whether such trivial modelling, which neglects the effect of the heat sink's increase in temperature, is enough to match the diode's temperature. The thermodynamic equilibrium of the diode's control volume (yellow area in figure 4.7) is given by:

$$\underbrace{P_{\text{ohmic}}}_{=I_D \cdot V_D} + \underbrace{P_{\text{exchange_D-HS}}}_{=-h_{\text{th_D-HS}} \cdot (T_D - T_{\text{HS}})} - \underbrace{P_{\text{Cv_D}}}_{=\frac{d}{dt} C_{\text{v_D}}(T_D) \cdot T_D} = 0. \quad (4.10)$$

(I) (II) (III)

with $P_{\text{exchange_D-HS}}$ being the power associated with the heat exchange between the diode and the heat sink, P_{ohmic} being the power associated with ohmic heating, and $P_{\text{Cv_D}}$ being the power absorbed or released by the diode's heat capacity. These terms are dependent on the current through the diode I_D , the voltage across the diode V_D , the temperature of the diode T_D , the temperature of the heat sinks T_{HS} , the diode's heat capacity at constant volume $C_{\text{v_D}}$ and the thermal conductance between the diode and its heat sinks $h_{\text{th_D-HS}}$.

Furthermore, a scaling factor of the diode's thermal capacitance $f_{\text{Cv_D}}$ is introduced, accounting for imperfections in the assumed heat capacity of the diode. As introduced in section 2.2.2, thermal powers are modelled by currents, while temperatures are modelled with voltages. This leads to the thermodynamic system's equivalent schematic, and its system equations¹ presented in figure 4.7. Within this model, the temperature-dependent heat capacity of the wafer was modelled with a look-up table for silicon, similar to the implementation presented in section 3.1.2. The volume of the wafer was mapped according to its dimensions with $V_{\text{wafer}} = 0.2 \text{ mm} \cdot \pi \cdot (37.5 \text{ mm})^2 \approx 0.884 \text{ cm}^3$.

4.1.2.2 Model 2 - considering heating of the heat sink

In a further step, a second, more refined model was developed to simulate the temperature of two components: the wafer and the heat sink. For this purpose, the heat sink's control volume

¹The term (III) is solved for $(T_D - T_{\text{HS}})$ as presented in section 3.5, in order to model the heat capacity using a voltage-controlled voltage source (PSPICE E-component).

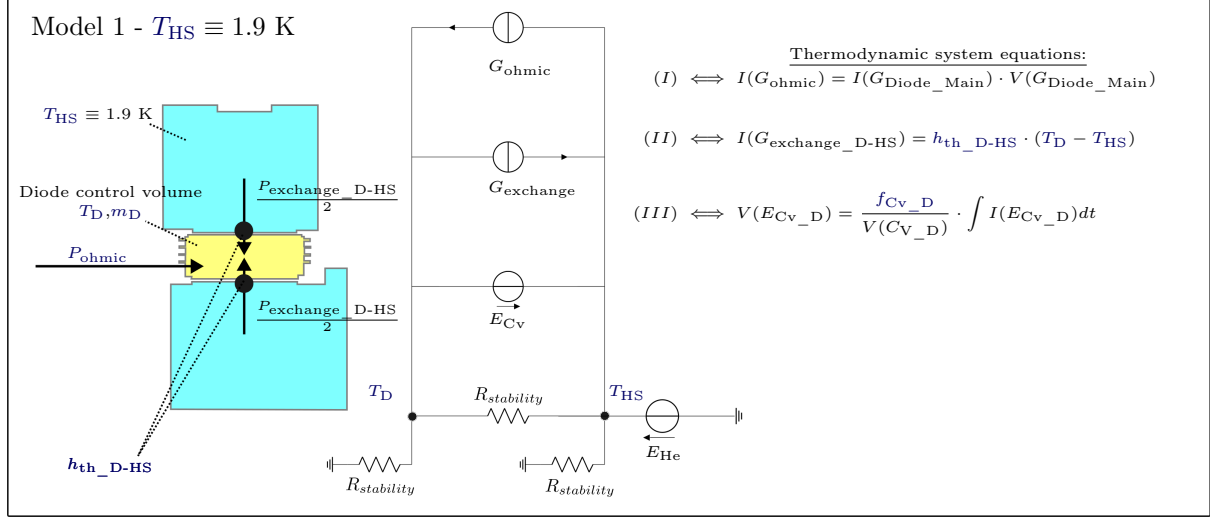


Figure 4.7: Derivation of the diodes' thermodynamic model, considering constant temperature of the heat sink. The direction of the power flow arrows corresponds to the sign convention of equation 4.10. The resistors shown, which were introduced to improve the numerical stability, have a resistance of $1 \text{ G}\Omega$. The voltage source E_{He} has a constant voltage of 1.9 V , a voltage corresponding to the temperature of superfluid helium. The voltage-controlled current source (PSPIICE G-component) $G_{\text{Diode_Main}}$ corresponds to the main component of the diode model, which defines the current through the diode (see equation 4.9). The parameter $f_{\text{Cv_D}}$ is a scaling factor, accounting for imperfections in the assumed heat capacity of the diode.

as shown in figure 4.8 was evaluated based on the first law of thermodynamics. This equilibrium can be written as:

$$\underbrace{P_{\text{exchange_D-HS}}}_{(II)} + \underbrace{P_{\text{exchange_HS-He}}}_{(IV)} - \underbrace{P_{\text{Cv_HS}}}_{(V)} = 0, \quad (4.11)$$

$$= -h_{\text{th_D-HS}} \cdot (T_{\text{HS}} - T_{\text{D}}) = -h_{\text{th_HS-He}} \cdot (T_{\text{HS}} - T_{\text{He}}) = \frac{d}{dt} C_{\text{V_HS}}(T_{\text{HS}}) \cdot T_{\text{HS}}$$

with $P_{\text{exchange_HS-He}}$ being the power associated with the heat exchange between the heat sink and the environment, and $P_{\text{Cv_HS}}$ being the power absorbed or released by the heat sink's heat capacity. These terms are dependent on the temperature of the heat sink's environment $T_{\text{He}} \equiv 1.9 \text{ K}$, the thermal conductance between the heat sink and its environment $h_{\text{th_HS-He}}$, and the heat capacity at constant volume of the heat sink $C_{\text{V_HS}}$. It is important to note that in this equation the power $P_{\text{exchange_D-HS}}$ has the same magnitude as in equation 4.10, but is of opposite sign. Furthermore a scaling factor of the diode heat sink's thermal capacitance $f_{\text{Cv_HS}}$ is introduced, accounting for imperfections in the assumed heat capacity of the heat sink. The thermodynamic system equations of the diode control volume remain similar to the ones introduced for model 1, with the only difference that the heat sink temperature is no longer constant, but dependent on heating and cooling effects, as described in equation 4.11. Here again, the analogies defined in section 2.2.2 are used to translate the five terms into an

Table 4.1: Parameters of the FPA event simulated. In the event under consideration, six individual magnets of the circuit quenched at different current levels.

Circuit current at beginning of quench I_Q (A)	Time after FPA t (s)
11493	-0.009
6509	55.911
6469	56.503
6444	56.854
4447	90.413
2251	150.101

4.2 Results

After the modelling was complete, the models free parameters were fitted³. In a first step, the diode model has been optimized with Dakota, using a very simplified thermal modelling of the diode (model 1). For this purpose, a FPA event with six individual magnet quenches was considered, whose characteristic parameters are given in table 4.1. Within the optimization using the thermal description of model 1, five distinct parameters were to be fitted. These included scaling factors of n , R_S and f_{Is} (f_n , f_{Rs} , f_{Is}) as well as a scaling factor for the diode's heat capacity f_{Cv_D} , and the thermal conductance between diode and heat sink h_{th_D-HS} . However, as the optimization of this model did not lead to a significant improvement in the objective function after 1000 iterations, the attempt to simulate the diode's temperature with the one-body model was rejected. Accordingly, model 2, which considers the heating and cooling of the heat sink, was favoured for modelling the diode's temperature. For the optimization of model 2, two additional parameters were added, including a scaling factor of the heat sink's heat capacitance f_{Cv_HS} , as well as the thermal conductance h_{th_HS-He} between the heat sinks and the helium reservoir. However, the increased amount of free model parameters made it infeasible to fit the model with Dakota in a practical amount of time. For this reason, the model was adjusted manually so that the modelled maximum temperatures are consistent with the temperature measurements carried out at ENEA [42]. In a first step, the parameters of the thermal coupling were manually adjusted to values that represented a realistic approximation of the possible diode behavior. After that, it was possible to determine a trendline relating the diodes' simulated maximal temperatures as a function of their maximum currents, as shown in figure 4.9. Based on this trendline, the diode's thermal capacitance, which in the model influences the maximum temperature the most, was adjusted. By reducing f_{Cv_D} to 0.7, the trendline obtained is coinciding with the 280 K maximum diode temperature measured at 13 kA at ENEA [42]. As presented below, the model parameters found (which can be used as a possible starting point for a deeper optimization in a future work) already allow promising qualitative

³To improve the results of the optimization, a filter for the measurement data was programmed. This is necessary, since the diode forward voltage measurement data obtained from PM always represents a mixture of the voltage across the diode with and without its current leads. Since the implemented diode model does not take the current leads into account, the measurement data must be filtered before the optimization, so that only the voltage directly at the diode is considered as a reference. This filtering is illustrated in figure A.4 of the appendix.




