Design of the High Luminosity Large Hadron Collider Beam Dump Vessel

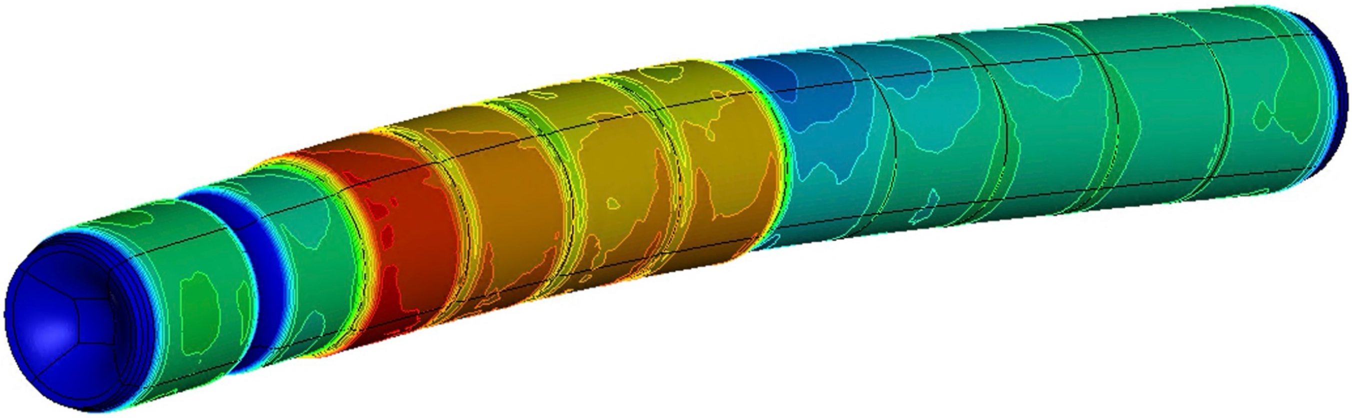

Asbjørn Lassen Lund

Master's Thesis, Spring 2026

Aalborg University, Denmark
Faculty of Engineering and Science
Mechanical Engineering, Design of Mechanical Systems

March 2026

This report was typeset with LaTeX using Overleaf. Charts are created using Microsoft Excel and Matlab and flowcharts and diagrams are created using draw.io. Figures from simulations are outputs from Ansys, LS Dyna or Python.

The title page illustration is taken from a simulation conducted in LS Dyna of the dynamic response of the High Luminosity beam dump subjected to a full-intensity beam impact. The image shows a schematic colour contour of the von Mises stress.

The report has a total of 142 000 characters, including space, corresponding to 59.2 normal pages (2 400 characters per page).

Abstract

The High-Luminosity (HL) upgrade of the Large Hadron Collider (LHC), scheduled for implementation in 2028, aims to significantly increase physics output of the accelerator machine. This upgrade results in stored beam energies of up to 710 MJ, imposing severe demands on the beam dump systems that must safely absorb this energy within 86 µs.

This work presents the design and thermo-mechanical assessment of a new vessel for the HL-LHC beam dumps capable of withstanding the increased operational loads. The response of the currently operational Run 3 beam dumps is first analysed using a combination of finite element simulations and experimental validation. The study shows that the existing design would not operate safely under HL-LHC conditions due to large oscillating stresses generated by the particle shower and associated energy deposition, which excite the structural vibration modes of the device. Based on these findings, a series of finite element based optimisation studies are performed to identify stress mitigation strategies and develop an improved beam dump design for HL.

The proposed design features a titanium Grade 5 vessel, leveraging the alloy's low stiffness, low coefficient of thermal expansion, low density, and high specific strength to reduce dynamic stresses induced during beam impact. The design is validated through thermo-mechanical simulations and detailed fatigue analyses, including assessment of the vessel, the electron beam welded joints, the beam windows, and their bolted flange connection.

The results demonstrate that the new design pf the beam dumps for HL maintains structural integrity under HL-LHC beam impact conditions and satisfies high fatigue safety margins. This work provides a validated design solution for the HL-LHC beam dumps and supports the initiation of their production in 2026.

Asbjørn Lassen Lund

Preface

This report presents the master's thesis project undertaken by Asbjørn Lassen Lund during the fourth semester of the Mechanical Engineering master's programme, specialising in the design of mechanical systems at Aalborg University, Denmark. The thesis was compiled during the autumn semester of 2025 and the spring semester of 2026 and addresses the design of the new-generation HL beam dumps for the LHC at CERN, with particular emphasis on the vessel design.

Reading Guide

The reader is assumed to have a solid foundation in mechanical engineering principles, including structural mechanics, the application of numerical methods, finite element analysis (FEA), and fatigue and fracture mechanics, as these constitute key aspects of the design and analysis process. In this report, FEA is used extensively to simulate physical phenomena, predict mechanical performance, and optimise designs.

The report complies with the ISO/IEC 80000 standard. All quantities are expressed using SI units unless explicitly specified otherwise. Numerical values are written using a period (.) as the decimal marker and a space to separate thousands. References are formatted in accordance with the IEEE citation style. Any figures without source references have been produced by the author. Figures, tables, and equations are systematically numbered by chapter and sequence of appearance. An overview of abbreviations is included prior to the table of contents.

Generative artificial intelligence was used to (i) produce MATLAB scripts for generating plots from raw simulation data, (ii) assist in the development of certain Python scripts (explicitly stated on a case-by-case basis), and (iii) paraphrase text exclusively to improve linguistic quality and sentence formulation.

The report begins with background information on CERN, particle accelerators, and the CERN accelerator complex. An introduction to beam-matter interaction is given followed by an introduction to the LHC beam dumps and the corresponding beam dump process. Previous research is briefly reviewed to establish the state of work and to identify the sources for selected inputs used in this thesis. The problem analysis addresses the HL upgrade of the LHC and its impact on the beam dumps. Expected operating conditions are defined using historical data and known failure mechanisms, and the relevant operational response categories and thesis scope are outlined. A system requirements definition serves as the problem statement, specifying functional and non-functional design requirements. The finite element model used to evaluate different design concepts is then presented, followed by a detailed assessment of the vessel, welds, and beam window design. The report concludes with a summary of the main findings.

Acknowledgements

The author wishes to express his gratitude to the supervisors and professors of the Design of Mechanical Systems programme at Aalborg University for their guidance and constructive feedback throughout the master's programme. Particular thanks are given to Erik Lund for his valuable contributions and insightful discussions related to this thesis. The author also acknowledges his supervisor at CERN, Nicola Solieri, for his guidance and valuable inputs to the project.

Nomenclature

Abbreviations

CERN	The European Organization for Nuclear Research
LBDS	Beam Dump System
LHC	Large Hadron Collider
HL	High Luminosity
LINAC 4	Linear Accelerator 4
PSB	Proton Synchrotron Booster
PS	Proton Synchrotron
SPS	Super Proton Synchrotron
FEA	Finite Element Analysis
HD	High-density
LD	Low-density
DR	Dynamic Relaxation
CFL	Courant–Friedrichs–Lewy
MSD	Extraction Septa
MKB	Extraction Kicker Magnets
LS2	Long Shutdown 2
CFD	Computational Fluid Dynamics
MKBH	Horizontal Dilution Magnet
MKBV	Vertical Dilution Magnet
WPS	Welding Procedure Specification
WPQR	Welding Procedure Qualification Record

Contents

Chapter

1

Introduction

The European Organization for Nuclear Research (CERN) is a European research organisation located at the border between France and Switzerland north of the city of Geneva. Its mission is to advance the understanding of the fundamental constituents of matter and the physical laws governing the universe [35]. At CERN, hadrons, primarily protons, are accelerated using electromagnetic fields generated by radio frequency cavities and guided inside a 27-km long underground ultra-high vacuum ring by superconducting electromagnets [31, 9]. Once the protons have been accelerated to close to the speed of light, they are distributed to various experiments across the facility for fundamental physics research [38].

The beam dumps of the LHC are dedicated devices designed to safely absorb the extremely high kinetic energy of the particle beam. These devices play a critical role in the safe and reliable operation of the accelerator for several reasons:

- **Damage prevention**: Once accelerated to their maximum velocity, the proton beams in the LHC carry an enormous amount of stored energy and can cause severe damage to sensitive accelerator components, including magnets, cavities, electronics, and vacuum beam pipes. Such damage may occur in the event of beam misalignment, equipment malfunction, or unintended beam-matter interactions within the accelerator. Beam dumps are therefore engineered to safely intercept and dissipate the beam energy, thereby protecting critical infrastructure [23, 24, 43].
- **Radiation protection**: Interactions involving high-energy particle beams produce significant radiation levels that may be harmful to both equipment and personnel. By intercepting and absorbing the particle beam within dedicated devices installed in shielded environments, beam dumps confine radiation to designated areas. Consequently, beam dumps are among the most radioactive components within the CERN complex [23, 24, 43].
- **Operational role**: Beam dumps are not limited to emergency scenarios. During normal operation, the number of protons circulating in the Large Hadron Collider (LHC) gradually decreases as particles are consumed in collisions at the experiments. When the beam intensity drops below the level required to sustain the desired collision rate, beam dumps allow the beam to be safely disposed prior to refilling the accelerator. In addition, beam dumps play a crucial role during commissioning and beam tuning phases, enabling verification of magnet and cavity settings without directing the beam toward experimental areas [23, 24, 43].

The HL upgrade represents the next evolutionary stage of the LHC, involving significant enhancements to the underground infrastructure to increase beam energy and collision rate. While these upgrades will substantially improve the physics output of the experiments, they will also impose considerably higher loads on the beam dumps. In HL, the beam dumps must safely absorb beam energies of up to 710 MJ which is deposited within 86 µs, presenting a substantial engineering challenge.

The initial problem statement is as follows:

What are the requirements and constraints for the new design of the High-Luminosity Large Hadron Collider beam dumps?

Chapter

2

Background

This chapter provides the necessary background to follow the problem analysis and the remainder of the report. It introduces CERN, particle accelerators, and their basic operating principles, and clarifies commonly used accelerator terminology. The chapter then addresses beam-matter interactions relevant to beam dumps and presents an overview of the CERN accelerator complex. The current operational beam dump design and the beam dumping process are introduced and finally, previous work is reviewed to establish the current state of development and illustrate and define the scope of this thesis.

2.1 The European Organization for Nuclear Research

CERN was founded in 1954 by twelve European countries, including Denmark, with the purpose of fostering international scientific collaboration in the post-war period [35]. Today, the organisation comprises 25 member states and is supported by an annual budget of approximately 1.4 billion Swiss Franc [6]. More than 17 500 scientists, engineers, and technical staff from over 110 nationalities contribute to CERN's research activities.

The primary scientific objective of CERN is to support research into the fundamental structure of matter and the forces governing particle interactions. CERN operates the world's largest particle accelerator complex and by accelerating protons to velocities approaching the speed of light and colliding them inside the dedicated experiment, kinetic energy is converted into mass, enabling the production and study of rare and short-lived particles. Measurements of the resulting particle interactions and decay processes contribute to research on topics such as the properties of the Higgs boson and physics beyond the Standard Model [56, 38].

The accelerator complex has evolved significantly since its initial construction and currently consists of seven accelerators and two decelerators, supplying particle beams over a wide range of energies to various experiments. The physics programme has also evolved to include research areas beyond high-energy particle physics such as nuclear physics, antimatter studies, and applications in medical imaging and radiation therapy [28, 38].

2.2 Particle Accelerators

Particle accelerators are sophisticated systems developed to accelerate and steer charged particles to extreme velocities for experimental purposes. At CERN, these particles are predominantly protons, organised into tightly confined beams that travel through the accelerator complex. Familiarity with the physical characteristics of proton beams is fundamental to understanding the design principles of beam interception and beam dump systems [23, 24, 43].

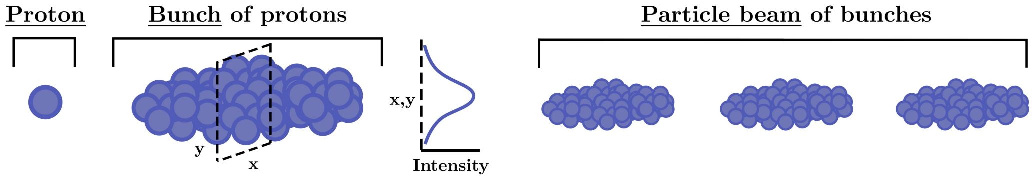

Figure 2.1: Schematic illustration of a proton beam, highlighting the organisation of protons into bunches used for acceleration and experimental applications [43].

Figure 2.1 illustrates the composition of a particle beam. Protons are produced from hydrogen gas through an ionisation process in which an applied electric field removes the electrons from the atom. This extraction procedure is repeated multiple times, with each cycle yielding a specific number of protons that are subsequently assembled into a bunch of protons [49]. The accelerator system then accelerates a sequence of such bunches. The use of bunched beams allows for efficient operation and precise control of the beam, as each bunch can be accurately guided toward different experimental areas or collision points. Under typical operating conditions, a single bunch contains on the order of 10^{10} to 10^{12} protons, a parameter referred to as the bunch intensity. Modern accelerator facilities are capable of simultaneously circulating hundreds to several thousand bunches within the machine simultaneously [16, 43].

In the transverse plane, the spatial distribution of protons within a bunch is commonly described by a Gaussian profile, as shown in Figure 2.1 [50, 43]. The particle density is highest along the beam axis and decreases toward the edges of the bunch. A compact transverse beam profile enhances focusing efficiency, increases the likelihood of particle collisions at the experiments, and reduces beam losses caused by interactions with the accelerator aperture.

The kinetic energy of the proton beam is increased gradually as the particles pass through a series of radio frequency cavities, which apply oscillating electric fields to accelerate the protons [31]. The energy of an individual proton is referred to as the proton energy, whereas the integrated energy of all protons represents the total stored energy of the particle beam. The latter is obtained by multiplying the energy per proton by the bunch intensity and the total number of circulating bunches.

Beam guidance and confinement are achieved using magnets distributed along the accelerator that operate in a superconducting state. This state is maintained by cryogenically cooling the beam pipe and surrounding structures with helium [9, 43]. In addition, the particle beam travels in ultra-high vacuum to limit particle interactions with residual gas molecules.

2.3 Accelerator Complex

The largest accelerator at CERN is the circular LHC, installed in a tunnel with a circumference of approximately 27 km at a depth of about 100 m below ground level [51]. Achieving the beam energies required for operation of the LCH relies on a chain of accelerators that progressively increase the energy of the proton bunches prior to injection.

An overview of the CERN accelerator complex is shown in Figure 2.2. The proton bunches extracted from hydrogen are initially injected into the Linear Accelerator 4 (Linac 4), which increases their energy to 160 MeV. The beam is subsequently transferred to the Proton Synchrotron Booster (PSB), followed by further acceleration in the Proton Synchrotron (PS) and the Super Proton Synchrotron (SPS). Through this sequence, the proton energy is increased from 120 MeV to 440 GeV, after which the beams are injected into the two counter-rotating beam-lines of the LHC [43].

Inside the LHC, the protons are accelerated to 6.8 TeV before colliding at four main experimental interaction points hosting the ALICE, ATLAS, CMS, and LHCb detectors [38]. At these energies, the protons travel at velocities exceeding 99.99% of the speed of light, completing nearly 11 000 revolutions in the accelerator ring per second. During Run 3 operation, which refers to the current generation of the LHC, the stored energy of a full proton intensity beam reaches 539 MJ [51, 45]. For comparison, this amount of kinetic energy is in the same order of magnitude as an Airbus A310 at maximum take-off weight of 164 t travelling at a take-off speed of 290 km/h.

The LHC operates in defined periods known as Runs. As shown in Table 2.1, the machine is presently in Run 3, while Run 4, the HL run, is planned to begin in 2030 after a prolonged shutdown.

Table 2.1: Overview of the different LHC runs [?].

Run 1 (2009–2013)	Run 2 (2015–2018)	Run 3 (2022–2026)	Run 4 / HL-LHC (2030–2041)

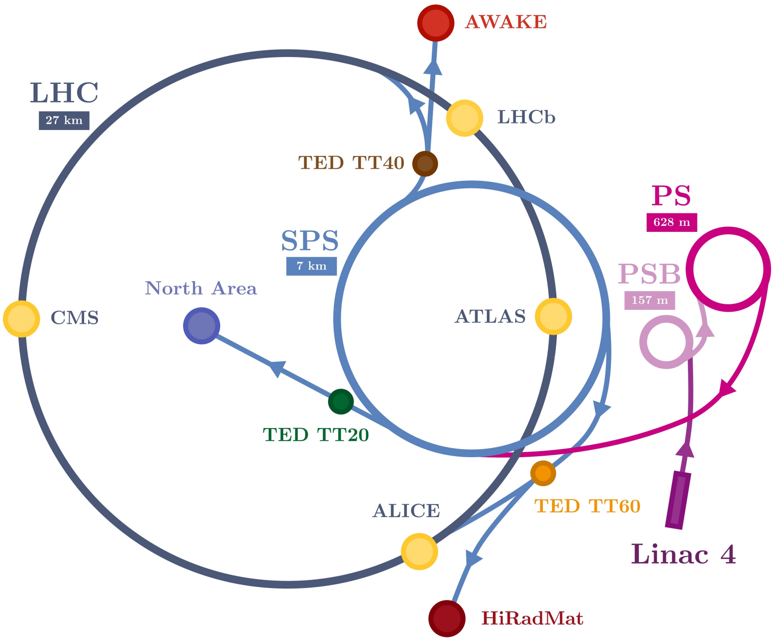

Figure 2.2: Schematic overview of the CERN Accelerator Complex, showing the accelerators, beamlines and experimental areas. The proton beam is initially accelerated in Linac 4, then passes sequentially through the PSB, the PS, and the SPS before being injected into the LHC ring [43].

2.4 Beam-Matter Interaction

During a beam dump event, the particle beam is intentionally directed into a dedicated absorbing material. Upon impact, a series of complex particle-matter interactions are initiated. As the particles propagate through the material, multiple physical processes occur that can alter the material properties and directly affect the functional performance of the beam dump system [43]. The principal interaction mechanisms are summarised below:

- **Angular scattering and atomic displacement:**
  Interactions between impacting protons and atomic nuclei in the target material result in Coulomb scattering, causing deflection of the protons and displacement of atoms from their lattice positions. These atomic displacements lead to radiation-induced material damage over time, which can progressively reduce mechanical properties [53].
- **Formation of residual nuclei and radioactive isotopes:**
  Proton impacts can result in nuclear reactions in the impacted material that generate unstable residual nuclei. The subsequent decay of these nuclei produces radiation. Consequently, material selection for beam dump devices must consider not only mechanical and thermal properties, but also the activation behaviour and the minimisation of long-lived radioactive isotopes [53].
- **Secondary particle production and particle showers:**
  At beam impact, a cascade of secondary particles including photons, electrons, and positrons is generated, referred to as a particle shower. This process distributes the deposited energy over an extended volume, thereby influencing the internal thermal field in the impacted material. Figure 2.3 presents a simulated example of a particle shower produced by a single 450 GeV proton impacting an aluminium target. The

formation of the particle shower is clearly characterised by the scattering and propagation of a significant amount of secondary particles within the material [23].

- **Heat generation and thermal response:**
  The kinetic energy of the primary and secondary particles is deposited in the target material through particle–matter interactions, where it is ultimately converted into heat rather than transferred as momentum. This energy deposition leads to rapid and often highly localised temperature rises within the material. The resulting thermal gradients can induce severe thermal shocks, which may cause material degradation or failure. The energy density deposited in the material can be influenced by the material's density, as lower-density materials tend to spread particle interactions over a larger volume, thereby reducing the peak energy density. However, this also leads to less spatial confinement of the particle shower, representing a design trade-off. During repeated beam impacts, significant heat accumulation may occur, which is typically managed through the integration of active cooling systems. [23, 43].

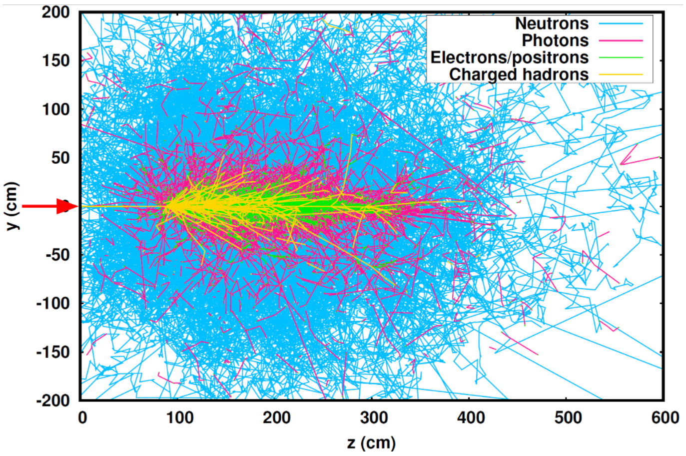

Figure 2.3: Simulated particle shower in aluminium impacted by a 450 GeV proton from the left, as indicated by the red arrow. The particle shower deposits energy in the material [53].

Beam-matter interactions represent the dominant loading mechanisms acting on beam dump devices. As demonstrated by the interaction categories discussed above, their design must account for a wide range of considerations throughout the design process, from material selection and geometric layout to thermo-mechanical response and thermal management.

2.5 Large Hadron Collider Beam Dumps

The beam dumps are critical components of the LHC Beam Dump System (LBDS), which ensures the safe extraction and absorption of the circulating particle beams under all operating conditions, including abnormal machine scenarios [21]. Two beam dump devices are installed, one for each of the counter-rotating beam.

The beam dumps are located in dedicated underground caverns and are connected to the LHC ring by tangential extraction lines approximately 700 m in length. At a distance of about 10 m upstream of the dump, the particle beam exits the ultra-high-vacuum extraction line and propagates freely through air before impacting the beam dump [45]. The front face of an installed beam dump is shown in Figure 2.4. Each beam dump has a diameter of 720 mm, a length of 8.5 m, and a mass of approximately 6.4 t. To limit radiation exposure, the beam dumps are surrounded by massive iron shielding blocks. Cooling of the device is provided by an air-cooling system that directs airflow onto the external surface of the device through nozzles distributed along its length [45]. Two air supply pipes connected to the air handling unit are visible beneath the beam dump in Figure 2.4.

The following section describes the beam dump design currently deployed in the LHC during Run 3. In addition to the beam dump, the system includes a dedicated support structure designed to support the weight of the device. Both the dump block and the support structure are detailed in the subsequent sections.

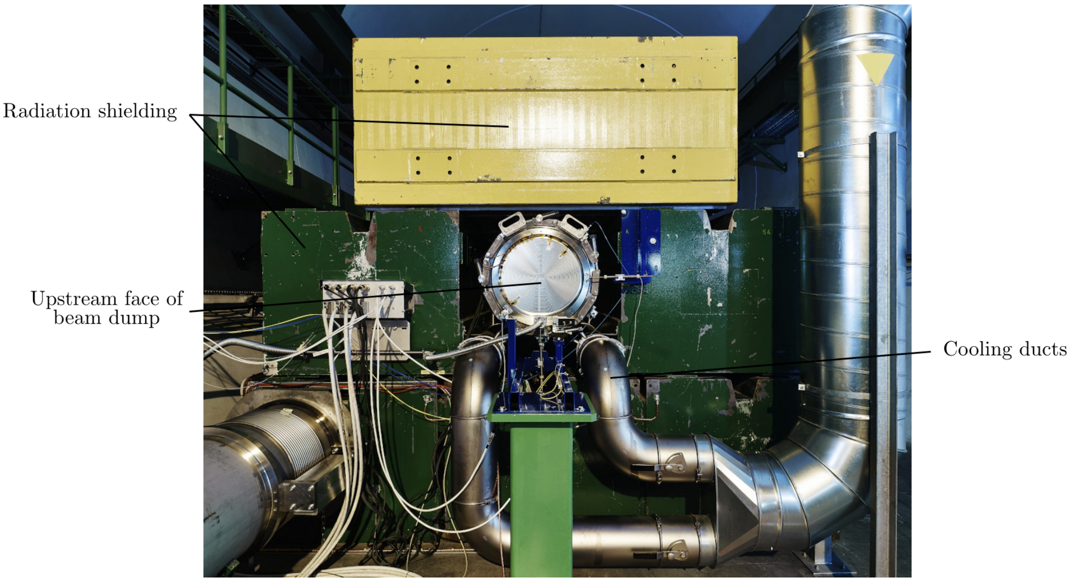

Figure 2.4: Front view of one of the Run 3 LHC beam dumps installed inside thick iron shielding in the underground cavern.

2.5.1 Dump Block

A schematic cross-sectional representation of the Run 3 beam dump design is provided in Figure 2.5.

The pressure vessel is fabricated from duplex stainless steel grade 318LN and consists of upstream and downstream High-density (HD) sectors connected to a central Low-density (LD) sector. The HD sectors are assembled from seamless cylindrical blanks, each approximately 0.7 m in length and 12 mm in wall thickness, which are joined circumferentially using partial-penetration butt welds with a throat thickness of 6.65 mm. The 3.5 m-long LD sector is formed from a bent plate of the same material, which is longitudinally welded, and subsequently connected to the HD sectors by circumferential welds. The vessel is sealed at both ends by 10 mm-thick Titanium Grade 5 windows. The upstream window is secured using a chain mechanism that compresses the window against the flange, while the downstream window is attached to the vessel through a bolted connection [55].

Internally, the device consists of carbon-based materials with varying densities to enable controlled energy absorption along its length [59, 60]. By varying the densities, the materials can be ensured not to be exposed to energy densities they cannot withstand. The upstream and downstream HD sectors contain a total of six 0.7 m-long blocks of isostatic graphite with a density of $1.73\,\mathrm{g/cm^3}$. These blocks are shrink-fitted into the vessel to promote thermal contact and enhance heat transfer from the core to the vessel. The central LD sector contains approximately 1 630 stacked expanded graphite sheets, each 2 mm thick and with a density of $1.2\,\mathrm{g/cm^3}$, resulting in a total length of 3.26 m. Owing to their limited geometrical stiffness, the sheets are not shrink-fitted but are maintained under slight axial compression between two extruded graphite discs, which are secured by retaining rings spot-welded into internal grooves of the vessel [55]. To prevent oxidation of the core materials during operation, the vessel is filled with nitrogen at a pressure of approximately 1.2 bar, thereby maintaining an inert internal environment [45].

2.5.2 Support Structure

The beam dump assembly is suspended from a dedicated support structure using Diepa H40 steel wire ropes with a diameter of 21 mm. This configuration allows the beam dump to vibrate and oscillate during beam impact without inducing additional stresses associated with rigid boundary constraints.

A schematic representation of the support structure, including a transverse cross-sectional view at the upstream

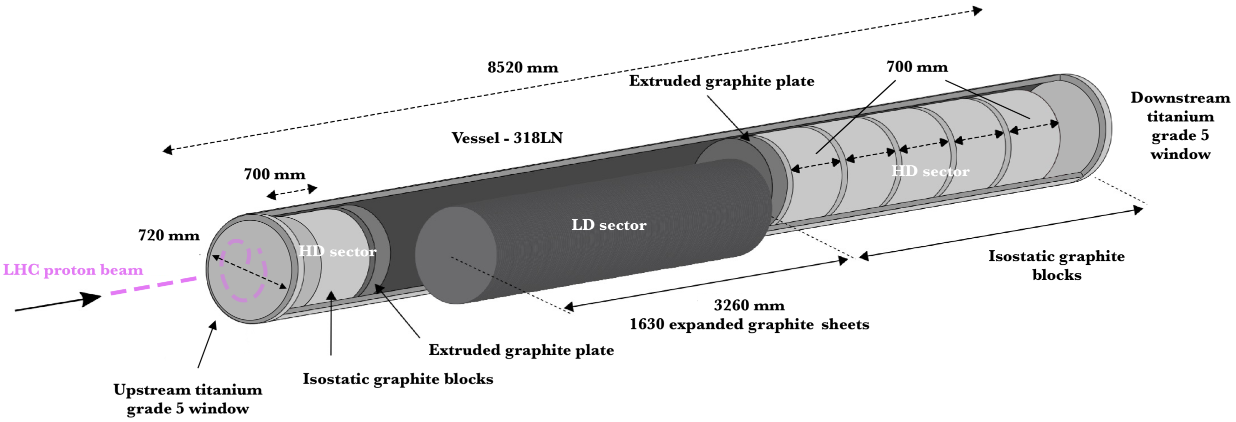

Figure 2.5: Schematic cross-sectional illustration of the current beam dumps in the configuration used for Run 3. In the direction of the beam, this includes: Upstream window of titanium Grade 5; a HD isostatic graphite block; a LD sector comprising 3.26 m of flexible graphite sheets, compressed at both ends by extruded graphite discs; five blocks of HD isostatic graphite; and a downstream titanium Grade 5 window. The core is encased in a 12 mm-thick duplex stainless-steel (318LN) vessel [45, 42]

suspension point, is shown in Figure 2.6.

The gravitational load of the beam dump is transmitted through the suspension ropes to the support structure via pin connections integrated into two cradles spaced 4 m apart. Each cradle has an internal diameter slightly exceeding the outer diameter of the beam dump, providing roughly 10 mm of radial clearance. The upstream and downstream cradles are interconnected by rectangular cross-section side beams that are bolted to the cradles, forming a rigid support structure. The cradles are manufactured from structural steel grade S460ML, while the side beams are fabricated from S235JR [46].

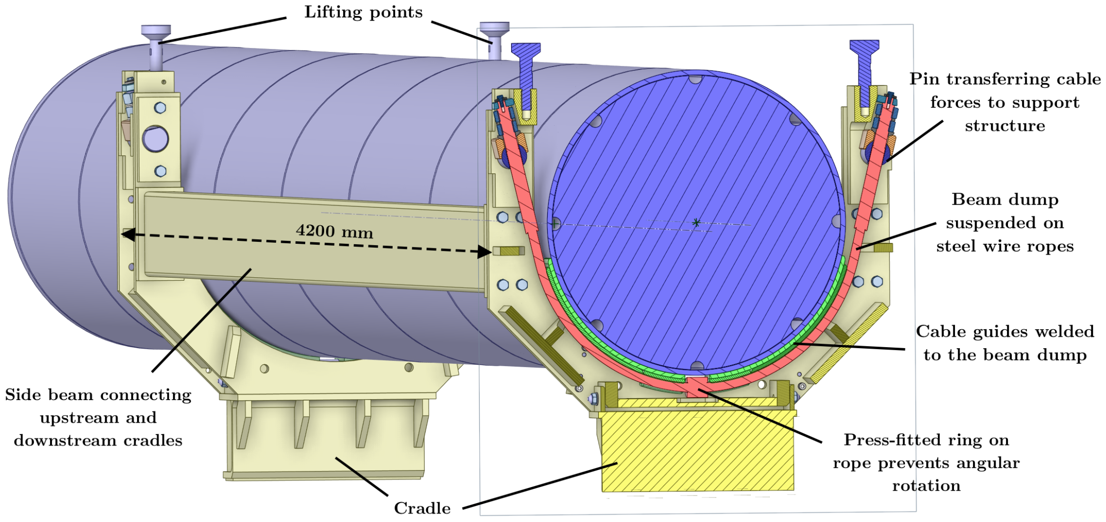

Figure 2.6: Schematic view of the Run 3 beam dump support structure with a transverse cross-section at the upstream suspension point. The illustration highlights the suspension ropes (red), the cable guide (green), the pin load transfer mechanism, and the alignment feature formed by the press-fitted ring engaging with the guide slot [46].

Angular positioning of the beam dump is controlled by a press-fitted ring mounted on each suspension rope. The ring engages with cable guides welded to the outer surface of the vessel and is seated within a machined slot in the guide. This configuration provides a mechanical constraint that limits rotational and longitudinal displacement of the beam dump over time. Although the flexibility of the suspension ropes allows vibration and swinging during beam impact, the alignment mechanism ensures that the assembly returns to a well-defined equilibrium position once oscillations are damped. This feature maintains consistent alignment between the support structure and the vessel during operation without introducing rigid boundary constraints [46].

Once the beam dump is installed on the suspension ropes, the complete assembly can be lifted using four designated lifting points located on the upper side of the two cradles. This feature facilitates installation and is particularly important during decommissioning, when the device may be highly radioactive. In such cases, the ability to transport the assembly without disassembly or direct human intervention is essential to limit personnel exposure.

2.6 Beam Dump Process

Beam dumps can be initiated under a variety of circumstances. These events can be grouped into two principal categories: protective dumps caused by machine faults observed in the accelerator and operational dumps, which are a natural consequence of the LHC cycle:

- **Protective beam dumps:**
  Protective dumps are triggered automatically by the LHC machine protection system when abnormal operating conditions are detected. Typical triggers include power supply or converter failures, magnet quenches, vacuum instabilities, beam instrumentation faults, or excessive beam losses. The protection system continuously monitors the status of thousands of datapoints in the machine equipment and beam parameters and initiates a rapid extraction of the beam from the accelerator if predefined thresholds are exceeded. Such dumps may occur at any beam intensity and stored energy level, depending on when the fault arises [45, 42].
- **Operational beam dump for physics optimization:**
  Beam dumps are also performed intentionally as part of standard operation. After injection and ramp-up to collision energy, the beams are brought into collision within the experiments. The beam intensity gradually decreases due to the loss of protons at the experiment collisions. When the intensity drops below a level considered efficient for physics production, the remaining beam is extracted and dumped to prepare the machine for a new fill at maximum intensity. These planned dumps therefore typically occur at reduced beam intensity and correspondingly lower stored energy [45, 42].

Once a beam dump is triggered or requested, a precisely synchronised sequence of magnetic kickers first extracts the beam from the LHC and subsequently dilutes it over a larger surface area of the beam dump. The layout of the extraction and dilution kicker magnets is illustrated in Figure 2.7. In brief, the two types of kicker magnets operate as follows:

- **Extraction kickers:**
  Fifteen horizontal kicker magnets are installed at the extraction point in the LHC ring (denoted MSD in Figure 2.7). Their extremely short rise time ensures that the circulating beam is deflected at the correct moment. The generated horizontal magnetic field steers the beam out of its nominal circular trajectory of the LHC and into the dedicated extraction tunnels leading toward the beam dumps [20, 42].
- **Dilution kickers:**
  Further downstream in the extraction line, a set of dilution kickers (MKB in Figure 2.7) dilute the bunches of the particle beam on the beam dump front face. The system comprises four horizontal and six vertical kickers operating in combination to produce the dilution referred to as a sweep pattern [21]. By spreading the bunches of the beam over a larger impact area, the dilution system reduces the peak energy density in the absorbing materials and mitigates the risk of localised damage from thermal shock or sublimation [42]. An example of the resulting sweep pattern for Run 3 conditions is shown in Figure 2.8.

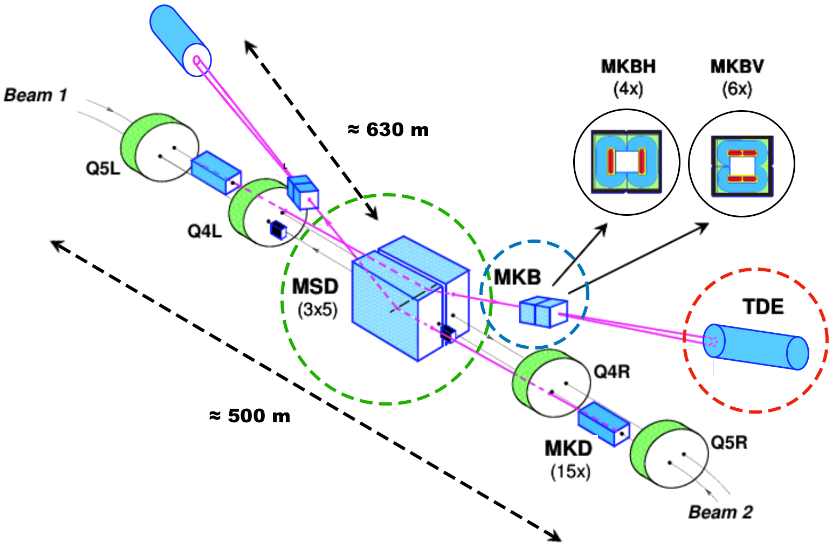

Figure 2.7: Schematic overview of the LHC beam dumping system. Fast horizontal extraction kickers deflect the circulating beam out of the LHC machine and into the extraction lines. Subsequent dilution kickers generate a controlled sweep pattern to reduce peak energy density in the beam dump. After approximately 630 m of transport, the beam reaches the beam dump. Each beam direction is equipped with an independent extraction and dilution system [29].

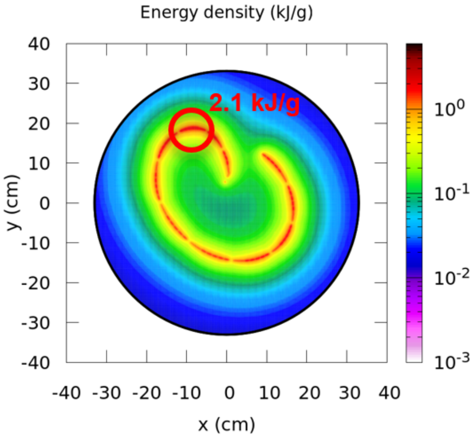

Figure 2.8: Transverse energy density distribution (kJ/g) at a depth of 290 cm within the beam dump under Run 3 conditions. The observed pattern results from the action of the dilution kickers, which sweep the beam across the beam dump surface over the 86 µs extraction duration. The location of peak energy density of 2.1 kJ/g is highlighted. The first bunch impacts near the centre of the pattern, while subsequent bunches follow the sweep trajectory [32].

The transverse energy density distribution shown in Figure 2.8 changes with penetration depth as the particle shower develops within the device. Near the upstream entrance, the deposition is very concentrated, resulting in pronounced gradients perpendicular to the sweep direction. With increasing depth, secondary particle production leads to an enlarged spatial distribution and consequently smoother transverse gradients.

2.7 Previous Published Work

The LHC beam dumps are under the responsibility of a dedicated engineering team at CERN. Their mandate includes both the operation and maintenance of the current Run 3 beam dumps installed in the accelerator, as well as the design and development of the new beam dumps for the HL-LHC era (Run 4). This section provides an overview of the aspects of the beam dumps that have already been addressed in previous publications, with emphasis on the specific topics and methodologies covered in those works. The objective is to clearly define the scope of the present report by identifying which results, models, and assumptions are adopted as input for the simulations and calculations performed here. In doing so, a well-defined starting point for the current study is established.

2.7.1 Modifications and Upgrades During Long Shutdown 2 for Run 3 [46]

During Long Shutdown 2 (LS2, 2018–2022), the LHC beam dump support system was upgraded. In Run 2, nitrogen leaks revealed operational damage while the beam dumps were rigidly fixed in the cavern. Investigations showed that beam-induced thermal shocks excited structural vibrations, leading to mechanical degradation at the window interfaces. To mitigate constraint-induced stresses, the rigid support was replaced with a suspension system (configuration described in Section 2.5.2) in which the beam dump is mounted on steel wire ropes, allowing controlled vibrational motion during operation.

2.7.2 Simulation Benchmarking of Run 3 Response and Inclusion of Support Structure [22]

The finite element model of the Run 3 beam dump was originally developed by the author to investigate concerns regarding the fatigue resistance and structural integrity of the 6.25 mm circumferential partial-penetration welds (out of a 12 mm wall thickness) used to join the vessel segments. The model was employed to determine the far-field stress state in the weld region, enabling subsequent sub modelling and detailed weld assessment. Under the author's supervision, Braken later used the model to benchmark the dynamic response of the Run 3 beam dump against a real beam dump event for which operational measurements were available. In addition, Braken implemented the modelling approach for the steel wire ropes of the support structure and developed a method to simulate the pre-stressed configuration of the dump block resting on the wire ropes prior to beam impact.

2.7.3 Design and Validation of Support Structure under HL Conditions [52, 36]

Bergmann and Kappelgaard performed detailed investigations of the structural integrity of the beam dump support structure under HL-LHC conditions, with the objective of assessing potential design upgrades and stress mitigation strategies. Bergmann concentrated on overall structural response of the support structure and potential design mitigations to reduce stress levels, while Kappelgaard carried out a more detailed evaluation of weld performance and fatigue behaviour.

2.7.4 Design and Validation of Enhanced Cooling System for HL Using Computational Fluid Dynamics [19]

Notarianni carried out detailed computational fluid dynamics (CFD) analyses of the beam dump cooling system. The work first benchmarked the numerical model against temperature measurements from Run 3 operation and subsequently investigated improved cooling configurations for HL-LHC implementation. The study focused on the heat transfer path from the absorber core to the vessel and cooling system, with particular emphasis on the temperature build-up resulting from consecutive beam impacts separated by approximately 2.5 hours.

2.7.5 High Irradiation to Materials Experiment [42, 40]

To qualify core materials for the HL-LHC beam dumps, an experimental campaign was conducted at a dedicated CERN facility in which candidate carbon-based materials were subjected to repeated beam impacts at energy densities representative of HL conditions. The author of this thesis contributed to the design of the experimental setup, which is documented in [42]. The experimental results, including post-irradiation examination, are presented in [40].

Chapter

3

Problem Analysis

This chapter presents the problem analysis forming the basis of the present work. It begins with an overview of the HL upgrade of the LHC and its impact on beam parameters and total stored energy. The historically observed distribution of dumped beam intensities is then examined to characterise the load spectrum experienced by the beam dumps and to extrapolate expected conditions for HL operation. Potential dilution failure scenarios resulting from malfunction of kicker magnets are subsequently described, as these represent critical design cases whose probability of occurrence must be considered. The thermo-mechanical response of the beam dump to beam impact is then outlined to define the loading conditions addressed within the scope of this work. The integration constraints are studied to explore the potential design space for the future dumps. Finally, a comparison of energy deposition in key components between Run 3 and HL conditions is presented to highlight the primary design challenges associated with future operation.

3.1 The High-Luminosity Large Hadron Collider Upgrade

The HL upgrade of the LHC marks the fourth operational phase of the accelerator, designated as Run 4 and planned for the period 2030–2041. This upgrade programme comprises extensive modifications to the accelerator complex, with the objective of substantially increasing physics performance [46].

Successive operational phases of the LHC have been characterised by progressively more demanding beam parameters. From Run 1 (2009–2013) through Run 3 (2022–2026), increase in beam intensity and stored energy have imposed growing requirements on critical systems, including the beam dumps. The transition to HL operation will further amplify these demands. In particular, the total stored energy per beam is projected to rise by approximately 30 %, increasing from 539 MJ in Run 3 to 710 MJ in Run 4 [32]. An overview of the primary beam parameters across the different operational phases is presented in Table 3.1.

Table 3.1: Overview of beam parameters during different LHC runs [15].

	Run 1 (2009–2013)	Run 2 (2015–2018)	Run 3 (2022–2026)	Run 4/HL-LHC (2030–2041)
Beam dumps/year	220 days × 2 = 440			
Energy/proton (TeV)	4.0	6.5	6.8	7.0
Nominal bunch intensity (protons/bunch)	1.7×10^{11}	1.2×10^{11}	1.8×10^{11}	2.2×10^{11}
Maximum number of bunches/beam	1380	2556	2748	2760
Maximum stored energy (MJ)	150	320	539	710

The beam dumps currently in operation during Run 3 are based on a design originally developed for Run 1 and subsequently adapted to accommodate progressively more demanding beam conditions. Although incremental upgrades have been implemented, the transition to HL operation necessitates continued research and development to ensure safe and reliable performance [37].