Figure 4.9: Maximum diode temperature as a function of the maximum diode current. The trendline of the simulated maximum temperature is extrapolated and coincides with the maximum diode temperature of 280 K, measured at 13 kA at ENEA [42].

matches between the measurements and simulation results.

The obtained forward voltages over time for this parameter configuration are shown in figure 4.10. It can be noted that the simulated opening voltage of all diodes lies in a range between 5 and 6 V, which is in good agreement with the measurements at 4.2 K, shown in figure 4.3. It can further be observed that despite experiencing similar maximum currents, the diodes' measured forward voltages over time may vary, as diode 29 exhibits a second increase in forward voltage after around 400 s, whereas no such increase can be observed for diode 126. The simulated forward voltages of the diodes 29, 31, 124 and 125 exhibit increased deviations from the measurement data of in some cases more than 20 %. For diode 30, which experienced an initial current of approximately 11493 A, and diode 124, which opened at approximately 2251 A, a better agreement between measurement and simulation is observed. A temporary drop below the measured forward voltages can be observed for all diodes shortly after the initial peak. This temporary undershoot could be due to an overestimation of the simulated diode wafer temperature. However, the rough trend of the simulated voltage appears to be correct. Just like in the measured data, the diodes' initial voltage peak is followed by a period of smaller changes in forward voltage. Depending on how much the diode and its heat sink have heated up during the discharge, which depends on the current at which the respective magnet quenches, the diode may cool down enough for the forward voltage to rise again. Such a second peak in forward voltage can be observed in both the measured and simulated forward voltage of diode 124. The simulated second peak in forward voltage at $\hat{t} = 385$ s is delayed by about 10 s, and of slightly smaller magnitude in comparison to the measurements. Towards the end of the simulation, at around 420 s, an abrupt drop in the simulated forward voltage can be observed for all diodes, which differs greatly from the decay observed in the measurement. This could be traced back to errors in the circuit current, possibly caused by a bug in the programming of the circuit's power converters. It is possible that a more detailed model of the powerconverter's thyristors at very




Figure 4.10: Forward voltages of diodes that are connected to quenching magnets. The model parameters were fitted manually, so that the maximum diode temperature simulated would extrapolate to measurement data conducted at ENEA. [42].




Figure 4.11: Simulated temperatures over time after the FPA corresponding to diodes that were connected to quenching magnets. The model parameters were fitted manually, so that the maximum diode temperature simulated would extrapolate to measurement data from ENEA [42].

low currents could improve the simulation at late stages of this transient. However, reproducing this feature is outside the scope of this thesis, as the protection and performance of the circuit is not critical at such low currents.

The simulated temperatures of each diode and its respective heat sink are given in figure 4.11. For all simulated temperatures, an initial sharp rise up to a maximum temperature, followed by a steady decline can be observed. For low temperatures, the simulated decrease in temperature exhibits an approximately linear course over time. The heat sink temperature follows a course similar to the diode temperature, but with much lower rates of change due to the higher thermal capacity of the heat sinks.

Towards the end of each diode's recooling process, an abrupt drop in the diode and heat sink temperature can be observed. Considering the temperature curves of diode 124, it is noticeable that the time of this abrupt drop ($\hat{t} = 385$ s) coincides with the simulated second peak of the diode's forward voltage.

In the following, the simulated currents through the diodes and the magnets are presented, based on figure 4.12. The currents through the magnets before their quench, as well as the currents through the diodes after the magnet attached to them quenched, follow the course of the measured circuit current (obtained from PM) as expected. This match of the currents is a result of the improved energy extraction resistor model, introduced in chapter 3.

As can be observed for diode 124 in the period between 150 and 160 s, a redistribution of the current from the magnet to the diode takes place after the quench. Furthermore, a reversal of




Figure 4.12: Simulated currents through the diodes and the quenching magnets connected to them as a function of time after the FPA. The measured circuit current I_{MEAS} , obtained from the PM database is also plotted for comparison. The currents of the diodes are marked with the index D and the currents through the magnets with an index MB.




Figure 4.13: Simulated resistances of the diodes and the quenching magnets connected to them as a function of time after the FPA. The resistances of the diodes are marked with the index D and the resistance of the magnets with an index MB.

this process can be observed for this diode, in which the diode gradually closes and the current is transferred back to the magnet.

The current redistribution from the diode back to the magnet is completed exactly at $\hat{t} = 385$ s, i.e. the point in time at which the second peak in the diode's forward voltage was observed (see figure 4.10).

In the following, the redistribution processes between diode and magnet are further explained with the course of the resistance of these components, which is given in figure 4.13. The time and current-dependent magnet resistances were determined by cooperative simulations as part of an earlier work by Janitschke [38]. A reduction of the diode effective resistance (i.e. $\frac{V_D}{I_D}$) can be observed immediately after the diodes heat up. As the diodes cool however, their resistance is continuously increasing until it finally surpasses the resistance of the magnet⁴. The resulting reduction in diode current causes an even faster cooling of the diode, until its resistance finally increases very rapidly causing the diode to close completely (except for leakage currents). This point in time ($\hat{t} = 385$ s) coincides with the time of the simulated second peak of the diode forward voltage.

⁴It shall be noted that the diode effective resistance (i.e. $\frac{V_D}{I_D}$), would also increase, if T_D would be constant. This is due to the fact that the circuit current I_{MEAS} decreases semi-exponentially over time, while the forward voltage V_D remains almost constant over a wide current interval, once the respective magnet has quenched and the diode has heated up to a certain amount.

4.3 Summary

In this chapter, an electrothermal model for the LHC main dipole circuit's by-pass diodes was derived. To model the LHC diode's temperature-dependent I-V characteristic, the Shockley equation was fitted to measurement data recorded at different temperatures. As it was not possible to model the temperature dependency of the diode's series resistance due to its high scattering and the sparse measurement data available for temperatures below 77 K, the series resistance was initially neglected. Hence, the Shockley equation was first fitted to measurement data at lower currents in order to minimize the distortions associated with this simplification. Afterwards, fitting functions were established, based on the ideality factors n and reverse-bias saturation currents I_S determined at different temperatures in order to model their temperature dependence. As physical laws could not represent the temperature dependence of the data points found for n and I_S well enough, other suitable functions were used for fitting which however do not have a physical foundation. In the last step of the temperature-dependent I-V modelling, the diode's series resistance, was taken into account with a constant value. The modelled diode series resistance was manually adjusted so that the obtained I-V characteristic would always align with I-V measurement data. Although the correlations found might not be applicable outside the range of operating conditions considered in this thesis, they allowed the diode's modelling to a good approximation.

In order to model how the diodes are heating up due to the current flow through them, two different thermodynamic equivalent models of the diodes were considered: a one-body model, considering only temperature changes in the diode body itself and a two-body model which in addition also considers heating and cooling effects of the diode heat sinks.

The thermal behavior of the main dipole circuit's by-pass diodes after the respective parallel-connected magnet quenched, is characterized by initially rapid heating which is followed by comparatively slow cooling. While the diode's initial rapid heating can be easily modelled by a single body, its cooling is too complex to be modelled by the one-body model (model 1). Hence, the increase in temperature of the heat sinks must be modelled, in order to better match the cooling of the diode, which is more powerful at the beginning of the transient (when the heat sink is still cold) and less powerful in later stages of the transient (when the heat sink's temperature is increased). For this reason, the two-body model (model 2) was preferred.

As the increased amount of free model parameters could not be fit with Dakota in a practical amount of time, the model was adapted manually to measured data, based on the simulated maximum temperatures of the diodes.

In this way a diode model was derived that allows a qualitative and semi quantitative estimation of the LHC diodes' post-quench voltage, resistance, current, and temperature behavior. Additionally, the occurrence of post-quench recooling related peaks in the forward voltage of the diodes could be simulated and explained. It was observed that the occurrence of the peaks is accompanied by a sharp drop in temperature and a greatly increased resistance of the diodes.

In some cases the measured forward voltage profiles of diodes attached to magnets that quenched at similar times, were observed to vary significantly. This could indicate differences in the cooling of the diodes, which could be manufacturing-related or caused by flows of gaseous helium coming from adjacent cryostats.

For example, previously quenching magnets that are in an adjacent cryostat part of the same cryo cell as the diode influence the diode's temperature.

Overall, the existing [STEAM](#) model of the [LHC](#) main dipole magnets' by-pass diodes has been significantly improved. It will therefore contribute to an enhanced accuracy of magnet simulations and a more precise analysis of quench-related FPA events.

Conclusion and outlook

The **L**arge **H**adron **C**ollider (LHC) comprises eight individual main dipole circuits (RB), each containing 154 superconducting dipole magnets. These magnets are used to keep the particle beams on their circular paths and make up the largest part of the LHC in terms of length. Simulations of the main dipole circuits are an essential tool for the safe operation of the LHC. For example they are used, to analyze **F**ast **P**ower **A**bort (FPA) events, which are triggered in case a disturbance (such as a power-supply trip, equipment malfunctioning, or a magnet quench) in the main dipole circuit is detected. During an FPA event, the circuit's power supply is switched off and the circuit discharges, causing transient effects in the main dipole circuit to occur. Understanding these transient effects is important to verify the circuit performance and for the analysis of unexpected events occurring in these circuits. Since the physical mechanisms determining the behavior of superconducting magnet circuits span different scales, it is not practical to exclusively simulate them using conventional approaches like FEM, as the simulation would become computationally too expensive.

However, these challenging simulations can be carried out with the **S**imulation of **T**ransient **E**ffects in **A**ccelerator **M**agnets (STEAM) framework, developed at CERN. This simulation framework includes (in-house developed) software packages, APIs to automatically generate models and run simulations, and libraries of superconducting magnets and circuits.

The accuracy of simulations of the main dipole circuits is heavily dependent on the models of the circuits' individual subcomponents. For example, highly accurate simulations of individual magnets require accurate modelling of the current through and the voltage across the magnet. These two quantities are significantly influenced by two different types of circuit components: the energy extraction resistors and the magnets' by-pass diodes. The electrical behavior of these two components is temperature-dependent. However, the effect of their temperature was previously not taken into account in the component models, which led to significant deviations between simulated and measured circuit signals, particularly for FPA events that were triggered at high currents.

The aim of this thesis was therefore to model the electrothermal behavior of these components and to implement and use a new method to automatically fit the models' free parameters to measurement data.

To model the heating and cooling effects of these components, analog behavioral models were implemented in PSPICE netlists. In this approach, analogies between the thermal domain and the electrical domain are used to simulate thermodynamic systems using equivalent electrical circuits. Furthermore, an already existing software interface between STEAM and the parameter optimization toolkit Dakota (developed at Sandia National Laboratories) was further developed in order to fit the models to measurement data. To this end, the interface was extended to perform so-called closed-loop optimizations of the netlist models. In this approach, simulation results are iteratively compared with measurement data in a closed-loop fashion. The choice of parameters used in the next iteration's simulation is informed by the results of previous iterations, causing closed-loop optimization to be quick and efficient. Since several signals (e.g.,

voltages across different energy extraction units) are optimized simultaneously, i.e., minimizing several error metrics at once, the method applied in this thesis is referred to as closed-loop multi-objective optimization.

The new energy extraction resistor model was implemented with two distinct free parameters: h_{th_EE} , which is proportional to the resistor's cooling and f_{Cv_EE} , used to scale the resistors initially assumed heat capacity up or down. The model was tested by simulating three different FPA events, occurring at three different circuit currents: 11500 A, 5923 A and 3499 A.

By fitting the model's free parameters with Dakota, a significant improvement in simulation accuracy could be achieved, especially at high FPA currents. In the course of model refinement it was observed that the consideration of current-dependent magnet differential inductance due to saturation of the iron-yoke around the magnet coils (previously modelled as constant) causes a further improvement in simulation accuracy.

In all optimizations, the scaling factor f_{Cv_EE} was reduced by the algorithm, which is attributed to the effect of uneven heating of the resistor. The optimum thermal conductance h_{th_EE} notably scattered for different optimizations, but often converged to values corresponding to virtually zero cooling.

On the contrary, the optimum value $f_{Cv_EE} = 0.77$, determined at $I_0 = 11500$ A is consistent across different starting points of the optimization.

Since the effects of resistor heating were observed to be most pronounced at high FPA currents, the model was deployed with the optimum parameters found for the event that occurred at the highest circuit current. The optimum parameters of these events were $h_{th_EE} = 91.47$ W K⁻¹ and $f_{Cv_EE} = 0.77$.

With these model parameters, an improvement in the simulated resistance across all tested FPA current levels of up to 14 % could be achieved (errors are normalized to the maximum absolute value of the measured quantity). This also improved the accuracy of the simulated circuit current, which showed an improvement in simulation accuracy of up to 2 %.

Furthermore, the fitted model allows to reproduce the experimentally-derived temperature increase ΔT_{calc} , which is the temperature increase needed to account for the measured resistance increase of the energy extraction resistors (assuming homogeneous resistivity, heat capacity and temperature across the resistor). The measured resistor temperature can not be reproduced which is due to the model's underlying assumption of a homogeneous temperature distribution across the resistor. This is however beyond the scope of this thesis, since the resistor temperature does not directly affect the transients in the superconducting magnets.

In a future work, the model could be further refined in order to model temperature gradients in the resistor. This would allow a better match of simulated and measured temperatures than it is possible with the model derived in this thesis, which primarily focussed on the accurate simulation of electrical signals.

Furthermore, following the implementation of the improved energy extraction model, a new model of the RB by-pass diodes was derived. To model the diodes' temperature-dependent I-V characteristic, the Shockley equation was fitted to measurement data recorded at different temperatures. Afterwards, fitting functions were established based on the ideality factors n and

reverse-bias saturation currents I_S determined at different temperatures in order to model their temperature dependence. As physical laws could not represent the temperature dependence of the data points found for n and I_S well enough, other suitable functions were used for fitting which however do not have a physical foundation. The final model obtained was adapted to align with I-V measurement data. Although the correlations found might not be applicable outside the range of operating conditions considered in this thesis, they allowed the diode's modelling to a good approximation.

Two different thermodynamic equivalent models of the diodes were considered in order to model its heating: a one-body model, considering only temperature changes in the diode body itself, and a two-body model which in addition also considers heating and cooling effects of the diode's heat sinks. The thermal behavior of the main dipole circuit's by-pass diodes after the respective parallel-connected magnet quenched is characterized by initially rapid heating which is followed by comparatively slow cooling. While the diode's initial rapid heating can be easily modelled by a single body, its cooling is too complex to be modelled by the one-body model. Hence, the increase in temperature of the heat sinks must be modelled in order to better match the cooling of the diode, which is more powerful at the beginning of the transient (when the heat sink is still cold) and less powerful in later stages of the transient (when the heat sink's temperature is increased). For this reason, the two-body model was preferred. As the increased amount of free model parameters associated with the two-body model could not be fit with Dakota in a practical amount of time, the model was adapted manually to measurement data, based on a comparison of simulated and measured maximum temperatures of the diodes.

In this way a diode model was derived that allows a qualitative and semi-quantitative estimation of the LHC diodes' post-quench voltage, resistance, current, and temperature behavior.

The newly implemented diode model allowed to qualitatively explain the occurrence of post-quench recooling-related peaks in the forward voltage of the diodes. These were observed to be accompanied by a sharp drop in the diodes' temperature and a highly increased resistance of the diodes.

In some cases the measured forward voltage profiles of diodes attached to magnets that quenched at similar times, were observed to vary significantly. This could indicate differences in the cooling of the diodes, which could be manufacturing-related or caused by flows of gaseous helium coming from adjacent cryostats. For example, previously quenching magnets that are in an adjacent cryostat part of the same cryo cell as the diode influence the diode's temperature. In a future work, the thermal modelling of the diode could therefore be supplemented by considering such preheating effects. Furthermore, each diode could be characterized individually in a future work, by determining a diode specific set of simulation parameters, using the STEAM-Dakota interface.

As part of this thesis, the STEAM-Dakota interface was supplemented with regard to the possibility to optimize the models of circuit components. The interface was prepared for conducting closed-loop optimizations, where model parameters are iteratively improved to better align simulation results to measurement data. Furthermore, a strategy was developed for managing the risks of faulty or frozen simulations which could occur during the optimization process. This strategy consisted of introducing a timeout after which frozen simulations are aborted, and assigning a penalty value to parameter vectors that do not lead to convergence. Following this

approach, the interface’s capabilities for closed-loop optimizations were showcased in several applications. The coliny pattern search algorithm was identified as particularly robust for the applications presented in this thesis, as it does not rely on numerical gradients, which are lost in case of frozen simulations.