The primary challenges associated with HL parameters arise from increased proton energy and beam intensity:

- **Proton energy:**
  Higher proton energy alter the development of the particle shower. This can modify the internal energy

deposition distribution in the absorbing material and increase the risk of downstream radiation leakage, potentially requiring adjustments to the absorber layout for HL operation.

- **Beam intensity:**
  The increase in bunch intensity and number of circulating bunches affects the total stored energy of the beam and significantly increases the loading on the beam dumps.

Figure 3.1 presents the transverse peak energy density as a function of depth in the beam dump for the different operational runs, together with the projected values for HL conditions assuming the current beam dump design is retained. The results indicate a pronounced increase in peak energy density under HL operation, particularly within the LD sector. In this region, the peak value rises from approximately 2.1 kJ/g in Run 3 to about 3.4 kJ/g in HL, corresponding to predicted peak temperatures in the order of 1900 °C.

The HL operational phase is expected to span approximately 11 years, exceeding the duration of previous runs and thereby increasing cumulative damage effects. Such prolonged exposure may lead to progressive material degradation of the core or fatigue damage in the vessel without visible external indications. At the same time, the high activation levels following operation prevents immediate inspection or disassembly, limiting direct assessment of the condition of the device.

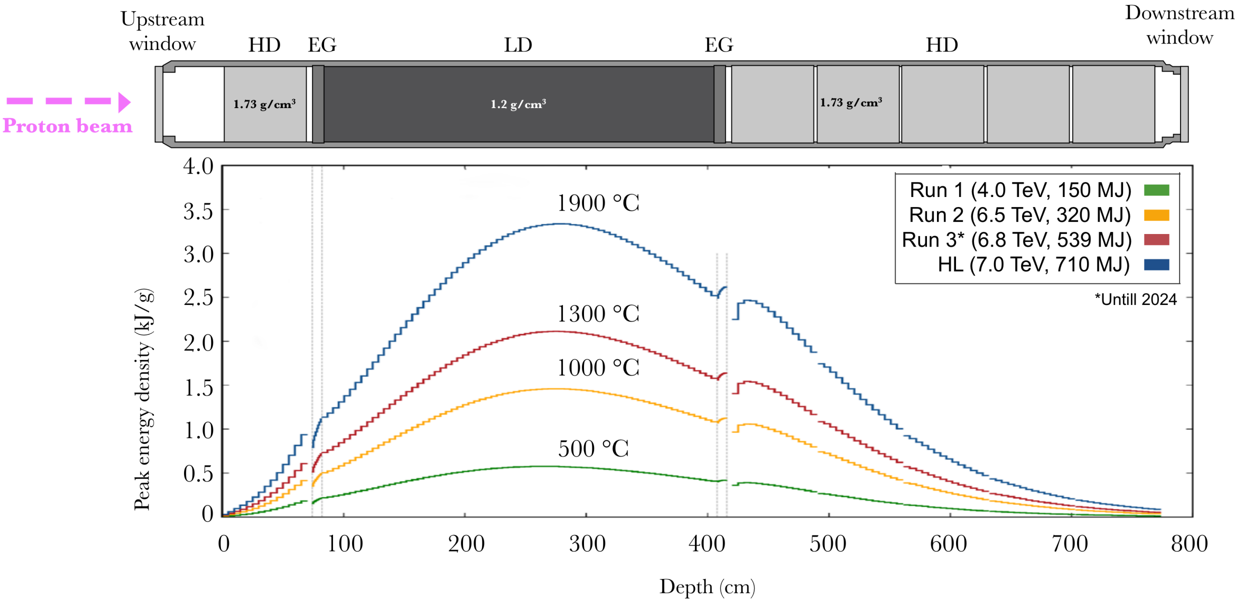

Figure 3.1: Overview of simulated energy deposition obtained from Monte Carlo analyses of particle–matter interactions from [32]. Peak transverse energy density as a function of longitudinal position along the beam dump for the different LHC runs. The maximum value occurs in the low-density (LD) sector at approximately 290 cm during HL operation, reaching 3.4 kJ/g, corresponding to an estimated temperature of approximately 1 900 °C.

3.2 Beam Dump Intensity Distribution

The loading experienced by the beam dumps is directly linked to the distribution of dumped beam intensities, as higher-intensity dumps contribute disproportionately to cumulative material degradation compared to low-intensity events. As previously discussed, beam dumps can be broadly classified into two categories: operational dumps, performed to optimise physics output during operation, and protective dumps, initiated by the machine protection system in response to detected faults or abnormal conditions.

To characterise the operational conditions and corresponding load spectrum experienced by the beam dumps, histograms of dumped energy / intensities for 2024 and for the period of 2025 up to and including August are presented in Figure 3.2. The histograms show the number of dumps as a function of dumped energy, expressed

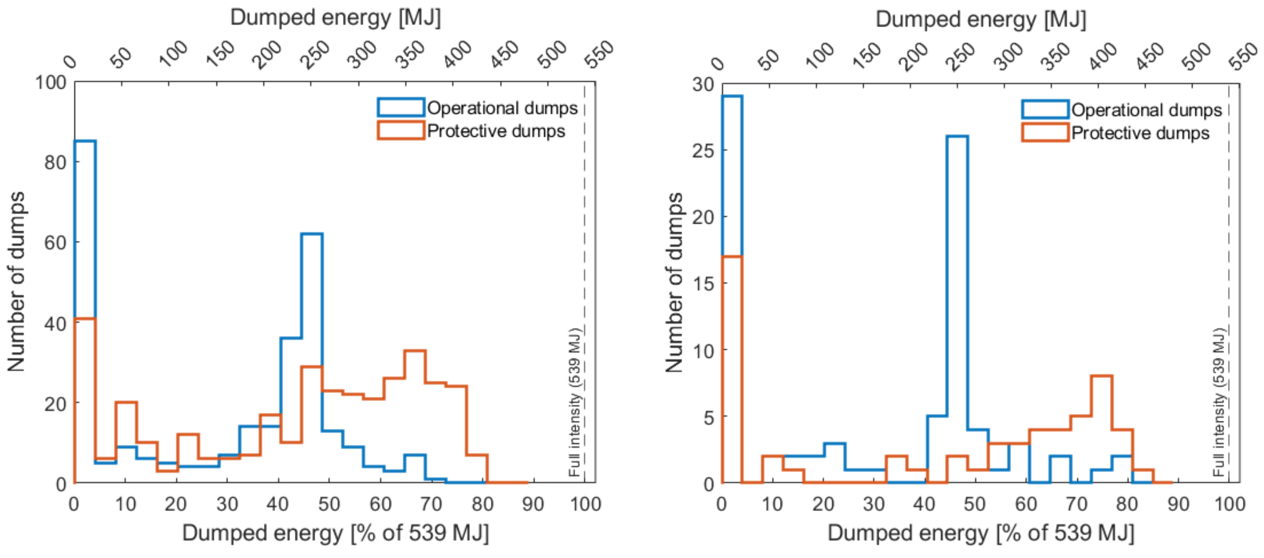

(a) 2024

(b) 2025 up to and including August

Figure 3.2: Overview of dumped intensities in Run 3 in year 2024 and 2025 based on type of beam dump.

as a percentage of full beam energy / intensity. The following general trends can be identified, and similar characteristics are observed for both operational and protective dumps across the analysed periods:

- A large number of very low-intensity dumps are observed in both categories. These events are most likely associated with commissioning and testing activities prior to full operation, during which accelerator components are verified at reduced beam intensity. Due to their low stored energy, such dumps are expected to contribute negligibly to cumulative damage in the beam dumps.
- Around 50% of full-intensity, a noticeable concentration of operational dumps is observed. This behaviour is consistent with the operational strategy described in Section 2.6, where the beam is intentionally dumped to allow reinjection of a fully populated beam and thereby optimise collision rate and physics output.
- At intensities exceeding 50% of full-intensity, the histograms are predominantly characterised by protective dumps. Owing to their high stored energy, these events are expected to contribute most significantly to cumulative damage in the beam dumps. It should be noted, however, that dumps at or near full-intensity are rare, with very limited recorded events above 80% of full-intensity.

It should be noted that Run 3 has been among the most successful operational periods in the history of CERN in terms of physics output and machine reliability. However, the extensive upgrades implemented in preparation for HL may alter future operational patterns. Consequently, historically derived load spectra may not be fully representative of HL conditions. Nevertheless, they provide a useful baseline for defining conservative design assumptions.

For the purpose of this work, the detailed load spectrum has been simplified into two representative intensity / stored energy levels with corresponding annual occurrence estimates for HL operation:

- **Half-intensity dumps:** All dumps below 50% intensity are conservatively assumed to occur at 50% of maximum intensity in HL. This corresponds to an estimated 340 dumps per year.
- **Full-intensity dumps:** All dumps between 50% and 100% intensity are conservatively assumed to occur at full-intensity in HL. This corresponds to approximately 100 dumps per year.

The definition of an appropriate load spectrum is essential for cyclically loaded components such as the beam dumps, where fatigue life predictions depend not only on the total number of load cycles but also on the magnitude of each cycle. The load spectrum defined above forms a central input to the fatigue assessments presented in this report.

3.3 Beam Dilution Failure Scenarios

Despite being designed for high operational reliability, the magnetic dilution kicker system cannot completely eliminate the possibility of faults [27]. Any loss of kicker functionality modifies the sweep pattern and can locally increase the energy density deposited in the beam dump core due to reduced transverse spacing between the bunches [42]. A reduction in dilution performance typically originates from electrical or triggering irregularities within the kicker system [29, 21]. Two principal failure mechanisms are identified:

- **Unintended (erratic) triggering:** Spurious activation of a kicker magnet without a dump request, often associated with charge accumulation effects. If such an event precedes a scheduled beam dump, the affected magnet may not restore its nominal operating charge prior to the beam dump, resulting in reduced kick strength during the actual extraction [29, 42].
- **High-voltage flashovers:** Electrical flashovers occurring during a beam dump, where high-voltage discharges generate plasma arcs between components such as feedthroughs. Because multiple kickers may share a common vacuum tank, such breakdowns can propagate and disable more than one magnet, thereby altering the beam sweep [29, 42].

Although rare, events involving the simultaneous loss of up to three dilution kickers have been recorded [21]. From a system safety perspective, the dilution system is classified as Safety Integrity Level 4 (SIL4) in accordance with IEC 61508 [26]. This rating corresponds to an expected probability of catastrophic failure, defined as complete loss of dilution resulting in full beam impact at a single location, of less than one event in 100 000 years, assuming approximately 440 beam dumps annually. Partial dilution failures has an estimated rate below one event per year.

An assessment of dilution-related failures recorded between 2010 and 2025 identified a total of eight events, seven of which occurred during the period 2014–2016. Of these, seven affected the horizontal kicker system (MKBH), while only a single event involved the vertical system (MKBV). Between 2018 and 2021, the dilution system underwent a series of upgrades aimed at mitigating the identified failure mechanisms. Since the completion of these improvements, no further dilution-related failures have been reported [21].

It is important to recognise that the energy density generated within the beam dump depends not only on the effectiveness of the dilution system, which determines how closely the beam is swept across the absorber surface but also on the total stored beam energy. The resulting local energy density is therefore governed by both beam intensity and the transverse spacing of successive bunch impacts. Historically, all recorded dilution failures have occurred during low-intensity dumps, and none have exceeded the peak energy density associated with a nominal sweep full-intensity Run 3 dump under a properly functioning sweep pattern.

Although statistically infrequent, dilution kicker malfunctions constitute a governing design case for the HL beam dump system. The LBDS incorporates four horizontal and six vertical dilution kickers, which together generate the beam sweep. According to the functional specification for the upgraded system, the most critical credible failure scenario corresponds to the loss of two horizontal kickers [21].

Figure 3.3 presents simulated sweep patterns for selected cases, including the nominal configuration (green box) and the accidental scenario with two horizontal kickers unavailable (red box). In the failure case, the beam dilution in the horizontal direction is significantly reduced, while the vertical sweep remains unaffected. The resulting sweep pattern produces a more concentrated impact region. The diagonal axis in the figure indicates the associated peak energy density for each configuration. A comprehensive overview of all evaluated dilution failure cases is provided in Appendix A.1.

The consequences of this failure mode under full-intensity HL conditions are quantified in Table 3.2, based on the peak energy densities obtained for both the LD and HD sectors (see Figure 3.1). In the event of two missing horizontal kickers, the peak energy density increases significantly. In the LD sector, the maximum energy density rises from 3.4 kJ/g to 5.7 kJ/g, corresponding to a 68% increase, while in the HD sector the increase is 57%, from 2.3 kJ/g to 3.6 kJ/g. These elevated energy densities translate directly into higher loading of the beam dumps. Consequently, this accidental scenario represents a critical scenario and must be considered for the HL design [32, 42].

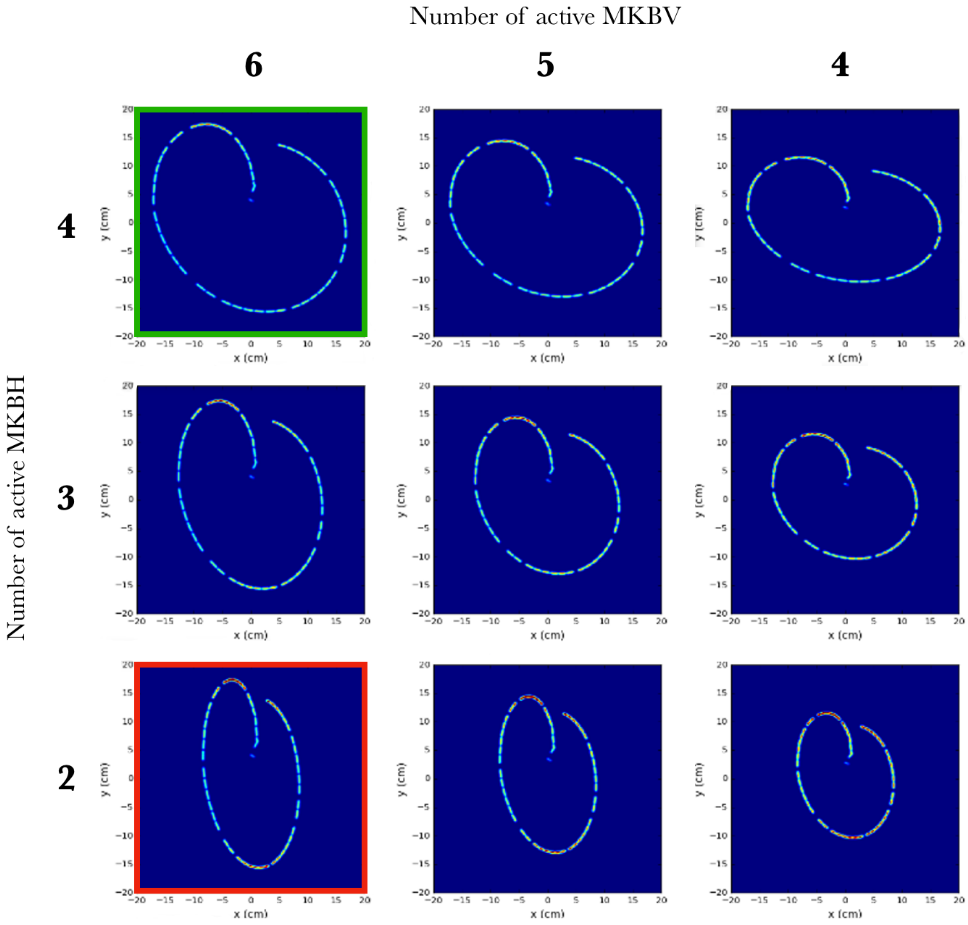

Figure 3.3: Subset of dilution sweep patterns showing the nominal case (green box) and a accidental scenario with two missing horizontal kickers (red box) [21, 42].

Table 3.2: Energy density in different sectors of the beam dumps under nominal and accidental scenarios [21].

Sector	Nominal [kJ/g]	Accidental [kJ/g]	Increase [%]
Low-density sector	3.4	5.7	68
Downstream high-density sector	2.3	3.6	57

3.4 Response Upon Beam Impact

Upon beam extraction, the full stored beam energy is deposited in the dump within 86 µs, corresponding to the time interval between the first and last proton impact. This extremely short deposition time results in a transient thermo-mechanical loading of the device. The structural response of the beam dump can be separated into two characteristic time scales:

- **Quasi-static response (hour scale):**
  During the impact, the kinetic energy of the beam is converted into heat, generating steep temperature gradients within the core materials. Peak temperatures in the impact region along the sweep approach 1300 °C in Run 3, while adjacent material remains near room temperature. Over time, thermal conduction redistributes the energy in the device and into the enclosing vessel, after which it is removed by the external air-cooling system. The resulting temperature rise in the vessel causes gradual thermal expansion of the device, leading to an axial elongation on the mm scale. This deformation is driven primarily by uniform heating and does not introduce significant stresses since the beam dump is free to expand on the wire ropes of the support structure [55, 45].
- **Dynamic response (millisecond scale:)**
  Approximately 4.6% of the total stored energy, corresponding to about 25 MJ in Run 3, is deposited in the vessel, while the vast majority of the remaining energy is absorbed by the core and shielding. The rapid energy deposition increases the vessel wall temperature up to 120 °C adiabatically within the 86 µs. Because the thermal expansion occurs on a very short time scale, it is initially constrained by the inertia of its own structure, leading to the development of a compressive stress state in the vessel. As this stored elastic energy is converted into motion, the natural vibration modes of the beam dump are excited [55, 45].

As the quasi-static response does not impose critical structural demands, the present work focuses on the dynamic behaviour of the beam dump, which has been identified as the governing load case with the most significant challenges to the structural integrity of the devices.

3.5 Energy Deposition Studies

The FLUKA code [3, 17, 18] is a general-purpose Monte Carlo simulation tool developed at CERN for modelling particle transport and interactions with matter. It implements microscopic physics models to describe hadronic and electromagnetic processes with high fidelity, ensuring that predicted responses are grounded in established physical principles [17]. The interaction mechanisms outlined in Section 2.4, including angular scattering, atomic displacement, production of residual nuclei, secondary particle cascades, and energy deposition, are inherently captured within the FLUKA framework together with their associated probability distributions.

In this work, all FLUKA energy deposition studies were performed by a dedicated section at CERN and subsequently provided to the author for interpretation and thermo-mechanical evaluation of the beam dumps.

Table 3.3 summarises the global energy deposition distribution for the present beam dump design. Approximately 80% of the total stored beam energy is absorbed within the dump itself, while around 16% is deposited in the surrounding shielding. Approximately 3–4% of the energy is dissipated through nuclear collisions, and roughly 1% is carried away by neutrinos. Only negligible fractions are transferred to the cavern, molasse, and air [32].

Table 3.3: Energy deposition distribution for the present beam dump design [32].

Region	**Energy fraction [%]**
Dump*	79.5
Shielding	15.8
Nuclear collisions	3.3
Neutrinos	1.0
Cavern**	< 0.2
Molasse	< 0.05
Air	< 0.02

* Dump includes the core materials and the vessel.
** Cavern refers to the surrounding underground infrastructure.

A more detailed breakdown of the energy deposited within the dump assembly is provided in Table 3.4. The fractions are expressed relative to the total stored beam energy. Approximately 73% of the energy is absorbed by the carbon-based core materials, whereas 4.6% is deposited directly in the vessel. The corresponding absolute energy values are shown for both Run 3 and HL.

Table 3.4: Energy deposition distribution within the present dump configuration. Run 3 total stored energy: 539 MJ. HL total stored energy: 710 MJ [32].

Component	Fraction [%]	Run 3 [MJ]	HL [MJ]
Upstream high-density sector	0.9	4.9	6.4
Low-density sector	40.3	217	286
Downstream high-density sector	32.9	177	234
Upstream window	< 0.01	< 0.05	< 0.07
Downstream window	0.04	0.22	0.28
Vessel	4.6	24.8	32.7
Support structure	0.7	3.8	5.0

A thorough understanding of the energy deposition distribution is essential when evaluating design modifications for the HL beam dumps. Even moderate shifts in deposition patterns can significantly alter local temperature fields and, consequently, the resulting thermo-mechanical response of the device.

3.6 Integration Constraints

A primary limitation on the design space for the HL beam dumps arises from the geometric constraints of the existing underground cavern. As the cavern is located approximately 100 m below ground level, enlargement is not a viable option.

Figure 3.4 illustrates the integration of the Run 3 beam dump within the shielding blocks, shown in both isometric and front views. While the upstream face of the dump remains unshielded to allow beam impact, all other surfaces, including the downstream end, are enclosed by shielding to confine radiation and activation to controlled regions. The limited spatial clearance is evident, particularly in the horizontal direction where the gap between the support structure and the shielding blocks is on the order of only a few millimetres. This effectively constrains any increase in outer dimensions of the dump or its support structure.

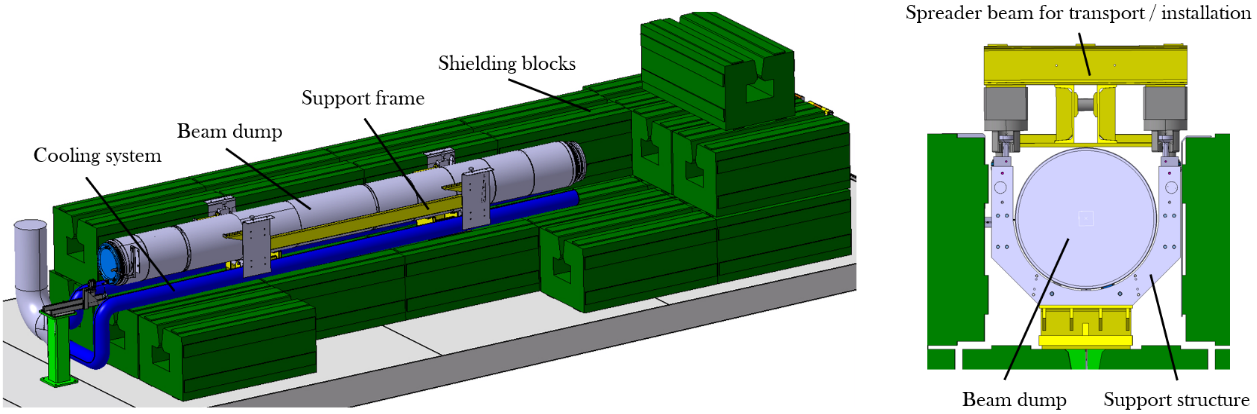

(a) Isometric view of the Run 3 beam dump installed within the shielding. Shielding blocks on the right-hand side and above the beam dump are hidden.

(b) Front view of the beam dump integrated within the shielding.

Figure 3.4: Integration of the Run 3 beam dump within the shielding configuration.

A broader overview of the cavern layout is presented in Figure 3.5. The beam exits the extraction line and travels through air before impacting the installed dump. The impact location is fixed due to the trajectory of the extraction line and the beam sweep generated by the dilution kicker system over approximately 600 m. Therefore, repositioning of the dump is not feasible. In addition, space must be reserved for a spare beam dump positioned on the left hand side of the operational device, and sufficient vertical clearance must be maintained to allow operation of the overhead crane running along the cavern ceiling. These constraints further restrict permissible modifications to the dump geometry.

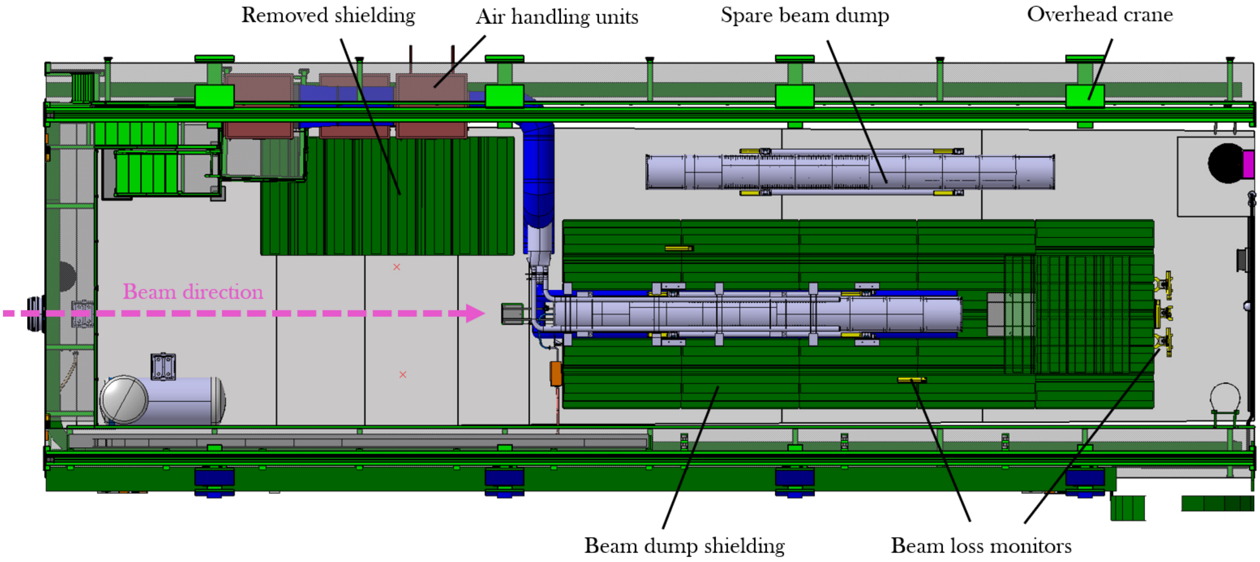

Figure 3.5: Top view of the beam dump cavern showing the installed beam dump, shielding configuration, and surrounding infrastructure. The top shielding blocks are temporarily displaced for visualisation purposes.

Based on these considerations, the following geometric constraints were defined for the HL beam dump design:

- The HL beam dump must comply with the existing shielding configuration.
- The overall width and height of the dump assembly, including its support structure, shall not exceed those of the Run 3 design.
- Modification of the shielding blocks is excluded due to activation and associated radiation exposure risks during handling.
- A modest extension in the longitudinal direction, on the order of 0.5 m, may be considered if required.

These spatial limitations significantly narrow the feasible design space and must be respected throughout the development of the HL beam dump configuration.

CHAPTER

4

System Requirements and Delimitations

This chapter consolidates the boundary conditions established in the problem analysis and translates them into system requirements for the HL-LHC beam dump design. The problem analysis defined the increased stored beam energy, characterised the expected load spectrum, identified critical dilution failure scenarios, and established the dynamic thermo-mechanical response as the governing design case. In addition, the energy deposition characteristics and geometric integration constraints have been determined and form the framework within which the upgraded HL design must be realised.

Functional Requirements

The functional requirements define the primary functions the HL beam dump system shall fulfil during operation.

- **Beam absorption:** The beam dump shall safely absorb the HL beam under both nominal and defined accidental conditions without loss of structural integrity.
- **Dilution failure tolerance:** The beam dump shall withstand the dilution failure involving the loss of two horizontal dilution kickers without catastrophic failure.
- **Radiation containment:** The system shall confine induced activation and radiation within the dump and shielding.
- **Energy dissipation:** The system shall transfer deposited energy to the cooling system without exceeding allowable material temperature limits.

Non-Functional Requirements

Non-functional requirements define the performance limits, lifetime criteria, and integration constraints necessary for reliable operation.

- **Lifetime requirement:** The beam dump shall withstand 10 years of HL operation corresponding to:
  - 340 beam dumps per year at half-intensity (355 MJ) with nominal sweep.
  - 100 beam dumps per year at full-intensity (710 MJ) with nominal sweep.
  - 1 beam dump per year at full-intensity with accidental sweep.
- **Fatigue resistance:** The vessel shall satisfy fatigue life requirements under the cyclic thermo-mechanical stresses induced by its dynamic response.
- **Temperature limits:** Peak temperatures in structural components shall remain below material-specific allowable limits to prevent loss of strength and avoid potential damage.
- **Geometric compatibility:** The overall outer dimensions of the beam dump and its support structure shall not exceed those of the Run 3 configuration. Modifications to the existing shielding are excluded. A longitudinal extension of up to 0.5 m may be considered.

Delimitations

The scope of this thesis is limited to the structural and thermo-mechanical assessment of the beam dump vessel and its interfaces. The following limitations apply:

- The analysis focuses on the dynamic response induced by rapid energy deposition in the vessel and the associated excitation of structural vibration modes.
- The quasi-static thermal response of the core materials is not treated in detail, as previous studies have demonstrated that it does not govern structural integrity.
- Qualification and optimisation of the carbon-based core materials for HL-LHC conditions have been addressed in [42, 40] and are therefore outside the scope of this work.

These requirements form the foundation for the development of the HL beam dump design and define the governing criteria for the subsequent engineering work. They ensure that the upgraded system can safely absorb the increased stored beam energy, withstand the associated thermo-mechanical loads, and tolerate defined accidental scenarios, while remaining compatible with the existing cavern and shielding configuration.

The system requirements can be condensed into the following research question:

How can the HL-LHC beam dump vessel be designed to withstand 10 years of cyclic dynamic behaviour under HL operational conditions, without compromising structural integrity or violating existing integration constraints?

Chapter

5

Finite Element Model

This thesis is built upon a global finite element model of the beam dump, established to characterise its overall dynamic response upon beam impact. The model includes the vessel, beam windows, and the shrink-fitted core materials. For the present work, the support structure is not included, as the focus is placed on the beam dump design. The global model provides the reference framework for subsequent, more detailed sub-models developed to evaluate the behaviour of the circumferential vessel welds, beam windows and their interface with the vessel flange. These sub-models use derived loading conditions from the global model.

The finite element model was originally developed by the author to assess the Run 3 partial penetration welds. It was later used by Braken [22] to benchmark it against measurements from the operational Run 3 dumps. The modelling framework was subsequently extended for the support structure assessments by Bergmann [52] and Kappelgaard [36]. The contributions of these authors are referenced in the relevant sections of this chapter.

In this thesis, the Run 3 design is first simulated and validated against measurement data installed on the operating beam dumps. Next, the difference in dynamic response between Run 3 and HL beam parameters are assessed to clarify if the Run 3 beam dumps are suitable for operating in HL. This establishes a reference configuration for subsequent design variations, including geometry and material modifications. All variants are derived from the same global model simulation framework. The model for the Run 3 configuration is presented in this chapter, its dynamic response in Chapter 6, and subsequent HL design modifications in Chapter 7.

5.1 Finite Element Methodology

This section describes the finite element methodology and simulation workflow used in this work. It focuses on the thermo-mechanical coupling, time integration scheme, numerical stability criterion, and the balance between accuracy and computational cost. The global model is simulated using LS-Dyna with the support of Ansys and LS-PrePost for the setup of the model.

5.1.1 Governing Equations

Due to the dynamic nature of the beam dumps, inertia effects must be included, and both velocities and accelerations are considered. The transient non-linear equilibrium equation is written as:

$$[M]\,\ddot{U}(t) + [C]\,\dot{U}(t) + F_{\text{int}}\big(U(t), T(t)\big) = F_{\text{ext}}(t) \tag{5.1}$$

where $[M]$ and $[C]$ denote the mass and damping matrices, $\ddot{U}(t)$, $\dot{U}(t)$, and $U(t)$ are the nodal accelerations, velocities, and displacements, and $F_{\text{ext}}(t)$ is the external load vector. The internal force vector depends on the instantaneous displacement field and the evolving temperature field through temperature-dependent material properties.

The objective of the analysis is to determine the displacement field $U(t)$, from which strains and stresses are derived. Since Equation 5.1 represents a coupled system of second-order differential equations in time, it must be solved numerically by discretising the time domain and integrating the equation of motion incrementally.

5.1.2 Explicit Time Integration

For the highly dynamic beam impact problem, the mechanical response is solved using the explicit central difference scheme [34]. In this formulation, the state at each time step is computed from known quantities of the previous step without solving a global system of equations. A lumped mass matrix allows each degree of freedom to be updated independently, resulting in low computational cost per increment. The central difference scheme uses a staggered formulation, where displacements and accelerations are evaluated at full time steps and velocities at half steps. This avoids solving a coupled system between velocities and displacements and keeps the method fully explicit.

The explicit scheme is conditionally stable and must satisfy the Courant–Friedrichs–Lewy (CFL) criterion [25], which requires that stress waves do not travel more than one element within a single time step. The stable time step is limited by the smallest element size and the material wave speed and is commonly expressed as $\Delta t_{\max} \leq f\,\Delta h_{\min}/c$, where $\Delta t_{\max}$ is the maximum stable time step, f is a safety factor (taken as 0.9), $\Delta h_{\min}$ is the smallest characteristic element length in the mesh, and c is the material wave speed which for elastic materials with a constant bulk modulus can be calculated from the Young's modulus E, the material density ρ, and the Poisson's ratio ν by [34]:

$$c = \sqrt{\frac{E(1-\nu)}{\rho(1+\nu)(1-2\nu)}} \tag{5.2}$$

Mesh refinement thus reduces the stable time step in addition to increasing the number of elements. Consequently, finer meshes significantly increase the computational cost in explicit simulations.

5.1.3 Element Technology

The beam dump is discretised using linear eight-node hexahedral solid elements with reduced integration [34]). These elements employ a single integration point and assume constant stress within the integration domain. The reduced integration formulation significantly lowers computational cost compared to fully integrated elements. Reduced integration elements can exhibit zero-energy deformation modes, commonly referred to as hourglass modes. These non-physical patterns do not contribute to strain energy and are therefore controlled using the standard Belytschko–Bindeman hourglass formulation available in LS-Dyna [34]. Hourglass energy is monitored throughout the simulations and is observed to remain small relative to the total internal energy, indicating that the physical response is not affected.

5.2 Model Development and Discretisation

The model setup was performed using Ansys SpaceClaim for geometry preparation, Ansys Mechanical for meshing, and the Ansys LS-Dyna interface within Ansys Workbench for defining analysis settings. The LS-Dyna input file was generated within the Ansys environment and subsequently imported into LS-PrePost for customised material assignment using the specific LS-Dyna material models before submission to the CERN high-performance computing facility.

5.2.1 Geometry and Discretisation

The model development was based on the detailed CAD geometry provided by the design office, shown in Figure 5.1. This geometry corresponds to the full manufacturing configuration and includes fillets, holes, chamfers, weld bevels, and other small-scale features.

The geometry was simplified prior to analysis. The external support structure was removed, as it is outside the scope of this work. The circumferential partial penetration weld (6.65 mm bevel) was omitted, and the vessel was modelled with continuous 12 mm thickness. This removes the local weld discontinuity but does not

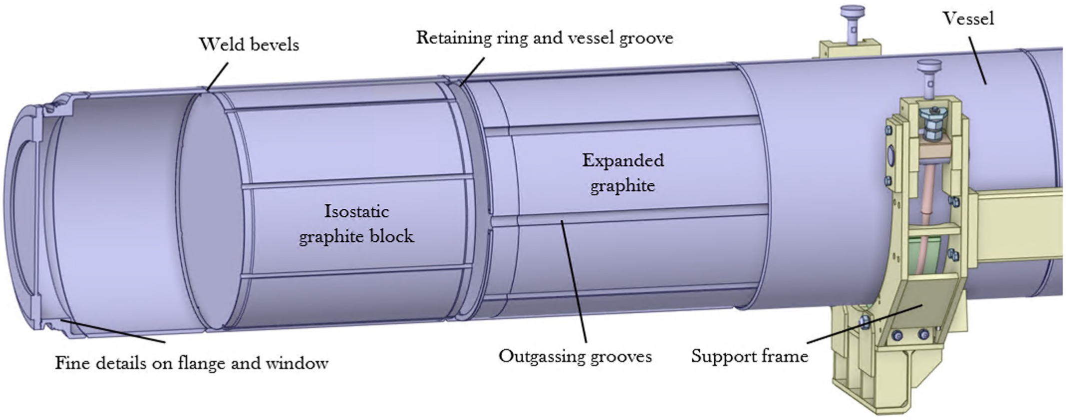

Figure 5.1: Detailed CAD geometry of the global model as provided by the design office. The view illustrates the upstream region, with selected vessel surfaces removed to expose the internal components.

affect the global stiffness or inertial response, which is the objective of the model. Local weld behaviour is assessed separately in Chapeter 7 for Run 3 and Section 10.2 for HL.

Small-scale features, including outgassing grooves, shrink-fit chamfers, and minor fillets in windows and flanges, were removed to facilitate meshing. The LD sector, originally consisting of 2 mm expanded graphite sheets, was represented as a continuous solid over its full 3.26 m length. A diameter interference of 1.2 mm was introduced for the IG blocks such that the outer diameter of the core exceeds the inner diameter of the vessel. The total mass of the simplified model was verified against the original CAD geometry with negligible difference, ensuring preservation of the inertial response.

As shown in Figure 5.2, the simplified geometry was extensively partitioned to control the mesh generation. Shared topology was enforced in the vessel, in the windows and in each core segment to prevent the introduction of contact interfaces. Longitudinal and circumferential partitions were introduced to ensure exact nodal matching between the outer surface of the core and the inner surface of the vessel.

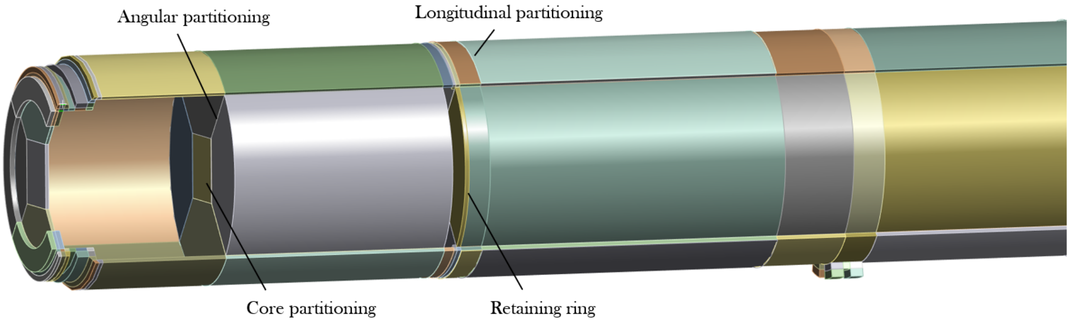

Figure 5.2: Simplified and partitioned geometry prepared for mesh generation. Partitions are introduced to support the meshing algorithm while maintaining conformal topology between components.

The discretised model is shown in Figure 5.3. The mesh consists purely of solid element with approximately 970k nodes and 870k hexahedral elements. The circumferential discretisation has 100 elements, while a longitudinal spacing of approximately 20 mm is used along the vessel axis. Four elements are introduced through the 12 mm vessel thickness. This through-thickness refinement is required to capture the through-thickness variation in energy deposition and resulting thermal gradients across the vessel wall. Considering the large diameter-to-thickness ratio of the vessel, four elements are deemed sufficient to capture bending behaviour. The windows were meshed coarsely in the initial model, as their local energy deposition was not expected to influence the global dynamic response, which is governed predominantly by the vessel.

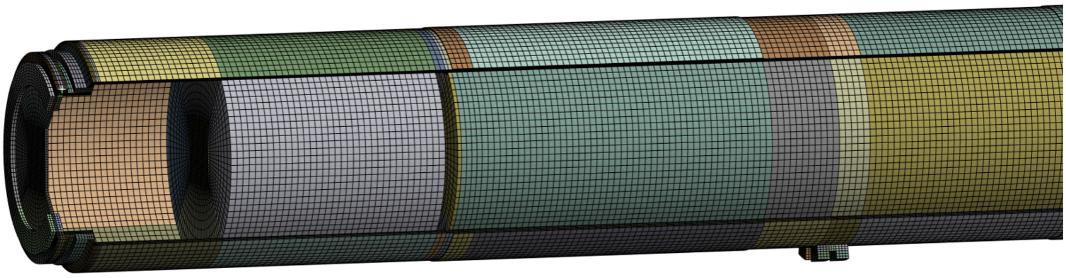

Figure 5.3: Final discretised global finite element model. Mesh consists of linear solid hexahedral elements.

The maximum stable time step was evaluated to 0.45 μs, or approximately 2 200 time steps to simulate 1 ms of physical response. Some additional information on the mesh is provided in Appendix A.3.

5.2.2 Materials

The material models included in the global model are defined using different material formulations. Titanium Grade 5 and SS318LN are modelled using temperature-dependent isotropic elastic formulations to account for the variation in properties with temperature. Isostatic graphite and flexible graphite are modelled using linear elastic isotropic material formulations. For the scope of the present global analysis, temperature dependence or other non-linearities are not included in the material model for these materials.

Table 5.1 summarises the materials used in the model and provides references to the corresponding appendix sections where detailed material properties are documented.

Table 5.1: Overview of material models used in the global model and corresponding appendix references.

Material	Component	Appendix reference
Titanium Grade 5 (Ti-6Al-4V)	Beam windows	Appendix A.4.1
Stainless steel (318LN)	Vessel, flanges, retaining rings and cable guides	Appendix A.4.2
Isostatic Graphite	High-density sectors	Appendix A.4.4
Extruded Graphite	Low-density sector retaining plates	Appendix A.4.4
Flexible Graphite	Low-density sector	Appendix A.4.5

5.2.3 Contact Modelling

Figure 5.4 and Table 5.2 provide an overview of the implemented contacts. Contact definitions are introduced at material interfaces where load transfer cannot be represented through shared topology. In addition, selected interfaces such as the vessel–cable guide and vessel–retaining ring are intentionally modelled using contact formulations to enable extraction of contact forces for subsequent sub-modelling.

The interference fit between the vessel and the isostatic graphite blocks is modelled using a special contact interference formulation. This contact type allows initial geometric penetration between bodies, which is required to represent the shrink-fit configuration. During an initial stabilisation phase, the contact stiffness is gradually ramped using a load curve until penetration is eliminated and the desired interference contact pressure is established.

All contact interfaces in the global model are implemented using penalty-based formulations. In frictional contacts, limited penetration is permitted and resisted by normal penalty stiffness combined with a Coulomb friction model. Tied contacts enforce displacement compatibility between surfaces through a high penalty stiffness, thereby representing bonded or bolted interfaces.

Further details on the contacts can be found in Table A.13 of Annex A.5 while individual contact algorithms referenced can be found in the LS-Dyna manual [7].

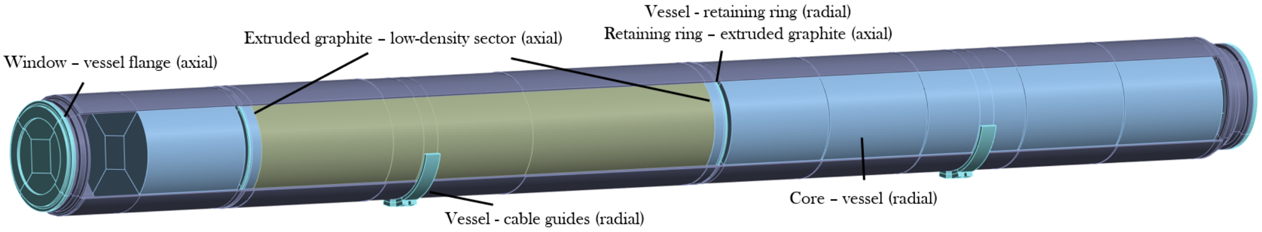

Figure 5.4: Overview of contacts implemented in the model. Contact definitions referred in Table 5.2.

Table 5.2: Overview of contact interfaces implemented in the global model.