In a future work, the STEAM-Dakota interface could be used for the characterization of individual circuit components, by fitting the same model to measurements from physically different components (e.g., diodes that are connected to different magnets). In this way, each component physically present in the circuit could be provided with its own set of model parameters. This would not only further improve the simulation accuracy, but also enable the identification of outliers, indicating potentially faulty behavior of the components.

The circuit component models developed as part of this thesis significantly improved the simulation accuracy of the LHC main dipole circuits. The models are based on a more complete set of physical effects, and are thus expected to be more robust for simulating unexpected events for which experimental data are currently not available. Furthermore, the new models provide approximations of unmeasured quantities of the circuit components, such as the diode’s temperature during an FPA, and can thereby contribute to a better understanding.

The application of the STEAM-Dakota interface for optimizing STEAM circuit component models was showcased for the first time. The good transferability of the applied method to other models ensures that this approach will be a promising option for further improving circuit simulation accuracy in future works.

In the future, the STEAM framework will thus be more accurate and reliable for research and development related tasks like hardware commissioning, the analysis of FPA events and the further development of the [LHC](#) superconducting magnet circuits.

Appendix

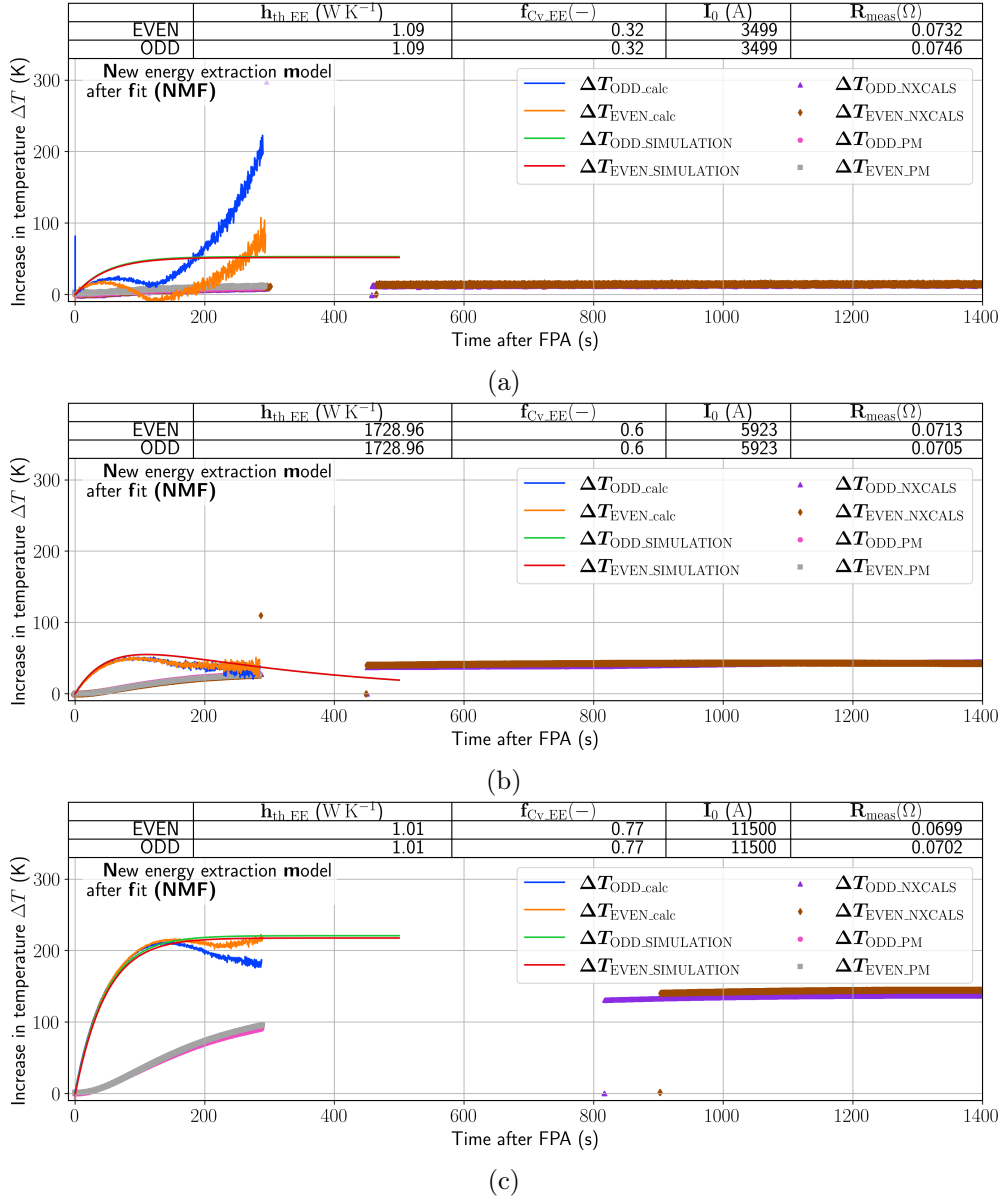


Figure A.1: Increase in temperature as simulated with the ABM in comparison to measurement data. The plots correspond to different FPA events, which occurred at different current levels I_0 .

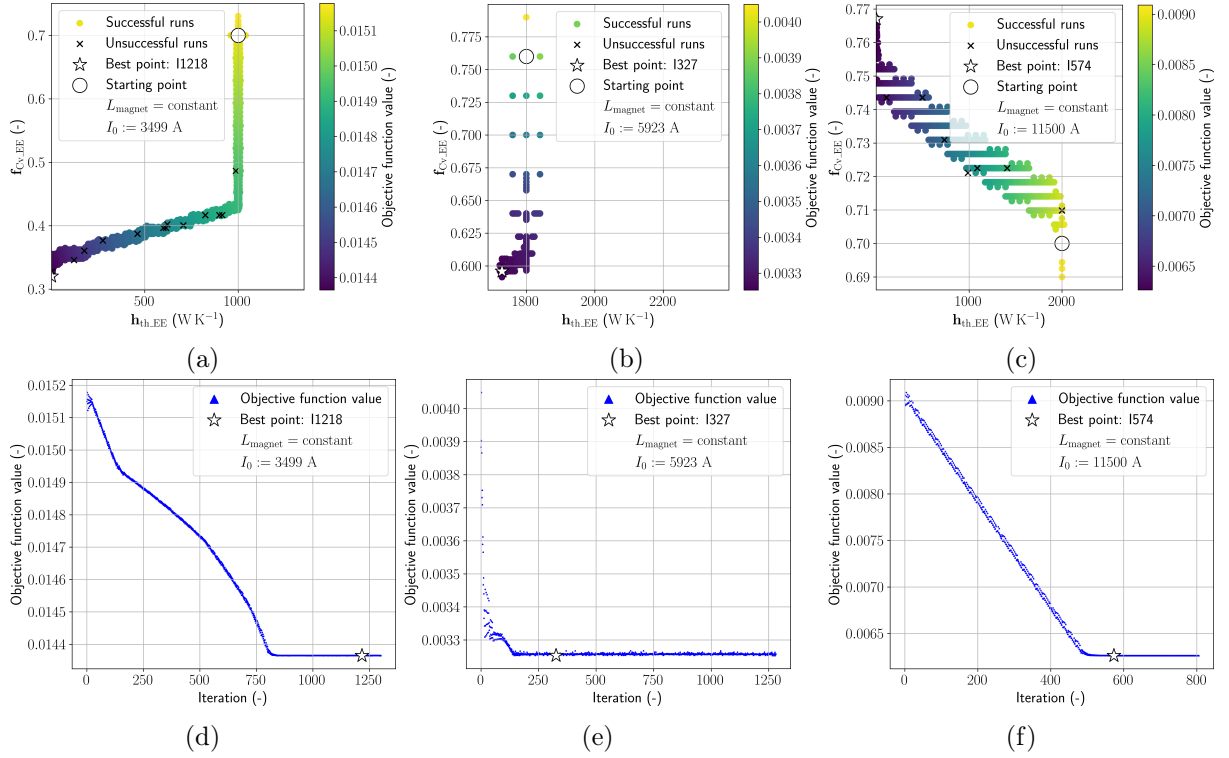


Figure A.2: Parameter optimization of the newly developed EE system model using a pattern search algorithm. The plots correspond to different FPA events, which occurred at different current levels I_0 . a - c) Objective function landscapes showing the initial points and the best points of the optimizations as well as the movement of the tested points within parameter space. d - f) Value of the objective function over the course of the optimizations. The magnet differential inductance was considered as constant in these optimizations.

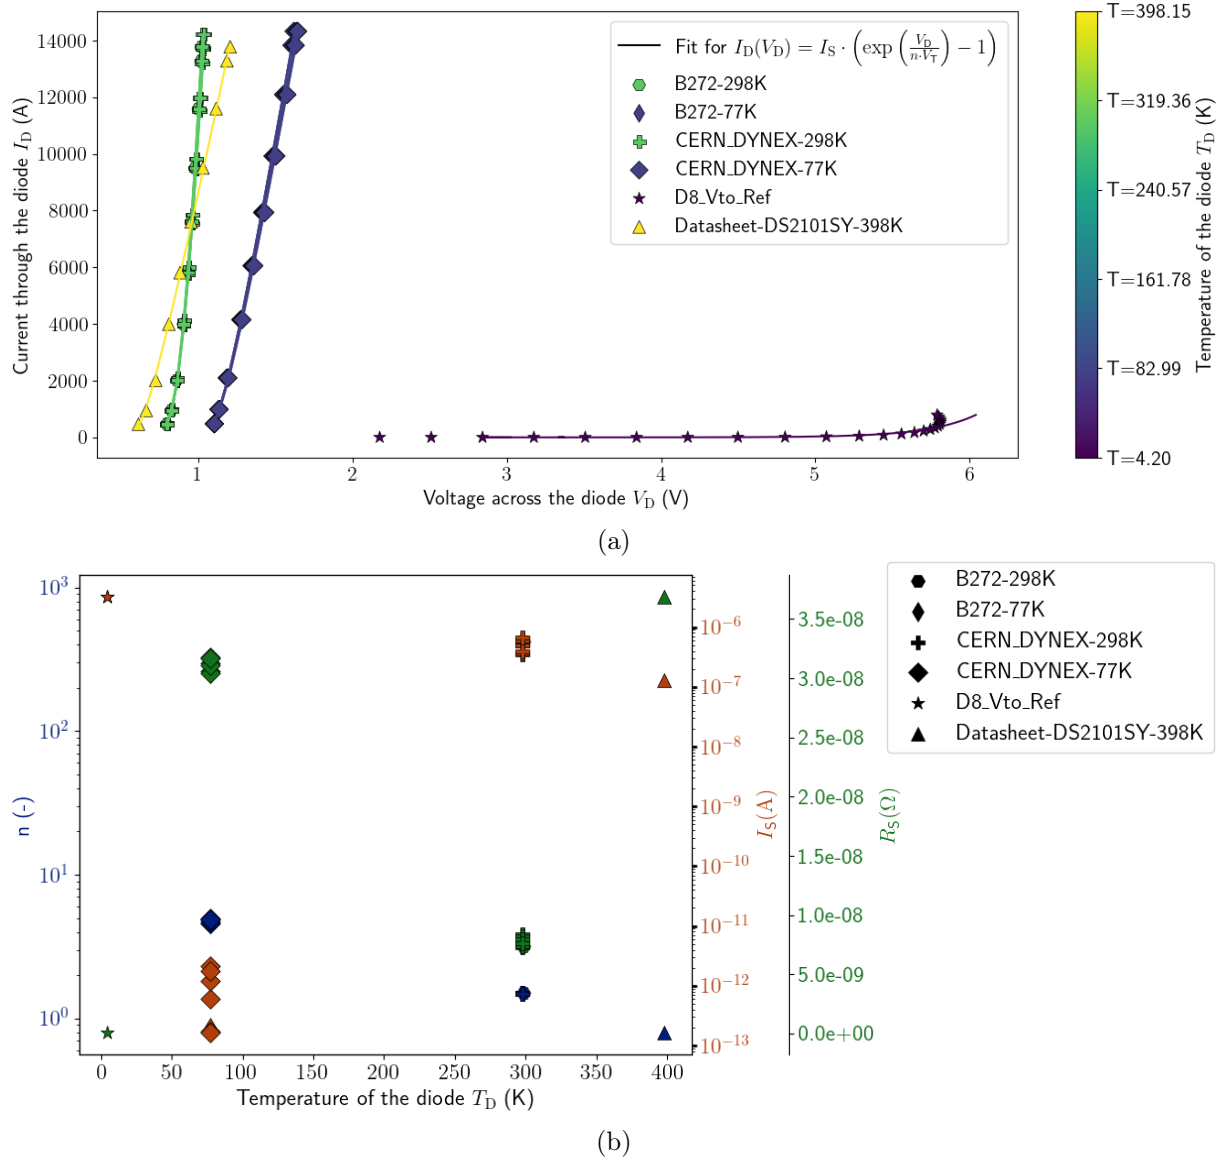


Figure A.3: Fit of equation 4.5 to measurement data of the LHC diode's forward characteristic. a) The fits to the measurement data exhibit a well representation. b) The fitting parameters found, show non-monotonic, seemingly arbitrary temperature dependence, which is why this approach was discarded for modeling.




Figure A.4: Filter applied to PM diode forward voltage measurement data. The PM measurement data alternates between measurements of the voltage across the diode with and without current leads. Since the implemented diode model does not take the current leads into account, the measurement data must be filtered for the optimization so that only the voltage directly at the diode is considered as a reference. This was achieved by repeatedly skipping certain data intervals. After about 372 s, the current in the circuit is so low that the voltages hardly differ. Hence, after 372 s the filter was disabled, in order to be able to better map the subsequent peaks.

Bibliography

- [1] O. Brüning *et al.*, “LHC Design Report,” CERN, Geneva, Tech. Rep., 2004. DOI: [10.5170/CERN-2004-003-V-1](https://cds.cern.ch/record/782076). [Online]. Available: <https://cds.cern.ch/record/782076>.
- [2] J. Schukraft, “Little bang at big accelerators: Heavy ion physics from AGS to LHC,” in *Nuclear Physics B - Proceedings Supplements*, vol. 75, ELSEVIER, Mar. 1999, pp. 46–53. DOI: [10.1016/S0920-5632\(99\)00214-5](https://www.sciencedirect.com/science/article/pii/S0920563299002145). [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0920563299002145>.
- [3] F. Gianotti and T. S. Virdee, “The discovery and measurements of a Higgs boson,” *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, vol. 373, Jan. 2015. DOI: [10.1098/rsta.2014.0384](https://royalsocietypublishing.org/doi/10.1098/rsta.2014.0384). [Online]. Available: <https://royalsocietypublishing.org/doi/10.1098/rsta.2014.0384>.
- [4] CERN, *High-Luminosity LHC Website*. [Online]. Available: <https://home.cern/science/accelerators/high-luminosity-lhc> (visited on 02/24/2024).
- [5] B. Schmidt, “The High-Luminosity upgrade of the LHC: Physics and Technology Challenges for the Accelerator and the Experiments,” *Journal of Physics: Conference Series*, vol. 706, 2016. DOI: [10.1088/1742-6596/706/2/022002](https://iopscience.iop.org/article/10.1088/1742-6596/706/2/022002). [Online]. Available: <https://iopscience.iop.org/article/10.1088/1742-6596/706/2/022002>.
- [6] L. Rossi and O. Brüning, “Progress with the High Luminosity LHC project at CERN,” in *10th International Particle Accelerator Conference*, JACoW Publishing, 2019, pp. 17–22, ISBN: 978-3-95450-208-0. DOI: [10.18429/JACoW-IPAC2019-MOYPLM3](https://cds.cern.ch/record/2696270?ln=de). [Online]. Available: <https://cds.cern.ch/record/2696270?ln=de>.
- [7] CERN, *The CERN accelerator complex*. [Online]. Available: <https://www.home.cern/science/accelerators/accelerator-complex> (visited on 02/24/2024).
- [8] CERN, *The HL-LHC project*. [Online]. Available: <https://hilumilhc.web.cern.ch/content/hl-lhc-project> (visited on 02/24/2024).
- [9] P. Mangin and R. Kahn, *Superconductivity An introduction*. Springer International Publishing AG, 2017, ISBN: 9783319505275. DOI: [10.1007/978-3-319-50527-5](https://link.springer.com/book/10.1007/978-3-319-50527-5).
- [10] W. Demtroeder, *Experimentalphysik 3*, 5th ed. Berlin, Heidelberg: Springer Spektrum, 2016, ISBN: 978-3-662-49093-8. DOI: [10.1007/978-3-662-49094-5](https://link.springer.com/book/10.1007/978-3-662-49094-5). [Online]. Available: <https://link.springer.com/book/10.1007/978-3-662-49094-5>.
- [11] H. Kamerlingh Onnes, K. Gavroglu, and Y. Goudaroulis, *Through Measurement to Knowledge*. Dordrecht, The Netherlands: Kluwer Academic Publishers, 1991, ISBN: 9789401074339. DOI: [10.1007/978-94-009-2079-8](https://link.springer.com/book/10.1007/978-94-009-2079-8).
- [12] W. Buckel and R. Kleiner, *Superconductivity: Fundamentals and Applications*, 2nd ed. Weinheim: WILEY-VCH Verlag GmbH & Co. KGaA, 2004, ISBN: 3-527-40349-3.
- [13] Y. Slimani and E. Hannachi, *Superconducting Materials: Fundamentals, Synthesis and Applications*. Springer Nature Singapore Pte Ltd, 2022, ISBN: 978-981-19-1211-5. DOI: [10.1007/978-981-19-1211-5](https://link.springer.com/book/10.1007/978-981-19-1211-5). [Online]. Available: <https://link.springer.com/book/10.1007/978-981-19-1211-5>.