Interface	Type	Purpose / Remarks
Core – Vessel (radial)	Interference (frictional)	Represents the shrink-fitted configuration between the IG core and the vessel, establishing the initial interference pressure.
Window – Vessel Flange (axial)	Bonded (bolted)	Represents the bolted connection and enables extraction of interface forces for submodelling of the window–vessel joint.
Vessel – Cable Guides (radial)	Bonded (bolted)	Allows extraction of interface forces for validation of cable guide welds (outside the scope of this thesis).
Vessel – Retaining Ring (radial)	Bonded (bolted)	Enables force extraction for submodelling and structural dimensioning of the retaining ring (outside the scope of this thesis).
Extruded Graphite – LD Sector (axial)	Frictional	Allows relative sliding between graphite sectors during dynamic loading.
Retaining Ring – Extruded Graphite (axial)	Frictional	Allows relative sliding between the retaining ring and graphite sector.

5.2.4 Loads and Boundary Conditions

The beam load is introduced as volumetric heat generation and not as a mechanical impulse. Although the protons carry substantial kinetic energy, the energy is converted into internal energy through particle–matter interactions as described in Section 2.4, and does not transfer momentum to the device. The energy deposition is imported to the analysis from three-dimensional energy deposition maps obtained from FLUKA simulations, which are Monte-Carlo based simulations that characterize the energy deposition in accordance with the description of beam-matter interaction provided in 2.4. More details on the FLUKA code is provided in Appendix A.2 and in [3, 14, 17]. These maps are provided as discretised volumetric bins containing locally averaged energy density values.

Energy deposition in the core is confined to the narrow beam sweep region, while the surrounding material remains only moderately heated (see Figure 2.8). The resulting thermal expansion is locally constrained by the surrounding colder material, keeping the induced stresses confined to the heated zone. As heat diffuses through the core on a time scale much longer than the structural dynamic response, these stresses relax and do not generate motion that affects the vessel. Consequently, energy deposition in the core does not influence the dynamic behaviour. Excluding it also avoids the need for a substantially refined mesh to resolve the localized sweep pattern, thereby reducing computational cost. This assumption was confirmed through the sensitivity analyses in Section 5.4. In contrast, the vessel cross-section is heated more uniformly and can expand freely, as it is not constrained by surrounding material like the core. The rapid thermal expansion initially generates compressive stresses, which are subsequently released into motion and excite the structural vibration modes. Based on these considerations, the energy deposition is applied to the vessel, flanges, and windows, but not to the core.

Since FLUKA data do not include temporal information, the deposited energy is applied in the model as a linear ramp over the 86 µs beam duration. The bin-based data are imported into Ansys using a dedicated CERN-developed Python extension. Because the FLUKA binning rarely coincides with the finite element mesh, spatial interpolation is required. The interpolation procedure is verified through:

- Visual inspection of bin positioning and interpolated power density to verify spatial alignment and detect artefacts. Appendix A.6 provides some figures of the energy interpolation with some interpretation.
- Quantitative comparison of peak energy density and total energy between the original FLUKA data and the interpolated Ansys results.

Table 5.3: Comparison of peak power density by components between FLUKA and Ansys for Run 3 beam parameters.

Component	**FLUKA [MW/cm^3]**	**Ansys [MW/cm^3]**	**Diff [%]**
Upstream window	4.18	1.60	-89.3
Upstream flange	0.22	0.21	-7.2
Vessel	4.68	4.56	-2.7
Downstream flange	1.51	1.48	-2.1
Downstream window	4.11	3.71	-10.3

Table 5.4: Comparison of total deposited energy by components between FLUKA and Ansys for Run 3 beam parameters. Percentage share only considers the energy share between the presented components.

Component	**FLUKA [kJ]**	**Ansys [kJ]**	**Diff [%]**	**Share [%]**
US Window	9.8	10.4	6.3	0.04
US Flange	26.9	26.2	-2.7	0.1
Vessel	24 800	24 859	0.22	97.2
DS Flange	415.1	412.4	-0.64	1.6
DS Window	269.1	258.3	-4.1	1.0
Total	25 521	25 563	0.16	100

Table 5.3 shows a component-wise summary of the quantitative comparison of the peak power density between FLUKA and the interpolated results on the Ansys model. Table 5.4 shows the corresponding integrated energy deposition. The following observations can be made from the tables:

- The peak energy density in the windows is not accurately captured due to the coarse mesh. Since the windows are assessed separately, resolving their peak values is not required for the global dynamic response. In contrast, the peak energy density in the vessel is captured within 3%.
- The total deposited energy is dominated by the vessel, which accounts for 97% of the energy among the components considered. The integrated vessel energy is captured within 0.2%. Since the vessel is expected to govern the dynamic response, this accuracy is sufficient. Although the peak density in the windows is not well resolved, their integrated energy is captured within 6%.

In addition an internal pressure of 0.2 bar is applied to all inner surfaces of the beam dump to account for the pressure of the internal atmosphere. The support structure is not included in the global model, and gravitational loading is neglected. Given the high structural stiffness of the assembly, self-weight effects are negligible compared to the transient thermal loading. This assumption was also confirmed through the sensitivity analyses presented in Section 5.4. No mechanical boundary conditions are imposed, as the only applied load is internal heat generation. Result quantities must therefore be interpreted in a relative sense, since small rigid body motion may occur.

5.3 Simulation Phases

The simulation is carried out in two sequential phases within a single analysis. First, a dynamic relaxation phase is used to establish the pre-stressed state caused by the shrink fitting of the IG core. This is followed directly by a transient analysis of the beam impact.

- **Dynamic relaxation phase:**
  Dynamic relaxation (DR) is used to establish the initial stress state of the shrink fitting prior to the

transient beam impact analysis. DR provides a quasi-static solution using explicit time integration. Artificial damping is applied through nodal velocity scaling at each time step, progressively dissipating kinetic energy while allowing the system to approach equilibrium. Convergence is assessed using the ratio of kinetic energy to internal energy. When the kinetic energy becomes negligible relative to the internal energy, the system is considered to have reached static equilibrium. The dynamic relaxation phase is solved in pseudo-time. After convergence, the resulting nodal displacements, strains, and contact states are automatically transferred to the transient phase, while all nodal velocities are reset to zero [34].

- **Transient thermo-mechanical phase**
  After dynamic relaxation, the beam impact is simulated by applying the energy deposition. The power input is applied to each element during the first 86 µs, corresponding to the beam impact duration. The thermo-mechanical analysis is fully coupled, meaning that the thermal and structural fields are solved within the same simulation and continuously exchange information at each thermal timestep. The thermal problem is solved implicitly with an initial timestep of 0.5 µs, a minimum timestep equal to the structural timestep, and a maximum timestep of 100 µs. The maximum allowable temperature change per thermal timestep is limited to 1 °C to ensure accurate temperature evolution and coupling to the mechanical response. Within each thermal increment, the structural solver advances with explicit timesteps until the end of the thermal interval, after which the thermal field is solved. If the temperature change is below 1 °C, the step is accepted; otherwise, the timestep is reduced and the structural solution is rolled back to the beginning of the thermal increment. This results in small thermal timesteps during the 86 µs beam impact and larger timesteps afterwards when the temperature field is nearly constant. The structural response is then simulated for 50 ms to capture the vibrational behaviour [34].

5.4 Sensitivity Analyses

Several sensitivity analyses were performed to establish the final modelling approach that was described in the previous sections:

- **Mesh sensitivity:** A mesh sensitivity study was conducted by systematically refining and coarsening the global element density in two alternative model configurations to assess its influence on the dynamic response. Increasing the mesh density resulted in only negligible changes in the monitored scalar output quantities, indicating that the baseline mesh is of sufficient density. In contrast, mesh coarsening, particularly through a reduction in the number of elements across the vessel wall thickness, led to reduced resolution of through-thickness energy deposition gradients. This loss of resolution introduced some variations in the output scalar values and was therefore deemed undesirable. Consequently, the mesh described in Section 5.2.1 was adopted as sufficiently resolved for the objectives of the present analysis.
- **Energy deposition in the core:** A model variant featuring a refined mesh in the core region was developed to enable the application of energy deposition within the core materials. Comparison with the baseline model showed no significant change in the global dynamic response, supporting the conclusion that the stress state resulting from the energy deposition in the core remains confined to the sweep region of the core. The surrounding regions of the core remain comparatively cold, thereby preventing the translation of internal stresses into motion that could influence the overall dynamic behaviour. Since the inclusion of energy deposition in the core was found to have no measurable effect on the dynamic response, and because its exclusion permits a substantially coarser mesh and reduced computational cost, energy deposition in the core was omitted from the final model in Section 5.2.1.
- **Support structure coupling and gravity:** Simulations with and without the support structure were performed to assess potential coupling and dependency effects between them. Owing to the flexibility of the suspension wire ropes and the small vibration amplitudes (on the order of millimetres), the reaction forces transmitted through the ropes that could influence the dynamic response of the beam dump were found to be negligible. The inclusion of gravitational loading to pre-stress the suspension ropes and support structure, and to account for self-weight effects, likewise showed no measurable influence on the dynamic response. Consequently, the support structure and gravitational loading were excluded from the global model.

Chapter

6

Dynamic Response of the Run 3 Beam Dump and Model Validation

This chapter presents the finite element results for the Run 3 beam dump design under Run 3 full intensity beam parameters. Model validation is addressed throughout the chapter to demonstrate the accuracy and representativeness of the finite element model.

6.1 Shrink-Fit Induced Pre-Stress

Figure 6.1 shows the pre-stressed state after dynamic relaxation, with contours of equivalent von Mises stress. The shrink-fit interference between the upstream and downstream HD sectors and the vessel is fully resolved, resulting in increased vessel stresses. The LD sector, which is not shrink-fit, remains unstressed.

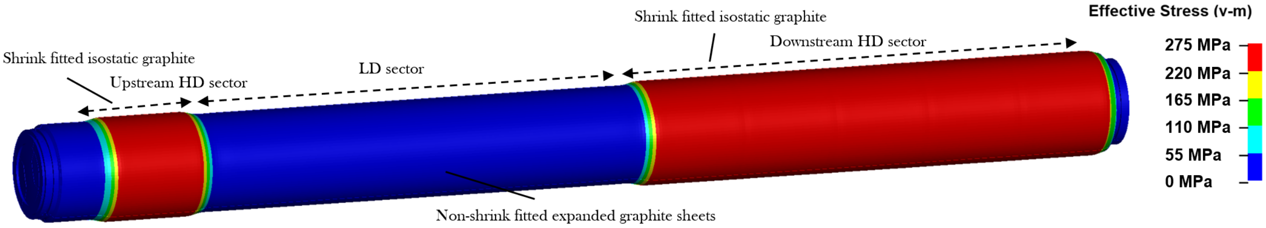

Sh

Figure 6.1: Pre-stress state after dynamic relaxation showing resolved shrink-fit interference with von Mises stress contours. Deformations scaled by 50 to visualise vessel expansion.

The vessel stress from shrink fitting was benchmarked against implicit analysis in Ansys Mechanical and analytical interference-fit calculations according to ASME/ANSI B4.1 [2] (Appendix A.7). The comparison focuses on the vessel wall von Mises stress and is summarised in Table 6.1.

Tool	**Vessel von Mises stress** [MPa]	**Difference [%]**
LS-DYNA explicit (DR)	261	–
Ansys Mechanical implicit	256	-1.92
ASME/ANSI B4.1 analytical	266	1.92

Table 6.1: Comparison of vessel von Mises stress due to shrink fitting of the isostatic graphite core in the 318LN vessel.

The agreement within ± 2% confirms that the dynamic relaxation procedure reproduces the correct interference state.

6.2 Transient Response and Mode Identification

After the dynamic relaxation phase, the transient phase starts. The thermal load corresponding to a full intensity Run 3 beam dump is applied to the vessel, flanges, and windows over 86 µs. The resulting temperature field is shown in Figure 6.2. The peak vessel temperature is 123 °C and occurs at the inner surface in the longitudinal mid-section, toward the top left side, where the beam sweep passes closest to the vessel. The temperature decreases toward both ends of the beam dump.

During the following milliseconds, the structural dynamic response develops. Figure 6.3 shows the von Mises stress field at 13 ms after beam impact. The response is a superposition of several natural vibration modes excited by the thermal shock. This results in highly oscillatory stresses and cyclic loading of the vessel.

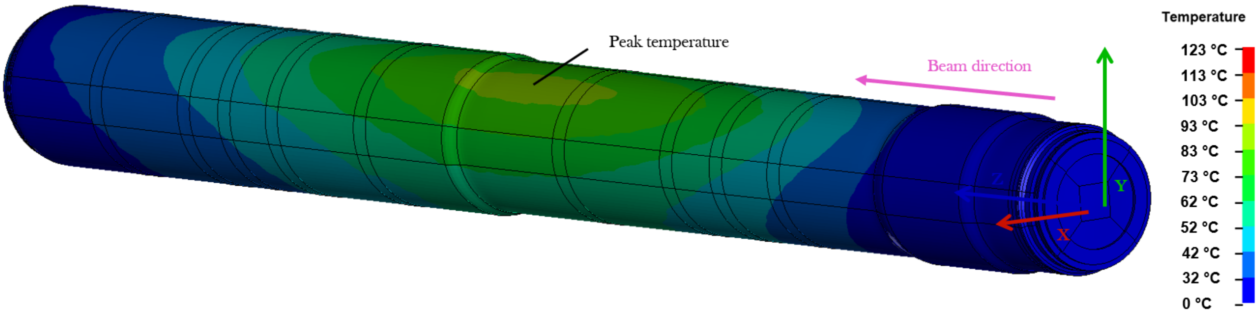

Figure 6.2: Temperature field after beam impact. The peak temperature occurs in the longitudinal mid-section toward the top left inner surface, where the beam sweep is closest to the vessel. Deformations are scaled by a factor of 50 to visualise the shrink-fit state.

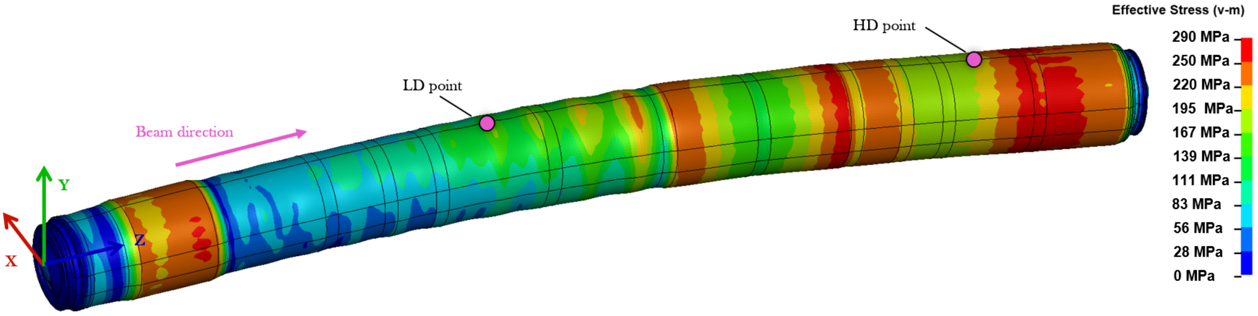

Figure 6.3: Von Mises stress at 13 ms after beam impact. The stress field results from the superposition of several excited vibration modes, producing cyclic loading of the vessel. Deformations are scaled by a factor of 50.

Figures 6.4, 6.5, and 6.6 identify the dominant vibration modes from nodal displacements and relate them to the local stress response. The evaluation is performed at two points shown in Figure 6.3. Both points are located on the top surface of the beam dump on the y-axis. One point is slightly upstream in the LD sector, and the other is near the downstream end in the HD sector.

Figure 6.4a shows the displacement components in the horizontal (x), vertical (y), and longitudinal (z) directions at the LD point. A distinct longitudinal mode appears in the z-direction. In the x- and y-directions, a low-frequency mode corresponds to a flexural mode, while a high-frequency mode, most pronounced in the y-displacement, corresponds to a radial (breathing) mode.

Figure 6.4b shows the corresponding stresses at the same location. The radial stress (y-stress) is negligible and therefore omitted, as radial displacement at this position is converted into hoop stress, captured in the x-direction. The longitudinal mode is clearly reflected in the z-stress. Immediately after beam impact, the z-stress shifts to compression because thermal expansion is constrained by inertia. As the motion develops, the z-stress oscillates with the longitudinal vibration. The x-stress (hoop stress) shows a very high amplitude response with stress ranges of 350 MPa and is related to the radial mode of vibration that is introduced as the vessel expands and contracts around its circumference.

Figures 6.5a and 6.5b show the same quantities for the HD point. The longitudinal mode (z-direction) appears at the same frequency as in the LD sector but with opposite displacement sign, since this point expands in the opposite direction. The flexural (x–y) and radial (y) modes have lower amplitudes. This is reflected in Figure 6.5b where the longitudinal stress (z-stress) has lower amplitude than in the LD sector, and the hoop stress (x-stress) oscillates with smaller amplitude but around a higher mean value. These differences between the LD and HD points are explained by three factors:

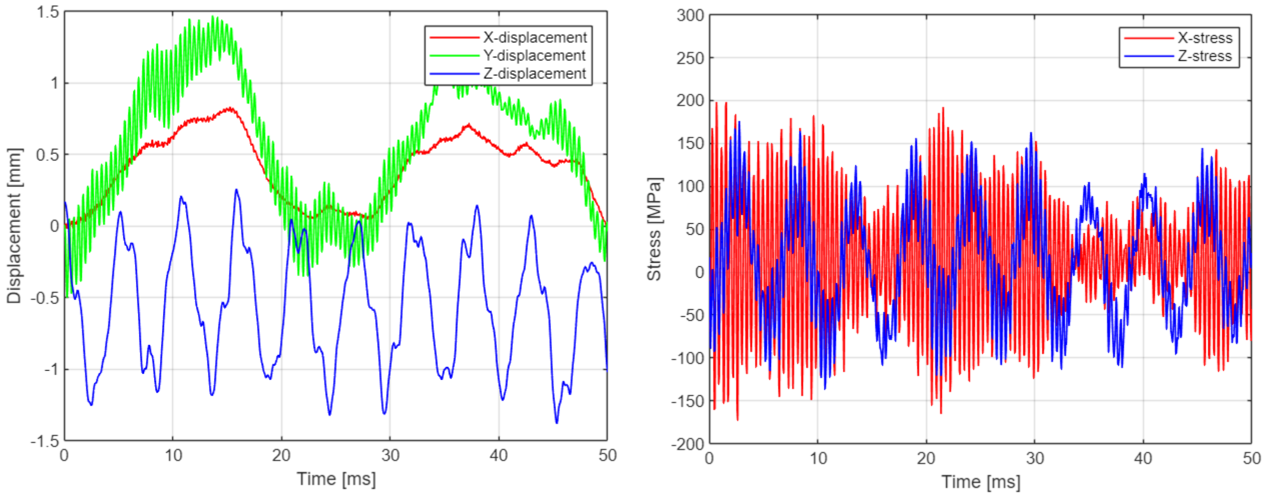

(a) Displacement components.

(b) Stress components.

Figure 6.4: Extracted transient results from the LD-point location from Figure 6.3. For both subfigures, the initial values are not zero because the structure is already deformed from the DR pre-stress phase.

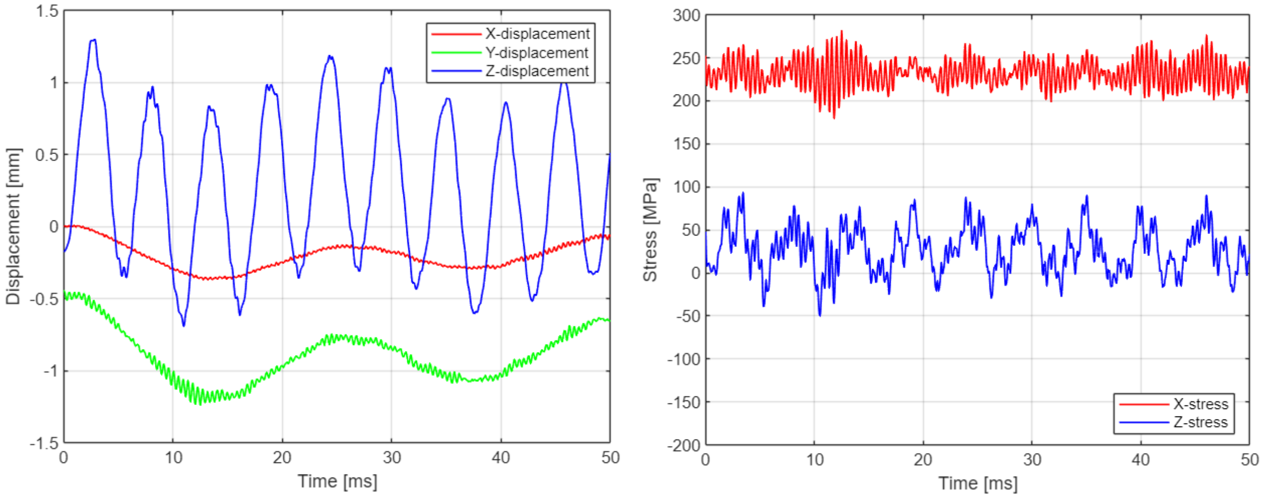

(a) Displacement components.

(b) Stress components.

Figure 6.5: Extracted transient results from the HD-point location from Figure 6.3. For both subfigures, the initial values are not zero because the structure is already deformed from the DR pre-stress phase.

- **Shrink-fitted configuration:** The isostatic graphite block introduces a tensile mean hoop stress (x-stress) through the initial interference fit and constrains radial (y) motion. This increases the mean stress level and reduces vibration amplitudes.
- **Non-uniform thermal field:** The thermal load varies with longitudinal (z) and circumferential position. The LD point is close to the peak energy deposition and therefore experiences a stronger excitation, particularly of the circumferential (radial) mode. The HD point is exposed to lower energy deposition and shows smaller stress amplitudes as a consequence.
- **Longitudinal position of the extraction point:** For longitudinal vibration (z-direction) of a beam with traction-free ends, the stress is zero at the ends and maximal at mid-span. The longitudinal centre of the beam dump (LD-point) therefore exhibits the highest z-stress amplitude, whereas the HD-point has lower amplitude of longitudinal stress.