- [14] H. Kamerlingh Onnes, "Further experiments with liquid helium," in *Proceedings of the Royal Netherlands Academy of Arts and Sciences (KNAW)*, vol. 16, Amsterdam, 1914, pp. 987–992. [Online]. Available: <https://dwc.knaw.nl/DL/publications/PU00012935.pdf>.
- [15] R. G. Sharma, *Superconductivity: Basics and Applications to Magnets*. Springer International Publishing Switzerland, 2015, ISBN: 978-3-319-13713-1. DOI: [10.1007/978-3-319-13713-1](https://doi.org/10.1007/978-3-319-13713-1).
- [16] A. A. Abrikosov, "On the Magnetic Properties of Superconductors of the Second Group," *SOVIET PHYSICS JETP*, vol. 5, no. 6, pp. 1174–1182, 1957.
- [17] D. Meschede, *Gerthsen Physik*. D. Meschede and Springer Spektrum, 2015, vol. 25, ISBN: 978-3-662-45977-5. DOI: [10.1007/978-3-662-45977-5](https://doi.org/10.1007/978-3-662-45977-5).
- [18] I. A. Parinov, *Microstructure and properties of high-temperature superconductors*. Springer-Verlag Berlin Heidelberg, 2007, ISBN: 9783540709763. DOI: [10.1007/978-3-540-70977-0](https://doi.org/10.1007/978-3-540-70977-0).
- [19] J. Livingston, "Magnetic Properties of Superconducting Lead-Base Alloys," *Physical Review*, vol. 129, no. 5, pp. 1943–1949, 1963. DOI: [10.1103/PhysRev.129.1943](https://doi.org/10.1103/PhysRev.129.1943).
- [20] M. R. Beasley, "A History of Superconductivity," in *Advances in Superconductivity: Proceedings of the 1st International Symposium on Superconductivity*, Springer, Tokyo, 1989. DOI: [10.1007/978-4-431-68084-0_1](https://doi.org/10.1007/978-4-431-68084-0_1).
- [21] J. Bardeen, L. N. Cooper, and J. R. Schrieffer, "Theory of Superconductivity," *Physical Review*, vol. 108, no. 5, pp. 1175–1204, Dec. 1957. DOI: [10.1103/PhysRev.108.1175](https://doi.org/10.1103/PhysRev.108.1175). [Online]. Available: <https://link.aps.org/doi/10.1103/PhysRev.108.1175>.
- [22] M. N. Wilson, "Superconductivity and accelerators: The good companions," *IEEE Transactions on Applied Superconductivity*, vol. 9, no. 2, pp. 111–121, 1999, ISSN: 1558-2515. DOI: [10.1109/77.783250](https://doi.org/10.1109/77.783250). [Online]. Available: <https://ieeexplore.ieee.org/document/783250>.
- [23] L. Rossi, "Superconductivity: Its role, its success and its setbacks in the Large Hadron Collider of CERN," *Superconductor Science and Technology*, vol. 23, no. 3, 2010. DOI: [10.1088/0953-2048/23/3/034001](https://doi.org/10.1088/0953-2048/23/3/034001).
- [24] M. N. Wilson, *Presentation: Superconducting Magnets for Accelerators*, 2006. [Online]. Available: <https://cas.web.cern.ch/sites/default/files/lectures/zakopane-2006/wilson-lect.pdf> (visited on 07/03/2024).
- [25] E. Ravaioli, "CLIQ - A new quench protection technology for superconducting magnets," PhD Dissertation, University of Twente, 2015, ISBN: 9789036539081. [Online]. Available: <https://cds.cern.ch/record/2031159?ln=de>.
- [26] D. Larbalestier, A. Gurevich, D. Matthew Feldmann, and A. Polyanskii, "High-Tc superconducting materials for electric power applications," *Nature*, vol. 414, no. 6861, pp. 368–377, 2001. DOI: [10.1038/35104654](https://doi.org/10.1038/35104654).
- [27] Q. Guo *et al.*, "Study on High Residual Resistance Ratio (RRR) of Wire in Channel Superconductor Wire," *IEEE Transactions on Applied Superconductivity*, vol. 29, no. 5, 2019. DOI: [10.1109/TASC.2019.2908275](https://doi.org/10.1109/TASC.2019.2908275). [Online]. Available: <https://ieeexplore.ieee.org/document/8676359>.

- [28] G. Willering, “Stability of superconducting rutherford cables for accelerator magnets,” Dissertation, University of Twente, The Netherlands, 2009, ISBN: 9789036528177.
- [29] H. Bajas, D. Durville, D. Ciazynski, and A. Devred, “Numerical simulation of the mechanical behavior of ITER cable-in-conduit conductors,” *IEEE Transactions on Applied Superconductivity*, vol. 20, no. 3, pp. 1467–1470, Jun. 2010. DOI: [10.1109/TASC.2010.2042944](https://doi.org/10.1109/TASC.2010.2042944).
- [30] G. Moritz, “Rapidly-cycling superconducting accelerator magnets for FAIR at GSI,” in *Proceedings of the IEEE Particle Accelerator Conference*, 2007, pp. 3745–3749. DOI: [10.1109/PAC.2007.4440118](https://doi.org/10.1109/PAC.2007.4440118).
- [31] L. Bottura, “Cable Stability,” in *CAS - CERN Accelerator School: Superconductivity for Accelerators*, 2014. DOI: [10.5170/CERN-2014-005.401](https://doi.org/10.5170/CERN-2014-005.401).
- [32] A. P. Verweij, “Eletrodynamics of Superconducting Cables in Accelerator Magnets,” PhD Dissertation, University of Twente, 1995, ISBN: 90-90-08555-6. DOI: [10.3990/1.9789090085555](https://doi.org/10.3990/1.9789090085555). [Online]. Available: <https://research.utwente.nl/en/publications/eletrodynamics-of-superconducting-cables-in-accelerator-magnets>.
- [33] G. H. Morgan, “Theoretical Behavior of Twisted Multicore Superconducting Wire in a Time-Varying Uniform Magnetic Field,” *Journal of Applied Physics*, vol. 41, no. 9, pp. 3673–3678, Aug. 1970. DOI: [10.1063/1.1659491](https://doi.org/10.1063/1.1659491).
- [34] N. Magnusson, A. B. Abrahamsen, D. Liu, M. Runde, and H. Polinder, “Hysteresis losses in MgB2 superconductors exposed to combinations of low AC and high DC magnetic fields and transport currents,” *Physica C: Superconductivity and its Applications*, vol. 506, pp. 133–137, Nov. 2014. DOI: [10.1016/J.PHYSC.2014.06.012](https://doi.org/10.1016/J.PHYSC.2014.06.012).
- [35] M. N. Wilson, *Superconducting Magnets*. Oxford University Press, New York, 1986, ISBN: 0-19-854810-9.
- [36] F. Grilli, E. Pardo, A. Stenvall, D. N. Nguyen, W. Yuan, and F. Gömöry, “Computation of losses in HTS Under the Action of Varying Magnetic Fields and Currents,” *IEEE Transactions on Applied Superconductivity*, vol. 24, no. 1, 2014. DOI: [10.1109/TASC.2013.2259827](https://doi.org/10.1109/TASC.2013.2259827).
- [37] D. Pracht, “Multiphysics modelling of the LHC main quadrupole superconducting circuit,” Master’s Thesis, Hochschule Hannover University of Applied Sciences and Arts, 2019. [Online]. Available: <http://cds.cern.ch/record/2681097?ln=de>.
- [38] M. Janitschke, “Framework for automatic superconducting magnet model generation & validation against transients measured in LHC magnets,” Master’s Thesis, Technische Universität Berlin, 2021. [Online]. Available: <https://cds.cern.ch/record/2799810>.
- [39] M. Janitschke, M. Mentink, F. Murgia, D. Pracht, E. Ravaioli, and A. P. Verweij, “A Simplified Approach to Simulate Quench Development in a Superconducting Magnet,” *IEEE Transactions on Applied Superconductivity*, vol. 31, no. 5, Feb. 2021. DOI: [10.1109/TASC.2021.3059980](https://doi.org/10.1109/TASC.2021.3059980). [Online]. Available: <https://ieeexplore.ieee.org/document/9356128>.
- [40] M. A. Hilal and Y. M. Eyssa, “Self protection of high current density superconducting magnets,” *IEEE Transactions on Magnetics*, vol. 25, no. 2, pp. 1604–1607, 1989. DOI: [10.1109/20.92605](https://doi.org/10.1109/20.92605). [Online]. Available: <https://ieeexplore.ieee.org/document/92605>.