Figure 6.6 shows the Fast Fourier Transformations of the displacement and stress signals for the LD point

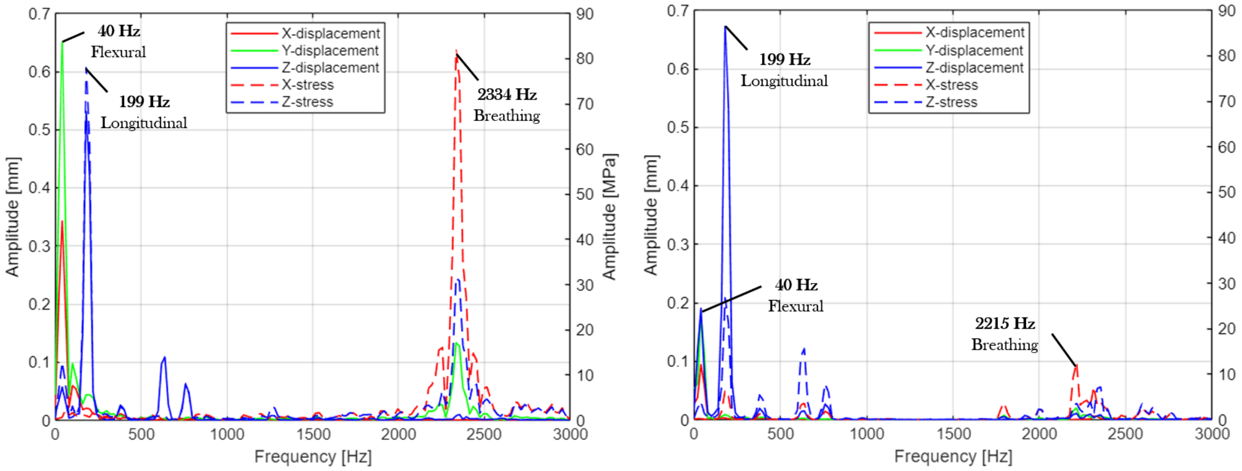

(a) LD point from Figure 6.4. **(b)** HD point from Figure 6.5.

Figure 6.6: Fast Fourier Transformations of displacement and stress components from Figure 6.4 (shown in Figure 6.6a) and Figure 6.5 (shown in Figure 6.6b).

(Figure 6.6a) and the HD point (Figure 6.6b). The following observations can be made for the LD point:

- **Flexural mode (40 Hz, x–y bending):** Clearly visible in the x- and y-displacements. Contributes only marginally to the longitudinal stress (z-stress).
- **Longitudinal mode (199 Hz, z-direction):** Dominant in the z-displacement and directly reflected in the primary oscillation of the z-stress.
- **Radial / breathing mode (2334 Hz, y-direction):** Produces significant hoop stress (x-stress), although the associated radial displacement (y-displacement) is moderate.

The following observations can be made for the HD point:

- **Flexural mode (40 Hz, x–y bending):** Less pronounced due to the downstream location. Negligible contribution to the stress response.
- **Longitudinal mode (199 Hz, z-direction):** Clearly present in the z-displacement but generates lower z-stress amplitudes than at the LD point.
- **Radial / breathing mode (2215 Hz, y-direction):** Frequency shift relative to the LD sector due to shrink-fit pre-stress. Both displacement and stress amplitudes are significantly smaller than at the LD point.

An experimental modal test campaign of the Run 3 beam dumps was performed by Guinchard [33] during Long Shutdown 2. The beam dump was instrumented with accelerometers and excited using an impact hammer. The experimentally identified natural frequencies were subsequently compared with the eigenfrequencies predicted by the finite element model. The results are presented in Table 6.2.

The good agreement in eigenfrequencies demonstrates that the global stiffness and mass representation of the finite element model are accurate.

Mode shape	Measurement [Hz]	Analysis [Hz]	Diff [%]
Flexural	37	40	8.11
Longitudinal	197	199	1.02
Breathing (LD-sector)	2360	2334	-1.10

Table 6.2: Comparison of predicted and measured natural frequencies of the Run 3 beam dump.

6.3 Dynamic Response Vibrational Magnitude

To validate the simulated vibration amplitude, Braken [22] simulated a real beam dump event with a stored beam energy of 402 MJ using the Run 3 model developed by the author. During this event, the axial velocity of the upstream flange was measured using a Laser Doppler Vibrometer. Figure 6.7 shows the measurement position on the upstream flange (left) and the measured longitudinal velocity compared with the simulated response (right).

The first two velocity peaks are reproduced accurately by the model. Subsequent oscillations show stronger damping in the measurements than in the simulation. The measured damping ratio is $\zeta = 8.5 \times 10^{-3}$, corresponding to approximately 5.2% amplitude decay per cycle. This damping is attributed to energy dissipation due to friction between the flexible graphite sheets and the vessel, which are difficult to capture in the simulation.

The agreement of the first oscillation confirms that the longitudinal vibration mode and its amplitude are accurately represented. Since the first oscillation has the largest magnitude, the simulated response can be used as a conservative basis for design, and the resulting vibration-induced stresses can be assessed with confidence.

It was not possible to validate the vibration amplitudes for the flexural (bending) and radial (breathing) modes due to the lack of instrumentation on the beam dumps for this purpose.

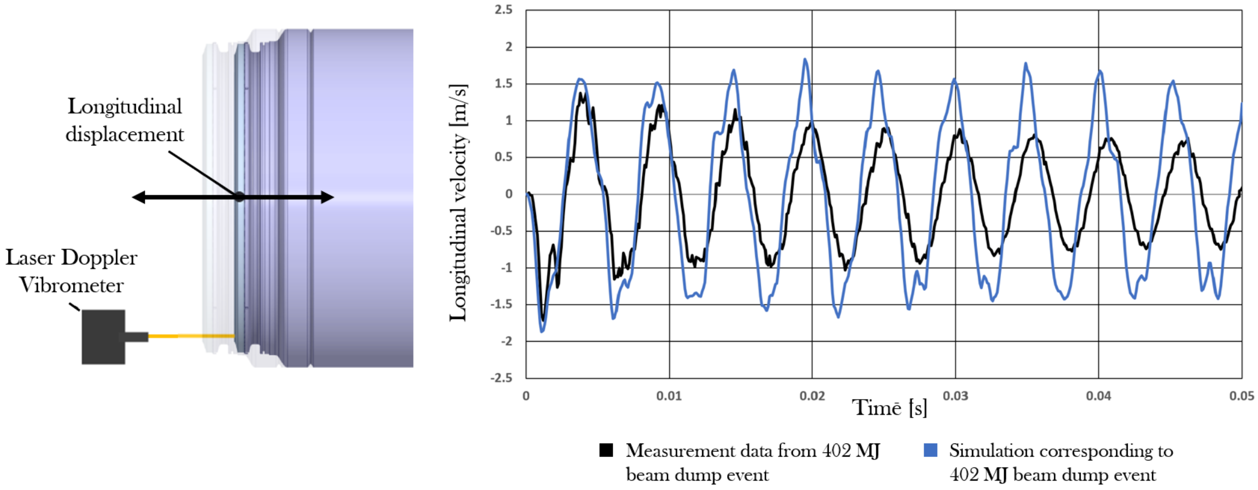

Figure 6.7: Left: Laser Doppler Vibrometer measurement position on the upstream flange. Negative velocity indicates vessel expansion and positive velocity vessel contraction. Right: Measured longitudinal velocity compared with simulated response for a 402 MJ beam dump event [22].

Chapter

7

Run 3 Beam Dump Performance Under High-Luminosity Conditions

Following the characterisation of the Run 3 beam dump response to Run 3 beam impact, the first objective is to assess the implications of operating the existing Run 3 beam dumps under HL conditions. As no damage has been observed on the Run 3 beam dumps to date, there is no technical reason to replace them, if their performance under HL conditions is acceptable. Reusing the existing dumps would significantly reduce the cost and effort associated with developing and manufacturing new dedicated HL beam dumps. This chapter outlines the operational differences between Run 3 and HL in terms of beam dump performance and evaluates the feasibility of installing the existing Run 3 beam dumps in HL.

The thermo-mechanical response of the beam dump is governed by the energy deposition from the particle shower during beam impact. Therefore, the comparison between Run 3 and HL starts with the change in beam parameters. As shown in Table 3.1, the number of protons per bunch increases by 22%, the number of bunches per beam increases from 2748 to 2760, and the energy per proton increases from 6.8 TeV to 7 TeV. As a result, the maximum stored beam energy increases by 32%, from 539 MJ to 710 MJ.

Table 7.1 shows the corresponding comparison of peak energy density and total energy between Run 3 and HL. For all components, the peak energy density increases by 28%. Consequently, the peak temperature in the vessel increases proportionally. Likewise for total energy, all components increase their energy deposition with 23%. This implies not only higher local temperatures but also a globally higher thermal field magnitude in HL compared to Run 3.

Table 7.1: Comparison of peak energy density and total energy deposition between full-intensity Run 3 and HL beam dumps from FLUKA simulations.

	Peak Energy Density [MW/cm^3]			Total Energy Deposition [kJ]		
	Run 3	HL	Increase [%]	Run 3	HL	Increase [%]
US Window	4.18	5.34	27.8	9.8	12.0	23.0
US Flange	0.22	0.28	27.8	26.9	33.0	23.0
Vessel	4.68	5.99	27.8	24 800.0	30 495.8	23.0
DS Flange	1.51	1.93	27.8	415.1	510.4	23.0
DS Window	4.11	5.26	27.8	269.1	330.9	23.0

Due to the complexity of the full post-processing described in Section 6.2, simplified response metrics are used to compare Run 3 and HL conditions, and to evaluate design variations in Chapter 8.

Figure 7.1a shows the relative longitudinal displacement (z-direction) between the upstream and downstream flanges. This quantity characterises the magnitude of the longitudinal vibration. From Run 3 to HL, the amplitude increases by 27%. Under HL conditions, the relative longitudinal expansion reaches approximately 7 mm at 140 Hz.

Figure 7.1b shows the von Mises stress history of the element experiencing the peak stress. This element is located at the transition between the LD sector and the downstream HD sector. The initial stress level originates from the shrink-fit pre-stress. Under HL loading, the stress oscillations have higher amplitude, reaching a peak of 370 MPa at approximately 22 ms. The temperature-dependent yield strength of 318LN is included to illustrate the reduction in material strength with temperature. The yield strength is 450 MPa at room temperature, 360 MPa at 100 °C, and 310 MPa at 200 °C. For the expected HL peak temperature of approximately 157 °C, the stress oscillations exceed the elastic limit of the material. This implies that the vessel would operate locally in the plastic regime at each full-intensity beam impact. Repeated plastic cycling could lead to accumulated damage, compromising the structural integrity of the vessel and potentially requiring replacement, resulting in significant LHC downtime and associated cost.

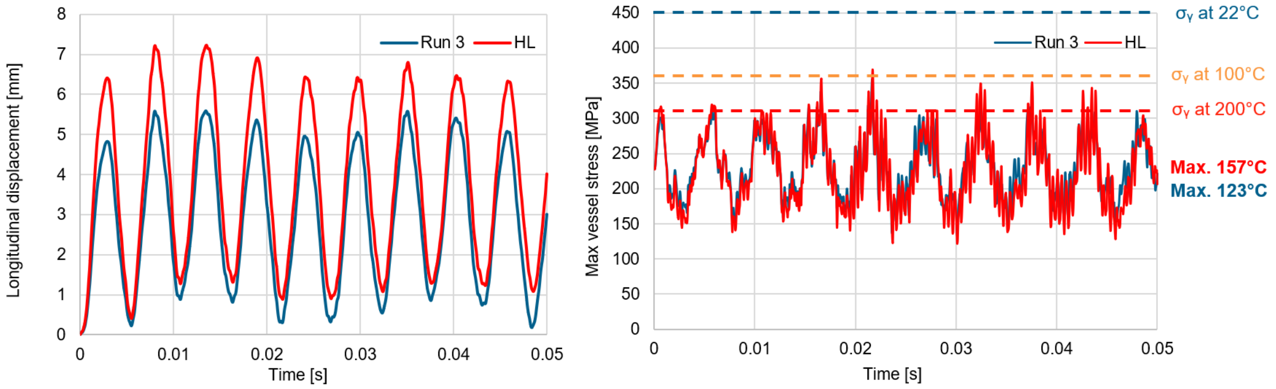

(a) Relative longitudinal displacement history measured between the upstream and downstream flange of the beam dump.

(b) Von Mises stress history of element experiencing peak stress level in the vessel. 318LN yield strengths at elevated temperatures from [1].

Figure 7.1: Comparison of longitudinal relative displacement and peak von Mises stress for Run 3 and HL conditions.

In addition, the Run 3 vessel segments are joined by partial penetration welds, as shown in Figure 7.2. This detail is not included in the global model, where the vessel is represented with a continuous wall thickness of 12 mm. Figure 7.2b shows a simplified 2D submodel of the detailed weld, based on the macroscopic weld cross-section in Figure 7.2a. The weld was analysed with a refined mesh near the root of the weld and subjected to axial stress. The results show elevated stresses and clear stress concentrations at the weld root during longitudinal (z-direction) vibration. Thus, if the vessel with a uniform 12 mm wall thickness would operate locally in the plastic regime under full-intensity beam impact, additional stress amplification should be expected in the weld regions due to geometric stress concentrations.

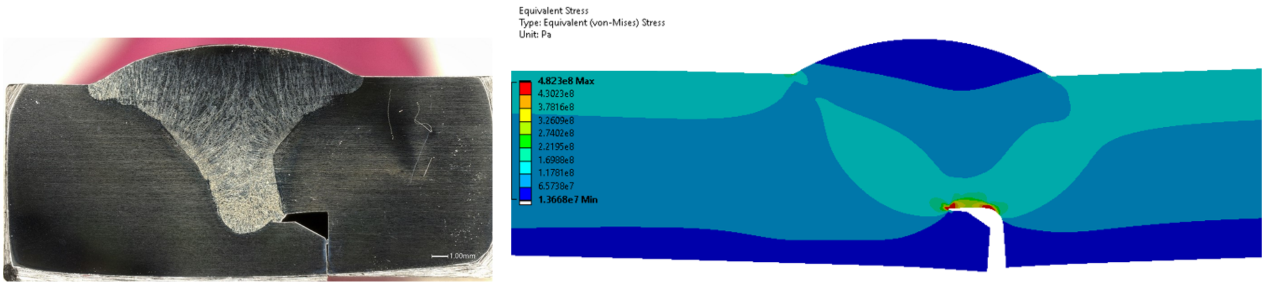

(a) Cross-sectional view of the partial penetration weld joint design.

(b) Simulated stress concentrations at the weld root due to the weld joint design.

Figure 7.2: Circumferential partial penetration weld joining the Run 3 vessel segments.

Based on these findings, the Run 3 beam dumps are not considered suitable for HL operation. The increased HL loading leads to cyclic stresses in the plastic regime at each full-intensity beam impact. Over repeated impacts during their service life, this may compromise structural integrity. Furthermore, partial penetration welds were deemed unsuitable due to inherent stress concentrations and the difficulty of reliably assessing their fatigue performance. These conclusions initiated dedicated design studies to ensure the development of beam dumps that can safely sustain HL operation, as described in the next chapter.

Chapter

8

Design Optimisation for the High-Luminosity Beam Dump

This chapter evaluates alternative beam dump configurations for HL aimed at reducing the dynamic stress levels caused by beam impact. The design space is constrained by the spatial limitations in the cavern as described in Section 3.6. Consequently, only minor modifications of the Run 3 design were considered. All simulations were performed with full intensity HL beam parameters.

8.1 Geometric Optimisation of the 318LN Vessel

The initial focus was on geometric modifications while retaining the 318LN vessel. This minimises manufacturing risk, as the materials, production and assembly remain similar to the current design. Two configurations were investigated and are described in the following sections.

8.1.1 Increase in Vessel Thickness

The first approach was to increase the vessel wall thickness from 12 mm to 18 mm. A larger cross-sectional area should reduce stresses by increasing stiffness and load-carrying capacity.

However, increasing the thickness by 50% also increases the interception of the particle shower. FLUKA simulations showed a 41% increase in total energy deposition, while the peak energy density remained nearly unchanged. The thermo-structural simulation with the updated energy deposition showed no improvement. The relative longitudinal displacement between the upstream and downstream flanges increased by 10%. The longitudinal frequency increased due to higher stiffness, and the maximum von Mises stress remained unchanged compared to the Run 3 design.

The increased load capacity from the thicker vessel wall is effectively offset by the corresponding increase in energy deposition. Consequently, this configuration does not provide significant stress reduction.

8.1.2 Reduction in Beam Dump Length

The second approach was to shorten the beam dump, assuming that a reduced quantity of vibrating mass would lower stress levels. Two downstream high-density blocks were removed, reducing the total length by approximately 1.4 m.

This configuration also showed no improvement. The relative longitudinal displacement remained essentially unchanged while the maximum von Mises stress increased by 14%, from 370 MPa to 424 MPa. The longitudinal vibration frequency increased due to the reduced length of the vessel. In addition, the downstream leakage, defined as the portion of the particle shower exiting the downstream window due to under-absorption, increased significantly after removal of the two isostatic graphite blocks. For the downstream window the peak energy density increased by 332%, and the total deposited energy increased by 273% to 1.24 MJ. A refined model, with local mesh refinement on the downstream window to accurately resolve the energy deposition, predicted peak temperatures of 440 °C and stress levels close to the yield strength of the titanium Grade 5 window material. In addition, activation of the shielding behind the dump would increase significantly.

This configuration does not provide viable stress mitigations in the vessel and introduces additional operational risks at the downstream end of the beam dump.

8.2 Vessel Material Change to Titanium Grade 5

The final option was to replace the vessel material with one offering more favourable mechanical and thermal properties. Higher-strength steels such as super duplex (2207) or martensitic grades (410, 420, 431) can reach yield strengths up to 1000 MPa at room temperature. However, their density, Young's modulus, and coefficient of thermal expansion are similar to 318LN, and all show significant strength reduction above 100 °C. Since the vessel operates at elevated temperatures, these steels offer limited improvement in overall performance.

Titanium Grade 5 (Ti-6Al-4V) was therefore evaluated. It is the most widely used titanium alloy, accounting for more than 50% of titanium alloys in service, which makes it well tested and commercially available [41]. Table 8.1 compares key properties of Titanium Grade 5 and SS 318LN. Titanium Grade 5 offers:

- 43% lower density, reducing interaction with the particle shower and total deposited energy.
- 30% lower thermal expansion coefficient, reducing thermally induced compressive stresses.
- 45% lower Young's modulus, reducing stress for a given vibration amplitude.
- 100% higher yield strength, significantly increasing the margin against plasticity.

Table 8.1: Comparison of SS 318LN [1] and titanium Grade 5 (Ti-6Al-4V) [41] at room temperature.

Property	SS 318LN	Ti-6Al-4V (Grade 5)	Difference
Density [kg/m^3]	7 800	4 430	-43%
Thermal Expansion Coefficient [K^{-1}]	1.24×10^{-5}	8.76×10^{-6}	-30%
Young's Modulus [GPa]	200	110	-45%
Yield Strength [MPa]	450	900	+100%

Titanium Grade 5 is also used in high-temperature applications, such as aerospace engines operating up to 350 °C. Table 8.2 shows the temperature-dependent yield strength based on ASME SB-318 minimum guaranteed values [13]. Even at 200 °C, the yield strength remains high at 590 MPa.

Table 8.2: Yield strength of titanium Grade 5 as a function of temperature from [47, 13].

Temperature [°C]	Yield Strength 0.2% [MPa]
100	702
200	590

Table 8.3 compares the energy deposition in the original 318LN vessel and the titanium Grade 5 vessel. The peak energy density in the vessel is reduced by 47%, and the total deposited energy decreases from 30.5 MJ to 17.1 MJ, corresponding to a reduction of 44%.

Table 8.3: Comparison of peak energy density and total energy deposition between 318LN and titanium Grade 5 with full-intensity HL beam parameters.

	Peak Energy Density [MW/cm^3]			Total Energy Deposition [kJ]		
	318LN	Titanium	Increase [%]	318LN	Titanium	Increase [%]
US Window	5.34	6.01	12.4	12.0	11.8	-1.7
US Flange	0.28	0.08	-72.2	33.0	11.1	-66.4
Vessel	5.99	3.18	-46.9	30 495.8	17 143.0	-43.8
DS Flange	1.93	1.06	-45.2	510.4	287.8	-43.6
DS Window	5.26	5.14	-2.2	330.9	341.7	3.3

The thermo-mechanical effect of replacing 318LN with titanium Grade 5 while maintaining the Run 3 geometry is shown in Figure 8.1a for the relative longitudinal displacement and in Figure 8.1b for the maximum vessel stress.

Figure 8.1a shows that titanium Grade 5 reduces the peak relative longitudinal displacement by 48%, demonstrating the combined benefit of lower density, lower thermal expansion, and lower stiffness. The longitudinal frequency decreases from 199 Hz to 144 Hz due to the reduced Young's modulus.

Figure 8.1b shows a 40% reduction in maximum von Mises stress, from 370 MPa (318LN) to 220 MPa (titanium Grade 5). The stress oscillation amplitudes are also lower, which improves fatigue resilliance. At approximately 150 °C, the yield strength of titanium Grade 5 is about 650 MPa, resulting in a safety factor of approximately 3 against yielding.

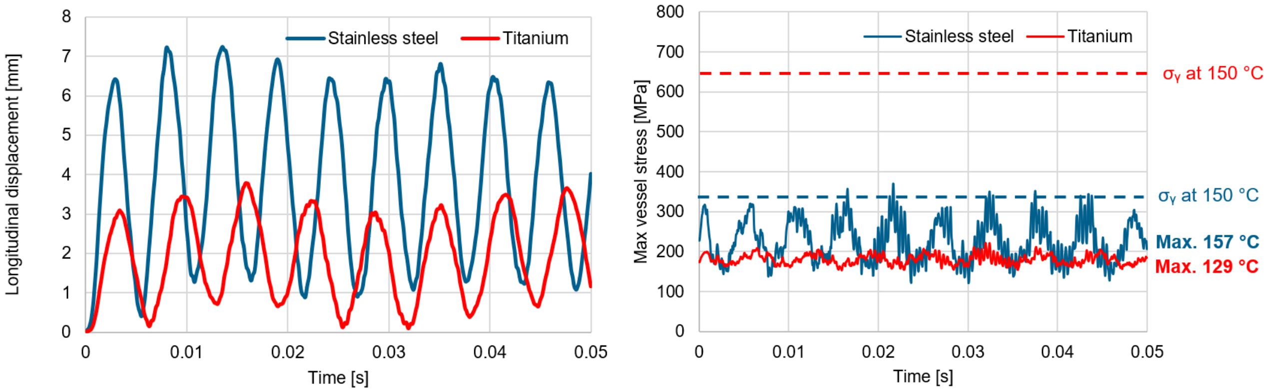

(a) Relative longitudinal displacement between upstream and downstream flanges between vessel of 318LN and Titanium Grade 5.

(b) Von Mises stress history of the element experiencing maximum stress. Yield strength from [1] (318LN) and [47, 13] (titanium Grade 5).

Figure 8.1: Comparison of longitudinal displacement and maximum stress under HL conditions for 318LN and titanium Grade 5 vessels.

Because titanium Grade 5 significantly reduces vibration amplitudes and stress levels and increase the safety margin against yielding, it was selected as the vessel material for the HL beam dumps.

Chapter

9

High-Luminosity Beam Dump Design and Performance Assessment

The dynamic response simulations formed a central input to the HL beam dump design, but they were only one part of the overall development process. In parallel, candidate core materials were subjected to an extensive qualification campaign. A dedicated experiment was installed at CERN to expose candidate materials to consecutive beam impacts representative of HL operation. In addition, cooling simulations were performed to evaluate heat extraction between successive beam dumps. Both activities strongly influenced the final design but are outside the scope of this work. Further details can be found in [40] and [19].

9.1 Design Upgrades for High-Luminosity

Figure 9.1 compares the Run 3 and HL beam dump designs. The main modifications are outlined below:

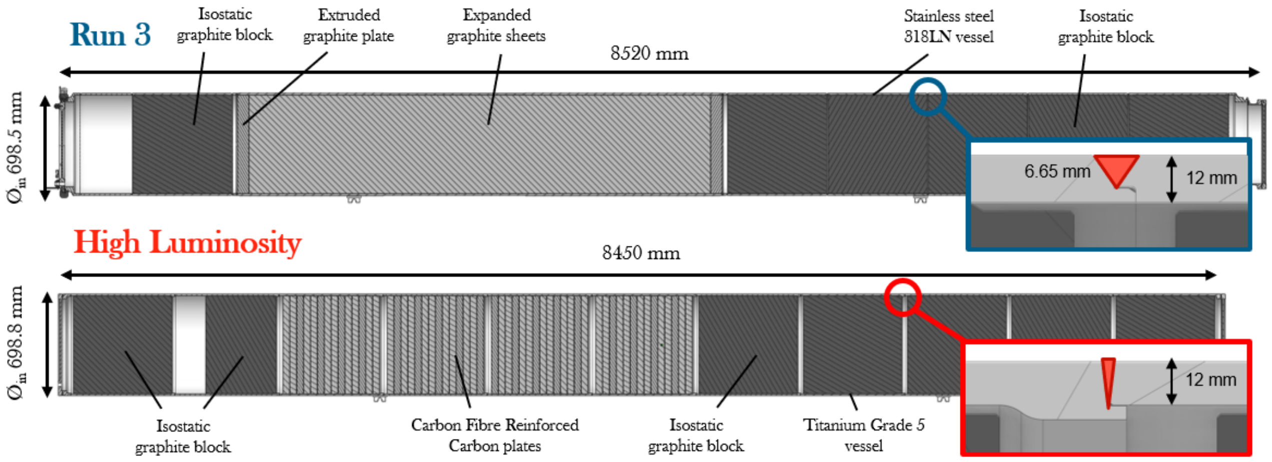

Figure 9.1: Schematic cross-sectional comparison between the Run 3 beam dump and the HL beam dump design.

- **Vessel material changed to titanium Grade 5:** The simulations showed that a titanium Grade 5 vessel is required to keep stresses within acceptable limits under HL loading. The wall thickness remains 12 mm. The inner diameter was slightly increased from 698.5 mm (Run 3, 318LN) to 698.8 mm to achieve the desired shrink-fit conditions. This adjustment balances sufficient contact pressure in operation with practical assembly constraints, taking into account the lower thermal expansion of titanium.
- **Full penetration welds:** The circumferential welds connecting the vessel segments were changed from 6.65 mm partial penetration welds to 12 mm full penetration welds to minimise stress concentrations. The new concept is based on electron beam welding. An integrated backing ring supports weld pool stability and ensures high weld quality. A total of 11 circumferential welds, one for each vessel segment with shrink fitted core, must be performed to assemble the full length of the beam dump. Fatigue performance of the new weld design is assessed in Section 10.2.
- **Revised window–flange joint:** The flanges are now integrated into the vessel segments, and the windows are attached using 72 M14 bolts. This concept is validated in Section 11.3.

- **Updated core configuration:** The 3.26 m LD sector of expanded graphite sheets is replaced by four sections of carbon fibre reinforced carbon (CFC) plates. The expanded graphite did not survive qualification testing, whereas CFC showed no visible damage. The use of CFC also enables shrink fitting of the LD-sector, since rigid plates are available in thicknesses up to 50 mm. This reduces radial vibration amplitudes and improves thermal contact between the core and vessel, enhancing heat extraction. The core layout was further optimised to limit the use of low-density CFC, which is approximately five times more expensive per unit length than isostatic graphite. CFC is therefore used only where it is necessary resulting in an extra 40 cm long isostatic graphite block being installed upstream of the LD sector.

9.2 Performance Assessment of the High-Luminosity Beam Dump Design

Due to the significant design modifications introduced for the HL beam dumps, the final configuration was re-evaluated to ensure that no unintended effects compromise performance. The global model was updated to the HL geometry, the titanium Grade 5 vessel, the CFC material in the LD sector, and the revised interference fit parameters. The configuration was assessed for full-intensity beam impact, dilution failure, and consecutive beam dumps.

9.2.1 Full Intensity Beam Dumps

Figure 9.2 shows the von Mises stress history of the vessel elements experiencing the maximum stress in the LD and HD sectors. The side view of the beam dump above the graph illustrates the shrink-fit stress state prior to beam impact.

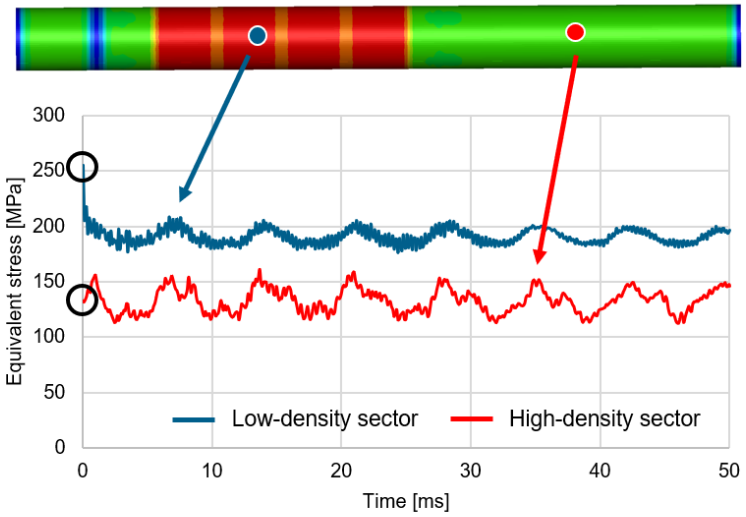

Figure 9.2: Transient von Mises history for elements in the LD and HD sector experiencing the maximum von Mises stress.

The LD sector introduces higher initial vessel stress due to increased shrink fitting interference. This is intentional since the LD sector experiences the highest temperature during beam impact, and a larger interference fit is required to prevent loss of contact during successive impacts with thermal build-up. The diameter interference fit is 1.8 mm in the LD sector and 1.2 mm in the HD-sector, corresponding to initial stresses in the vessel of approximately 250 MPa and 135 MPa, respectively.

Immediately after beam impact, the LD sector stress decreases due to circumferential thermal expansion, which partially relaxes the shrink-fit stress. This effect is less pronounced in the HD sector, where the temperature rise is smaller. Subsequently, the stress response is dominated by the longitudinal (z-direction) vibration mode at 158 Hz with stress oscillations of around 50 MPa in stress range. The flexural mode remains at 40 Hz, and the breathing mode at 2220 Hz; however, in the shrink-fit configuration of the HL beam dump, their contribution to the stress level is limited.

9.2.2 Dilution Failure

An important accidental scenario identified in Section 3.3 is beam dilution failure, in which one or more dilution kicker magnets malfunction and the beam is insufficiently diluted. The most severe case corresponds to two missing horizontal kickers.

Figure 9.3 compares the nominal beam sweep (left) with the accidental case (right). The absence of horizontal dilution creates a hotspot near the top of the sweep, where more proton bunches impact the same region. As shown in Figure 3.1, this increases the peak energy density deposited in the core materials in the LD sector by 68% and in the HD sector by 57%. While the impact on the core materials is critical, the effect on the global dynamic response must also be evaluated.

The global model was updated with the modified thermal load corresponding to two missing horizontal kickers. For the vessel, the total deposited energy decreased by 14%. This occurs because more protons are absorbed near the core centre, reducing the radial development of the particle shower toward the vessel. Although this is detrimental for the core, it is beneficial for the vessel. Consequently, the peak relative longitudinal expansion decreased by 11% compared to the nominal sweep, and the maximum vessel stress decreased by 2%.

For the windows, which are directly traversed by the beam, the dilution failure scenario is expected to produce more severe loading. This is assessed in Section 11.

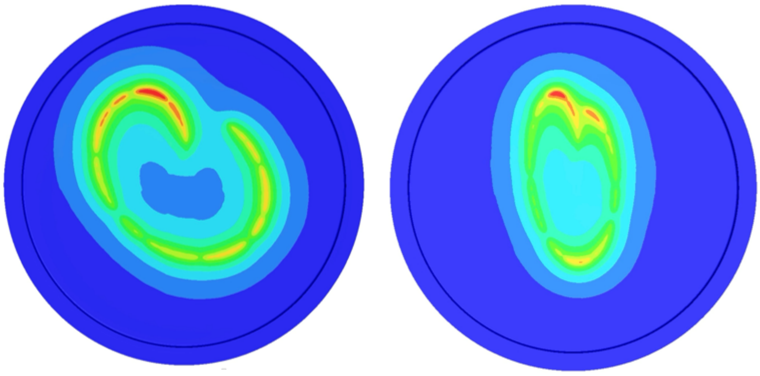

(a) Nominal sweep. **(b)** Two missing horizontal kickers.

Figure 9.3: Nominal and accidental beam dilution sweep patterns for HL operation with schematic contours of temperature.

9.2.3 Consecutive Beam Dumps

So far, the beam dump has been evaluated from an initial cold state at room temperature, representing normal conditions. In practice, however, the LHC can dump the beam every 2.5 hours. In this case, a second beam dump may occur before the cooling system has removed all residual heat from the first event.

To assess this scenario, two consecutive full-intensity beam dumps separated by 2.5 hours were simulated. An upgraded cooling system, developed by Notarianni [19], includes additional top ducts supplementing the original ducts beneath the beam dump to improve cooling performance. The thermal field remaining 2.5 hours after the first beam dump was extracted from Notarianni's results in Ansys Fluent and applied as the initial condition in the LS-Dyna model before the second impact.

A dynamic relaxation phase was performed prior to the second beam impact. This introduced the residual thermal field, allowed the associated deformations to develop, and established static equilibrium before applying energy deposition from the second beam impact. The resulting deformation and stress state were validated against implicit simulations in Ansys Mechanical, showing good agreement.

Figure 9.4 shows the thermal field applied to represent the state 2.5 hours after the first beam dump. The influence of the cooling ducts is visible in the non-uniform vessel temperature distribution. It can be observed that the vessel retains a non-negligible residual temperature following the first beam dump, with vessel peak temperatures approaching 170°C and extensive regions remaining above 100°C.

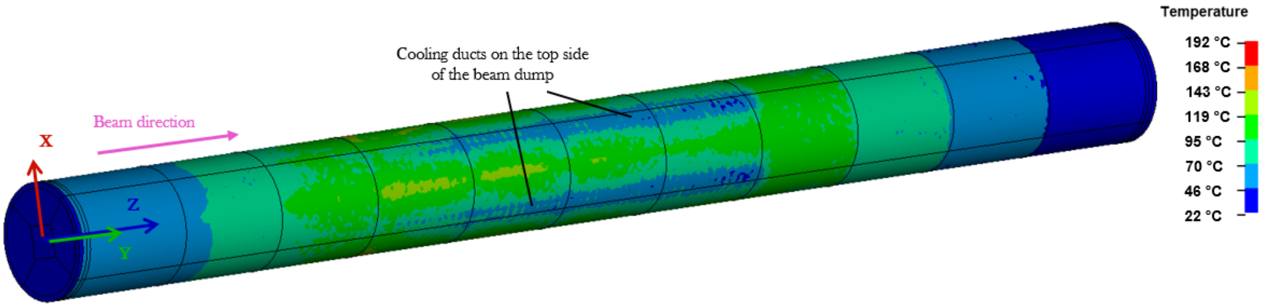

Figure 9.4: Top view of the thermal field representing conditions 2.5 hours after the first beam impact, immediately before the second impact. The influence of the cooling ducts is visible.

Table 9.1 compares the peak von Mises stresses for a single full-intensity beam dump and for the second dump occurring 2.5 hours later from an elevated temperature state. The maximum temperature increases by 57% to 217 °C in the LD sector and by 66% in the HD sector. Despite this temperature rise, the peak von Mises stress during the transient response decreases by 33% in the LD sector and 17% in the HD sector. This reduction is mainly due to thermal expansion of the vessel, which relaxes part of the hoop stress introduced by shrink fitting. Importantly, the stress reduction exceeds the corresponding decrease in yield strength, which drops by only 9.8% in the LD sector and 11.3% in the HD sector at the observed temperatures. The vessel stress therefore decreases at a greater rate than the material strength with increasing temperature, resulting in an increased safety margin with consecutive beam dumps.

Since multiple consecutive full-intensity dumps at 2.5-hour intervals are highly unlikely, the impact on the dynamic response of a third dump was not assessed. Two consecutive full-intensity dumps would already indicate abnormal machine operation, and further beam injection would typically be delayed until the cause is resolved, allowing sufficient cooling time.

Table 9.1: Maximum equivalent stress comparison in the vessel between the 1st and 2nd beam dump in the low-density and high-density sectors. Yield strengths linearly interpolated from Table 8.2.

Sector	Beam dump	Temperature [°C]	Max. von Mises [MPa]	Yield strength [MPa]	Safety factor [-]
Low density	1st	138	255	649	2.55
	2nd	217	170	585	3.44
	Difference	57%	-33%	-9.8%	
High density	1st	130	162	660	4.07
	2nd	216	135	585	4.33
	Difference	66%	-17%	-11.3%	

Chapter

10

Fatigue Assessment of the Vessel and Welds

So far, the dynamic response of the beam dump has been evaluated primarily in terms of peak stress values. However, fatigue life depends not only on the absolute stress level but also on the stress range. Elements subjected to large stress oscillations may therefore be more critical for fatigue than elements experiencing the highest peak stress but limited cyclic variation. Given the pronounced cyclic loading of the vessel, fatigue validation constitutes a key part of the structural assessment. It must be demonstrated that the beam dump can withstand HL operating conditions over a 10-year lifetime. This chapter presents the fatigue evaluation of the vessel base material and the electron beam welded joints circumferentially joining the vessel segments.

10.1 Vessel

nCode DesignLife is an Ansys extension for fatigue post-processing of transient finite-element results. It imports LS-Dyna result files and reads the full stress tensor history of each element. Further details are given in the DesignLife theory manual [11].

The workflow used for the vessel assessment is shown in Figure 10.1. The load spectrum is defined from LS-Dyna time-history results, and a duty-cycle feature enables combination of multiple load cases. The operational scenario corresponds to 10 years of operation and includes 1 000 full-intensity beam dumps and 3 400 half-intensity dumps (see Section 3.2). Accordingly, DesignLife was supplied with simulation results corresponding to both full-intensity and half-intensity conditions. Element temperatures are incorporated in the assessment, allowing the use of temperature-dependent S–N curves. During the assessment the stress tensor is first transformed into the principal coordinate system, and the maximum absolute principal stress is extracted as a scalar stress history. As fatigue cracks initiate at free surfaces, the evaluation is restricted to surface elements and out-of-plane stresses are neglected. Multiaxial indicators, including biaxiality and non-proportionality measures, are additionally computed and may be used for critical plane assessment. The resulting scalar stress history is decomposed using rainflow counting, after which fatigue damage is evaluated cycle-by-cycle based on a temperature-dependent S–N curve with Goodman mean-stress correction and Palmgren–Miner damage accumulation.

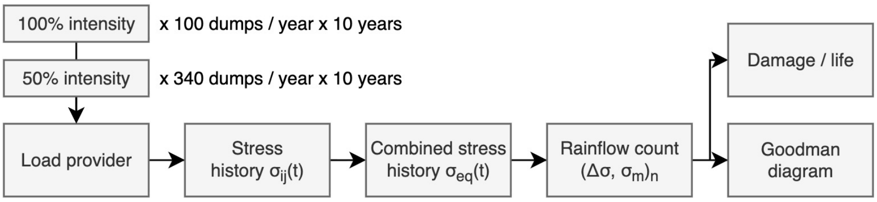

Figure 10.1: DesignLife fatigue analysis workflow.

In the fatigue assessment of the beam dumps, the stress state is biaxial, consisting of longitudinal vibrations in the axial direction and hoop stresses arising primarily from the shrink-fit, and secondarily from radial vibration modes. Consequently, all elements were evaluated using the critical plane method, which identifies the material plane experiencing the highest predicted damage. The normal stress is calculated and rainflow-counted on multiple planes, and the plane with maximum accumulated damage is defined as the critical one and used for the fatigue calculations. The stress on each plane from a 2D stress tensor on the surface is computed from:

$$\sigma_\phi = \frac{\sigma_{xx} + \sigma_{yy}}{2} + \frac{\sigma_{xx} - \sigma_{yy}}{2} \cos 2\phi + \sigma_{xy} \sin 2\phi \tag{10.1}$$

where ϕ is the angle between the two in-plane directions and is evaluated at $\phi = 0°, 10°, 20°, \ldots, 170°$.

For fatigue data, Tokaji [57] investigated the behaviour of titanium Grade 5 at room temperature and at 350°C and 450°C. The S–N curves shown in Figure 10.2 were implemented in DesignLife and allows the solver to interpolate between different temperatures. At room temperature, the infinite life limit is approximately 560 MPa, while at 350°C it decreases to approximately 240 MPa.

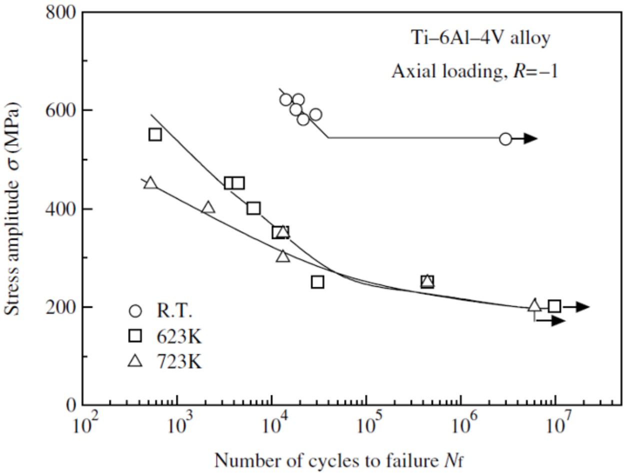

Figure 10.2: S-N diagram for titanium Grade 5 at different temperatures [57].

Figure 10.3 shows the spatial distribution of maximum stress amplitude in the vessel for the full load spectrum. The maximum amplitude is 75 MPa and occurs in the central region of the dump block, consistent with previous observations from simulations. In addition a significant mean stress is introduced due to the shrink fit of the core materials.

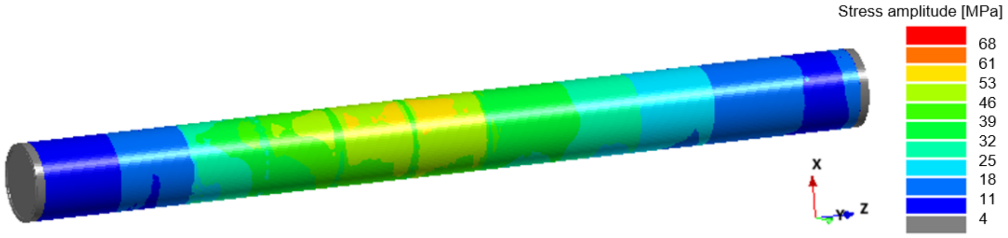

Figure 10.3: DesignLife fatigue assessment showing contour of absolute maximum stress amplitude.