- [41] R. Schmidt, C. Giloux, A. Hilaire, A. Ijspeert, F. Rodriguez-Mateos, and F. Sonnemann, "Protection of the Superconducting Corrector Magnets for the LHC," in *7th European Particle Accelerator Conference*, 2000, pp. 2163–2165. [Online]. Available: <https://cds.cern.ch/record/466521/?ln=de>.
- [42] A. Gharib *et al.*, "Cryogenic Testing of High Current By-Pass Diode Stacks for the Protection of the Superconducting Magnets in the LHC," in *International Cryogenic Materials Conference*, 2003. DOI: [10.1063/1.1774639](https://doi.org/10.1063/1.1774639). [Online]. Available: <https://cds.cern.ch/record/709119?ln=de>.
- [43] K. Dahlerup-Petersen, B. Kazmine, V. Popov, V. Sytchev, L. Vassiliev, and V. Zubko, "Energy Extraction Resistors for the Main Dipole and Quadrupole Circuits of the LHC," in *7th European Particle Accelerator Conference*, 2000. [Online]. Available: <https://cds.cern.ch/record/466523?ln=de>.
- [44] L. Coull, D. Hagedorn, V. Remondino, and F. Rodriguez-Mateos, "LHC Magnet Quench Protection System," *IEEE Transactions on Magnetics*, vol. 30, no. 4, pp. 1742–1745, 1994. DOI: [10.1109/20.305593](https://doi.org/10.1109/20.305593).
- [45] E. Ravaioli, V. I. Datskov, C. Giloux, G. Kirby, H. H. Ten Kate, and A. P. Verweij, "New, coupling loss induced, quench protection system for superconducting accelerator magnets," *IEEE Transactions on Applied Superconductivity*, vol. 24, no. 3, 2014. DOI: [10.1109/TASC.2013.2281223](https://doi.org/10.1109/TASC.2013.2281223).
- [46] F. Rodriguez-Mateos and F. Sonnemann, "Quench heater studies for the LHC magnets," in *Proceedings of the 2001 Particle Accelerator Conference*, vol. 5, IEEE, 2001, pp. 3451–3453. DOI: [10.1109/PAC.2001.988141](https://doi.org/10.1109/PAC.2001.988141). [Online]. Available: <https://ieeexplore.ieee.org/document/988141>.
- [47] E. Todesco, *Presentation: Unit 13 Protection systems*, 2020. [Online]. Available: <https://indico.cern.ch/event/940961/?print=1%7B%5C%7Dview=standard%7B%5C%7Dnumbered%7B%5C%7Dinline%7B%5C%7Dminutes> (visited on 02/21/2024).
- [48] E. Ravaioli *et al.*, "Protecting a Full-Scale Nb3Sn Magnet With CLIQ, the New Coupling-Loss-Induced Quench system," *IEEE Transactions on Applied Superconductivity*, vol. 25, no. 3, 2015. DOI: [10.1109/TASC.2014.2364892](https://doi.org/10.1109/TASC.2014.2364892).
- [49] E. Ravaioli *et al.*, "First Implementation of the CLIQ Quench Protection System on a Full-Scale Accelerator Quadrupole Magnet," *IEEE Transactions on Applied Superconductivity*, vol. 26, no. 3, Apr. 2016, ISSN: 1558-2515. DOI: [10.1109/TASC.2016.2529840](https://doi.org/10.1109/TASC.2016.2529840). [Online]. Available: <https://ieeexplore.ieee.org/document/7409944>.
- [50] E. Ravaioli *et al.*, "First Implementation of the CLIQ Quench Protection System on a 14-m-Long Full-Scale LHC Dipole Magnet," *IEEE Transactions on Applied Superconductivity*, vol. 26, no. 4, 2016. DOI: [10.1109/TASC.2015.2510400](https://doi.org/10.1109/TASC.2015.2510400).
- [51] S. M. Rowan, "LHC Main Dipole Magnet Circuits: Sustaining Near-Nominal Beam Energies," Dissertation, University of Glasgow, 2016. [Online]. Available: <https://cds.cern.ch/record/2229989?ln=de>.
- [52] E. Ravaioli *et al.*, "Impact of the Voltage Transients After a Fast Power abort on the Quench Detection System in the LHC Main Dipole Chain," *IEEE Transactions on Applied Superconductivity*, vol. 22, no. 3, 2012. DOI: [10.1109/TASC.2012.2183572](https://doi.org/10.1109/TASC.2012.2183572).

- [53] CERN, *Diagram of an LHC dipole magnet - CERN Document Server*, 1999. [Online]. Available: <https://cds.cern.ch/record/40524> (visited on 02/24/2024).
- [54] L. Bortot *et al.*, “A 2-D Finite-Element Model for Electrothermal Transients in Accelerator Magnets,” *IEEE Transactions on Magnetics*, vol. 54, no. 3, 2018. DOI: [10.1109/TMAG.2017.2748390](https://doi.org/10.1109/TMAG.2017.2748390).
- [55] A. Vitrano, M. Wozniak, E. Schnaubelt, T. Mulder, E. Ravaioli, and A. Verweij, “An Open-Source Finite Element Quench Simulation Tool for Superconducting Magnets,” *IEEE Transactions on Applied Superconductivity*, vol. 33, no. 5, 2023. DOI: [10.1109/TASC.2023.3259332](https://doi.org/10.1109/TASC.2023.3259332).
- [56] M. Wozniak, E. Ravaioli, and A. Verweij, “Fast Quench Propagation Conductor for Protecting Canted Cos-Theta Magnets,” *IEEE Transactions on Applied Superconductivity*, vol. 33, no. 5, Aug. 2023. DOI: [10.1109/TASC.2023.3247997](https://doi.org/10.1109/TASC.2023.3247997).
- [57] M. Mentink, J. Van Nugteren, F. Mangiarotti, M. Duda, and G. Kirby, “Quench Behavior of the HL-LHC Twin Aperture Orbit Correctors,” *IEEE Transactions on Applied Superconductivity*, vol. 28, no. 3, 2018. DOI: [10.1109/TASC.2018.2794451](https://doi.org/10.1109/TASC.2018.2794451).
- [58] L. Bortot *et al.*, “STEAM: A Hierarchical Cosimulation Framework for Superconducting Accelerator Magnet Circuits,” *IEEE Transactions on Applied Superconductivity*, vol. 28, no. 3, 2018. DOI: [10.1109/TASC.2017.2787665](https://doi.org/10.1109/TASC.2017.2787665).
- [59] M. Maciejewski, “Co-Simulation of Transient Effects in Superconducting Accelerator Magnets,” PhD Dissertation, Lodz University of Technology, 2018. [Online]. Available: <https://cds.cern.ch/record/2675039>.
- [60] E. Ravaioli, B. Auchmann, M. Maciejewski, H. H. ten Kate, and A. P. Verweij, “Lumped-Element Dynamic Electro-Thermal model of a superconducting magnet,” *Cryogenics*, vol. 80, pp. 346–356, 2016. DOI: [10.1016/J.CRYOGENICS.2016.04.004](https://doi.org/10.1016/J.CRYOGENICS.2016.04.004).
- [61] V. Chakravarty, “Automation of the analysis of Large Hadron Collider superconducting circuit families using PSPICE and XYCE in the STEAM framework at CERN,” Master’s Thesis, Indian Institute of Science Education and Research Bhopal, 2023. [Online]. Available: <https://cds.cern.ch/record/2884176>.
- [62] D. Mayr, “3D Modelling and Transient Simulation of the Superconducting High Order Corrector Magnets for the High Luminosity Upgrade of the LHC,” Master’s Thesis, Universität Innsbruck, 2024. [Online]. Available: <http://cds.cern.ch/record/2887378>.
- [63] *Website steam-sdk on gitlab*. [Online]. Available: https://gitlab.cern.ch/steam/steam_sdk (visited on 07/08/2024).
- [64] *Website steam-sdk on pypi*. [Online]. Available: <https://pypi.org/project/steam-sdk/> (visited on 07/08/2024).
- [65] *Notebook-Centric Workflows for Signal Monitoring, Hardware Commissioning and Operation Analysis*. [Online]. Available: <https://sigmon.web.cern.ch/about> (visited on 07/05/2024).
- [66] *Website SIGMON on gitlab*. [Online]. Available: <https://gitlab.cern.ch/LHCData/lhc-sm-analysis> (visited on 07/08/2024).

- [67] A. Laprade, S. Pearson, S. Benczkowski, G. Dolny, and F. Wheatley, “Application Note 7532 - A new PSPICE Electro-Thermal Subcircuit For Power MOSFETs,” Tech. Rep., 2003. [Online]. Available: <https://www.onsemi.cn/download/application-notes/pdf/an-7532.pdf>.
- [68] *Website Image STEAM_SDK_Dakota_structure*. [Online]. Available: https://gitlab.cern.ch/steam/steam_sdk/-/blob/master/documents/STEAM_SDK_Dakota_structure.svg?ref_type=heads (visited on 07/08/2024).
- [69] J. Barth *et al.*, “A Modernized Architecture for the Post Mortem System at CERN,” *13th International Particle Accelerator Conference (IPAC 2022)*, pp. 1557–1560, 2022. DOI: [10.18429/JACoW-IPAC2022-TUPOMS055](https://doi.org/10.18429/JACoW-IPAC2022-TUPOMS055).
- [70] J. Wozniak and C. Roderick, “NXCALLS - Architecture and Challenges of the Next CERN Accelerator Logging Service,” *17th Biennial International Conference on Accelerator and Large Experimental Physics Control Systems (ICALEPCS)*, pp. 1465–1469, 2019. DOI: [10.18429/JACoW-ICALEPCS2019-WEPHA163](https://doi.org/10.18429/JACoW-ICALEPCS2019-WEPHA163). [Online]. Available: <https://cds.cern.ch/record/2778529?ln=de>.
- [71] *Website lhc-sm-api on gitlab*. [Online]. Available: <https://gitlab.cern.ch/LHCData/lhc-sm-api> (visited on 07/08/2024).
- [72] I. Zaitsev, *The best format to save pandas data*, 2019. [Online]. Available: <https://towardsdatascience.com/the-best-format-to-save-pandas-data-414dca023e0d> (visited on 04/03/2024).
- [73] P. M. Attia *et al.*, “Closed-loop optimization of fast-charging protocols for batteries with machine learning,” *Nature*, vol. 578, pp. 397–402, Feb. 2020. DOI: [10.1038/s41586-020-1994-5](https://doi.org/10.1038/s41586-020-1994-5). [Online]. Available: <https://www.nature.com/articles/s41586-020-1994-5>.
- [74] B. M. Adams *et al.*, *Dakota User’s Manual*, 2022. [Online]. Available: <https://www.sandia.gov/app/uploads/sites/241/2023/03/Users-6.16.0.pdf> (visited on 04/03/2024).
- [75] *Website Dakota – Sandia National Laboratories*. [Online]. Available: <https://dakota.sandia.gov/> (visited on 07/08/2024).
- [76] M. Wozniak *et al.*, “Quench Protection of the HL-LHC Hollow Electron Lens Superconducting Solenoid Magnets,” *IEEE Transactions on Applied Superconductivity*, vol. 32, no. 6, 2022. DOI: [10.1109/TASC.2022.3162178](https://doi.org/10.1109/TASC.2022.3162178).
- [77] G. Zachou *et al.*, *Website STEAM Material Library*. [Online]. Available: <https://steam-material-library.docs.cern.ch/> (visited on 07/05/2024).
- [78] D. P. Bentz and K. R. Prasad, “Thermal Performance of Fire Resistive Materials I. Characterization with Respect to Thermal Performance Models,” National Institute of Standards and Technology, Gaithersburg, Tech. Rep., 2007. [Online]. Available: https://www.researchgate.net/publication/241211063_Thermal_Performance_of_Fire_Resistive_Materials_I_Characterization_with_Respect_to_Thermal_Performance_Models.
- [79] G. D’Angelo *et al.*, “Performance of the Large Hadron Collider’s Cryogenic Bypass Diodes over the First Two Physics Runs, Future Projects, and Perspectives,” *IEEE Transactions on Applied Superconductivity*, vol. 30, no. 4, 2020. DOI: [10.1109/TASC.2020.2970909](https://doi.org/10.1109/TASC.2020.2970909).

- [80] A. Will, “Damage mechanisms in superconductors due to the impact of high energy proton beams and radiation tolerance of cryogenic diodes used in particle accelerator magnet systems,” Dissertation, Karlsruher Instituts für Technologie (KIT), 2021. DOI: [10.5445/IR/1000134497](https://publikationen.bibliothek.kit.edu/1000134497). [Online]. Available: <https://publikationen.bibliothek.kit.edu/1000134497>.
- [81] A. Will *et al.*, “Characterization of the radiation tolerance of cryogenic diodes for the High Luminosity LHC inner triplet circuit,” *PHYSICAL REVIEW ACCELERATORS AND BEAMS*, vol. 23, 2020. DOI: [10.1103/PhysRevAccelBeams.23.053502](https://doi.org/10.1103/PhysRevAccelBeams.23.053502).
- [82] DYNEX, “Datasheet DS2101SY,” Tech. Rep., 2001.
- [83] G. P. Willering *et al.*, “Performance of the cold powered diodes and diode leads in the main magnets of the LHC,” *IOP Conference Series: Materials Science and Engineering*, vol. 101, 2015. DOI: [10.1088/1757-899X/101/1/012076](https://doi.org/10.1088/1757-899X/101/1/012076). [Online]. Available: <https://iopscience.iop.org/article/10.1088/1757-899X/101/1/012076>.
- [84] CERN, M. Brice, and C. Marcelloni, *A worker inside the LHC tunnel*, 2006. [Online]. Available: <https://cds.cern.ch/record/967185?ln=de> (visited on 06/20/2024).
- [85] W. Shockley, “The theory of p-n junctions in semiconductors and p-n junction transistors,” *The Bell System Technical Journal*, vol. 28, no. 3, pp. 435–489, 1949. DOI: [10.1002/j.1538-7305.1949.tb03645.x](https://doi.org/10.1002/j.1538-7305.1949.tb03645.x). [Online]. Available: <https://ieeexplore.ieee.org/document/6773080>.
- [86] P. E. Gray, D. DeWitt, A. R. Boothroyd, and J. F. Gibbons, *Physical electronics and circuit models of transistors*, 2nd ed. New York, London, Sydney: John Wiley & Sons, 1964, ISBN: 9780471323150.
- [87] D. Wollmann *et al.*, “Characterisation of the radiation hardness of cryogenic bypass diodes for the HL-LHC inner triplet circuit,” in *10th International Particle Accelerator Conference*, JACoW Publishing, 2019. DOI: [10.18429/JACoW-IPAC2019-THPTS067](https://doi.org/10.18429/JACoW-IPAC2019-THPTS067). [Online]. Available: <https://cds.cern.ch/record/2690323>.
- [88] H. Korkut, N. Yildirim, A. Turut, and H. Dogan, “Analysis of current–voltage–temperature characteristics and T0 anomaly in Cr/n-GaAs Schottky diodes fabricated by magnetron sputtering technique,” *Materials Science and Engineering: B*, vol. 157, no. 1, pp. 48–52, Feb. 2009, ISSN: 0921-5107. DOI: [10.1016/J.MSEB.2008.12.009](https://doi.org/10.1016/J.MSEB.2008.12.009). [Online]. Available: <https://www.sciencedirect.com/science/article/abs/pii/S0921510708006259>.
- [89] T. C. Lee, T. P. Chen, H. L. Au, S. Fung, and C. D. Beling, “Temperature Dependence of the Ideality factor of GaAs and Si Schottky Diodes,” *Physica Status Solidi (a)*, vol. 152, no. 2, pp. 563–571, 1995. DOI: [10.1002/pssa.2211520225](https://doi.org/10.1002/pssa.2211520225). [Online]. Available: <https://onlinelibrary.wiley.com/doi/abs/10.1002/pssa.2211520225>.
- [90] Z. Charifoulline and A. Verweij, “Hot spot temperatures in the MB and MQ diode stacks and leads,” CERN, Geneva, Tech. Rep., 2018, pp. 1–29.