The life predictions following the DesignLife results indicate infinite life for all elements, even with Goodman correction applied to account for mean stress corrections. To visualise the stress state, the maximum stress amplitude and corresponding mean stress were exported for each element in the model and plotted as a histogram and Goodman diagram in Figure 10.4 using a MatLab script.

The histogram in Figure 10.4a shows the number of elements from the finite element model at a given combination of mean stress and stress amplitude. Two clusters can be observed: one around a mean stress of 140 MPa which corresponds to the HD sector shrink fit and the other around a mean stress of 210 MPa which is due to the higher interference fit LD sector producing higher hoop stress. For both clusters, stress amplitudes remain below 75 MPa.

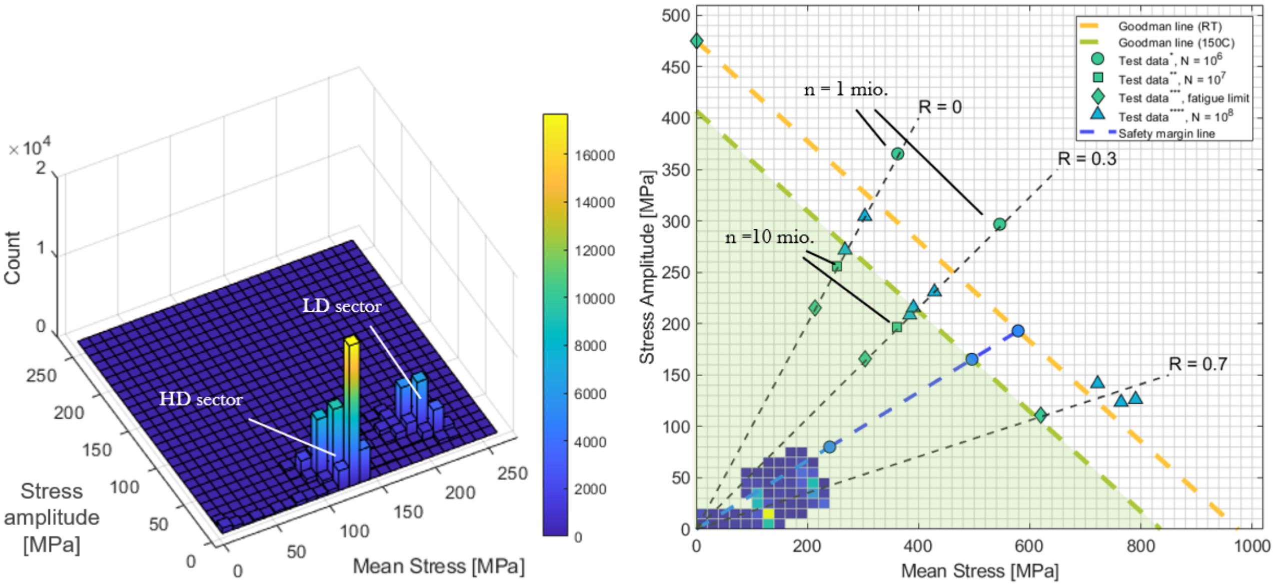

(a) Histogram of number of elements in the fatigue analysis that experience a given combination of mean stress and stress amplitude.

(b) Goodman diagram of the fatigue stress histogram, including Goodman lines and comparison to selected fatigue data from literature [30, 54]

Figure 10.4: Fatigue results from the DesignLife and MatLab post-processing.

Figure 10.4b shows the histogram from Figure 10.4a plotted in a Goodman diagram to illustrate the observed stress conditions in comparison to traditional defined Goodman limits and fatigue testing results found in literature. Goodman lines at room temperature and elevated temperature of 150°C as well as triangular test data were added to the diagram based on Furuya's study on *Gigacycle Fatigue Properties of Ti–6Al–4V Alloy under Tensile Mean Stress* [30]. Circular, square and diamond points are all fatigue test data from *Ensuring Fatigue Performance via Location-Specific Lifting in Aerospace Components Made of Titanium Alloys and Nickel-Base Superalloys* by Sangid [54]. The fatigue data are based on uniaxial specimens, where mean stress and stress amplitude act in the same direction. In the vessel, however, the mean stress arises from hoop stress, while the oscillating component is primarily axial, hence the loads are acting in perpendicular directions.

All combinations of mean stress and stress amplitude remain well below the Goodman limit, even at an elevated temperature of 150 °C. The total number of stress cycles over the beam dump lifetime depends on the assumed structural damping. Assuming a conservative damping ratio of 1% (representing 1/5 of the damping observed between successive oscillation cycles in the Run 3 beam dumps), approximately 5×10^5 cycles are accumulated over 10 years of operation. This estimate is based on 100 full intensity beam dumps and 340 half intensity dumps per year.

The test data (square datapoints) indicate endurance beyond 10^7 cycles at a mean stress of 360 MPa combined with a stress amplitude of 200 MPa acting in the same direction. In comparison, the worst element in the vessel has a mean stress of approximately 240 MPa and an amplitude of 80 MPa, acting in perpendicular directions, with a required lifetime of only 5×10^5 cycles.

Based on this comparison, the vessel is expected to exhibit infinite fatigue life with a substantial safety margin in both stress level and cycle count.

10.2 Electron Beam Welds

The detailed geometry of the circumferential electron beam welds that join the vessel segments is not implemented in the global finite element model. The vessel is modelled with a constant thickness along its full length. Consequently, local geometric effects of the weld are not captured in the global analysis and must be assessed separately. This section evaluates the fatigue performance of the electron beam welds under full-intensity HL operating conditions.

10.2.1 Weld Geometry and Loading Conditions

Figure 10.5 shows a cross-sectional view of the HL beam dump design. The orange lines indicate the circumferential electron beam weld locations. The plot below presents the relative longitudinal displacement between the upstream and downstream flange (red) and the resulting axial stress at the weld element experiencing the maximum axial stress during full-intensity HL beam impact.

The most critically loaded weld is located in the central region of the beam dump, between the LD sector and the downstream HD sector. The axial stress oscillates around a mean value of 0 MPa with a maximum amplitude of 80 MPa. These values represent the far-field loading acting on the weld. In the immediate vicinity of the weld, the local stress state is more complex due to geometric stress concentrations and non-uniform stress distributions.

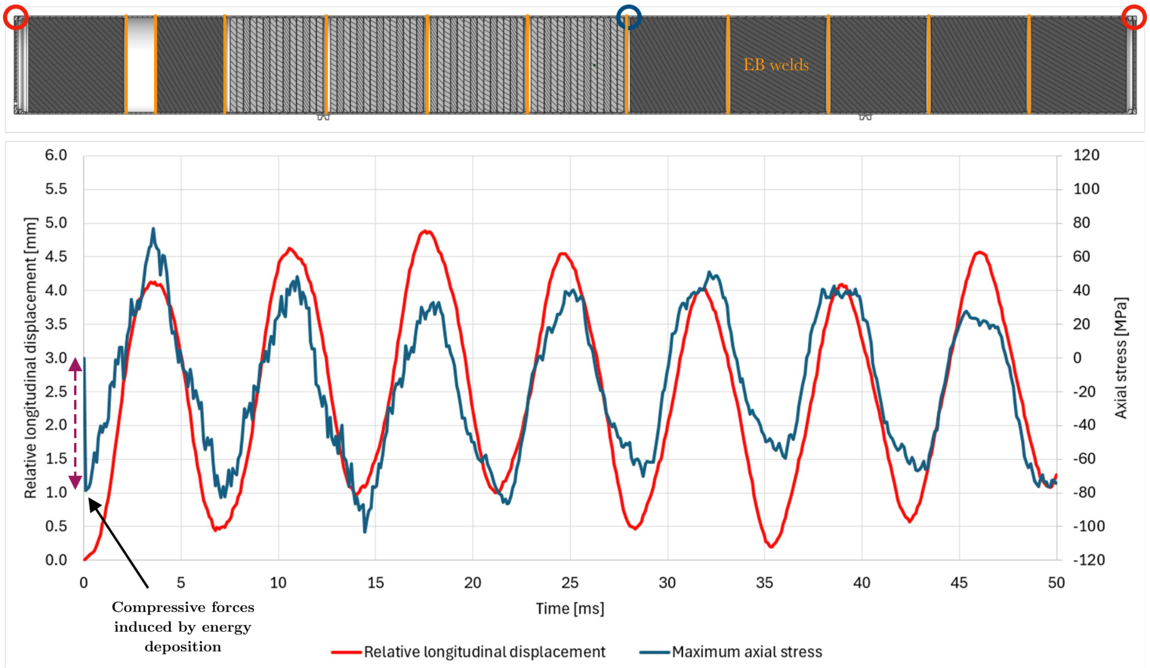

Figure 10.5: Top: Cross-sectional view of the HL beam dump. Orange lines indicate the locations of the circumferential electron beam welds. Bottom: The red curve shows the relative longitudinal displacement between the two flanges (positions marked with red circles in the cross-sectional view). The blue curve shows the axial stress history of the element experiencing the highest axial stress, located at the transition between the LD sector and the downstream HD sector[39].

Figure 10.6 shows a schematic cross-section of the weld joint. The weld fully penetrates the 12 mm vessel thickness. A backing ring is integrated and serves as the weld interface during the electron beam welding process. The implementation of the backing ring is motivated by the following considerations [39]:

- **Inaccessible weld root after assembly:**
  The carbon-based core is shrink-fitted before welding. After welding, the inner weld root is inaccessible for inspection or post-weld treatment. The backing ring therefore remains a part of the joint.
- **Sensitive weld penetration:**
  Without a backing ring, excessive penetration could melt through the wall and eject molten titanium onto the internal carbon materials. Penetration depth cannot be controlled precisely due to beam current fluctuations and dimensional tolerances. Full melt-through may also destabilise the weld pool

and create geometric imperfections such as undercuts or irregular solidification, reducing fatigue strength. Insufficient penetration would result in a lack-of-fusion defect at the weld root, effectively forming a crack perpendicular to the axial load. The backing ring stabilises the weld pool and tolerates penetration variations of 1–2 mm while ensuring full fusion across the 12 mm wall thickness of the vessel.

- **Integrated alignment function:**
  The backing ring also provides a male-female interface between the joining segments, enabling self-alignment during assembly. This ensures proper surface contact and accurate lateral positioning, which is beneficial for weld quality.

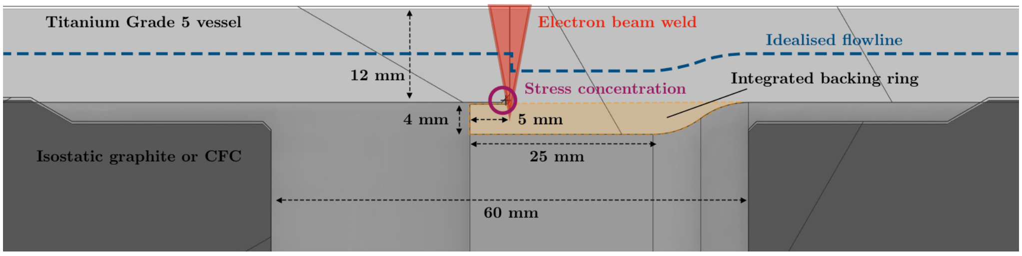

Figure 10.6: Schematic cross-sectional view of the circumferential full-penetration weld in the titanium Grade 5 vessel, including the integrated backing ring [39].

Although the backing ring is required, it introduces geometric asymmetry in the weld region. As shown in Figure 10.6, the stress distribution on the left side of the weld can be approximated as uniform through the 12 mm vessel thickness. On the right side, the backing ring increases the effective wall thickness, reducing the average axial stress across that section. At the interface between the two blanks, this thickness discontinuity forces the load path to transition from the thicker backing ring region to the nominal wall. A local stress concentration is therefore expected on the left side near the weld root.

A previous study by the author [39] assessed the weld design using fracture mechanics, evaluating the stress intensity factor at the notch and estimating fatigue life through Paris law crack growth calculations. Due to the difficulty of accurately modelling the complex geometry, heterogeneous weld material, and residual stresses, the results were associated with significant uncertainty. Therefore, a dedicated fatigue test campaign was conducted to validate the weld design.

10.2.2 Fatigue Life Assessment

For the weld fatigue assessment, damping was included to obtain a realistic stress range and number of cycles over the lifetime. Figure 10.7 shows the axial stress history of the most critical weld during the first 50 ms full intensity beam dump, identical to Figure 10.5. A damped oscillation function was fitted to the simulation data, ensuring accurate representation of the first and most damaging cycle (160 MPa stress range). The fitted response was extrapolated until the stress amplitude dropped below 10 MPa, such that all relevant damaging cycles were included. Damped oscillations were generated for three assumed damping ratios, $\zeta = 8.5 \times 10^{-3}$ (measured in Run 3), $\zeta = 4.25 \times 10^{-3}$ (conservative estimate for HL), and $\zeta = 1.6 \times 10^{-3}$ (sensitivity study). The same exercise was performed on half-intensity beam dumps to study their combined effect on the fatigue life of the welds. A detailed description of the computation of this data as well as the Matlab script used is provided in Annex A.9.

For each damping ratio, rainflow counting was applied to determine the number of cycles n_i at stress ranges $\Delta\sigma_i$. The equivalent stress range was then calculated using the slope of the S-N curve obtained from the dedicated fatigue campaign ($b_1 = 2.85$). Its derivation can be found in Appendix A.10. With $N_{\text{tot}} = \sum n_i$ denoting the total number of counted cycles, the equivalent stress range is given by:

$$\Delta\sigma_{\text{eq}} = \left(\frac{\sum n_i (\Delta\sigma_i)^{b_1}}{N_{\text{tot}}} \right)^{\frac{1}{b_1}} \tag{10.2}$$

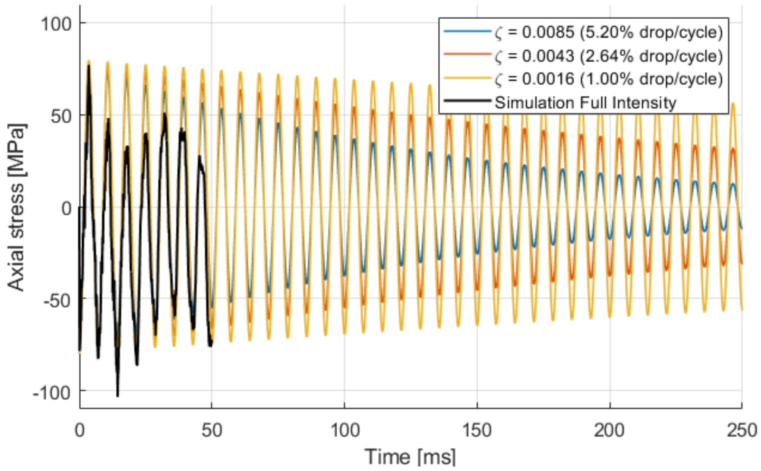

Figure 10.7: Extrapolation of axial stress response assuming different damping ratios.

.

Figure 10.8 presents the combined equivalent number of cycles resulting from 100 full-intensity and 340 half-intensity beam dumps per year over a 10-year period, evaluated for different assumed damping ratios. For comparison, all cases are normalised to an equivalent stress range of 64.5 MPa.

The results indicate that full-intensity beam dumps are approximately nine times more damaging than half-intensity dumps, despite occurring 3.4 times less frequently. A pronounced sensitivity to the assumed damping ratio is also observed. Reducing the damping from 5.2% to 1% increases the equivalent number of cycles by a factor of 5.3, from 97 000 to 517 300. This equivalent stress range and corresponding lifetime cycle count define the governing fatigue requirement for the welded joints. It must therefore be demonstrated that the welds can withstand this loading with an adequate safety margin.

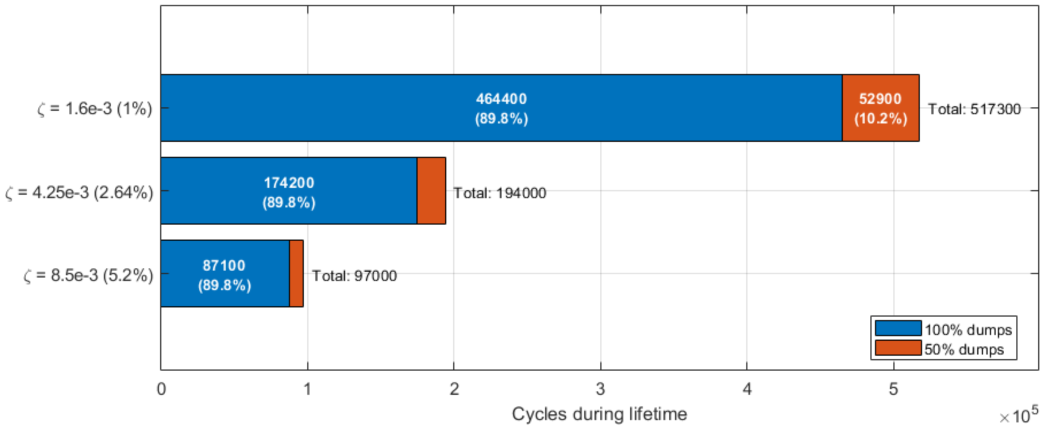

Figure 10.8: Combined equivalent number of cycles from 10 years of operation with 100 full-intensity and 340 half-intensity beam dumps per year for different assumed damping ratios.

10.2.3 Fatigue Test Campaign and Results

The objective of the fatigue test campaign was to characterise the performance of an electron beam weld representative of the final vessel design under realistic, yet conservative, service conditions. The specimens were manufactured from rolled titanium Grade 5 plates and incorporated the integrated backing ring and the weld-root notch present in the geometry. Several load levels representative of operational stresses were applied in order to establish a S–N curve. All tests were conducted at an elevated temperature of 150 °C, corresponding to the maximum expected operating condition.

A total of 20 specimens were tested using a resonance pulser with a maximum load capacity of 150 kN. The test frequency was set to 80 Hz. Specimens were loaded until failure, defined as a 2 Hz shift in resonance frequency, which in most cases corresponded to a crack propagating through more than half of the wall thickness. The specimen geometry included the integrated backing ring and a large gauge section to prevent buckling under compressive loading. Flatness between gripping points was carefully controlled to avoid unintended bending stresses. The welds were produced in accordance with the qualified Welding Procedure Specification (WPS) and Welding Procedure Qualification Record (WPQR) to be used for the welding of the beam dumps. The detailed specimen drawing is provided in Appendix A.8.

Two crack initiation mechanisms were observed. The dominant failure mode, occurring in 19 of 20 specimens, initiated at the weld-root notch and propagated through the wall thickness. Figure 10.9 shows a cross-sectional view of the specimen with the crack propagation through the thickness of the specimen. In two specimens, crack initiation occurred at the weld crown and propagated perpendicular to the loading direction. The predominance of root-initiated cracks confirms that the weld-root geometry introduces a stress concentration, making it the governing fatigue-critical feature of the design.

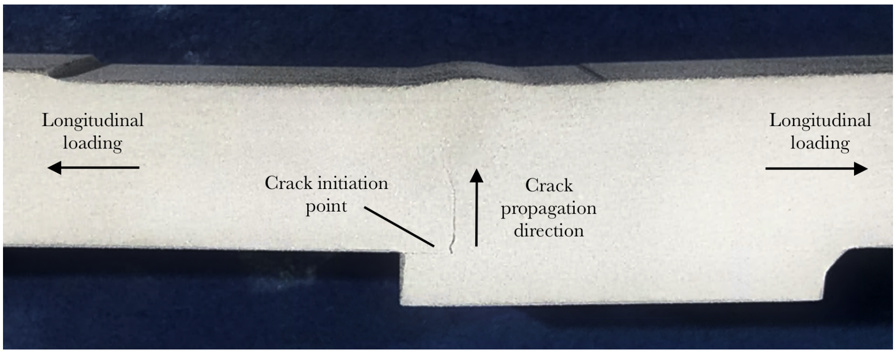

Figure 10.9: Side view of fatigue tested titanium Grade 5 electron beam welded specimen showing crack initiation from stress concentration at the weld root and propagation through the thickness perpendicular to the longitudinal loading.

The fatigue data were post-processed in accordance with ISO 12107:2012 [10]. Figure 10.10 presents the results in a double-logarithmic S–N diagram showing stress amplitude as a function of number of cycles. The red crosses indicate specimen failure points. The coloured lines represent S–N curves derived for different survival probabilities and confidence levels. The diamonds denote the operational conditions corresponding to 10 years of combined full- and half-intensity beam dumps, as previously shown in Figure 10.8.

ISO 12107:2012 assumes fatigue life to follow a log-normal distribution. Verification of this assumption is described in Appendix A.10. The experimental results were found to be reasonably log-normally distributed, allowing application of the standard methodology. Given the limited number of specimens and the inherent scatter in fatigue data, statistical treatment including confidence intervals and a specified survival probability was required to obtain a conservative design curve. The lower tolerance limit of the logarithmic fatigue life at a given stress level is defined as [10]:

$$\hat{Y}_{TL} = \hat{Y} - k_{(P,1-\alpha,v)}\,\hat{\sigma}\left[1 + \frac{1}{n} + \frac{(X-\bar{X})^2}{\sum_{i=1}^{n}(X_i-\bar{X})^2}\right]^{1/2} \tag{10.3}$$

Here, $Y = \log_{10}(N)$ denotes logarithmic fatigue life and X the stress level. The fitted regression line in log–log space is given by $\hat{Y}_i = b_0 + b_1 X_i$. $\hat{\sigma}$ is the estimated standard deviation of the residuals, i.e. the standard deviation of the differences between the experimentally observed logarithmic fatigue lives and the values predicted by the fitted regression line. n is the number of specimens, and $\bar{X}$ the mean stress level. The term $k_{(P,1-\alpha,v)}$ is the tolerance coefficient defined in ISO 12107:2012, which depends on the required survival probability P, confidence level $1-\alpha$, and degrees of freedom $v = n-2$. An overview of the $k_{(P,1-\alpha,v)}$ tolerance

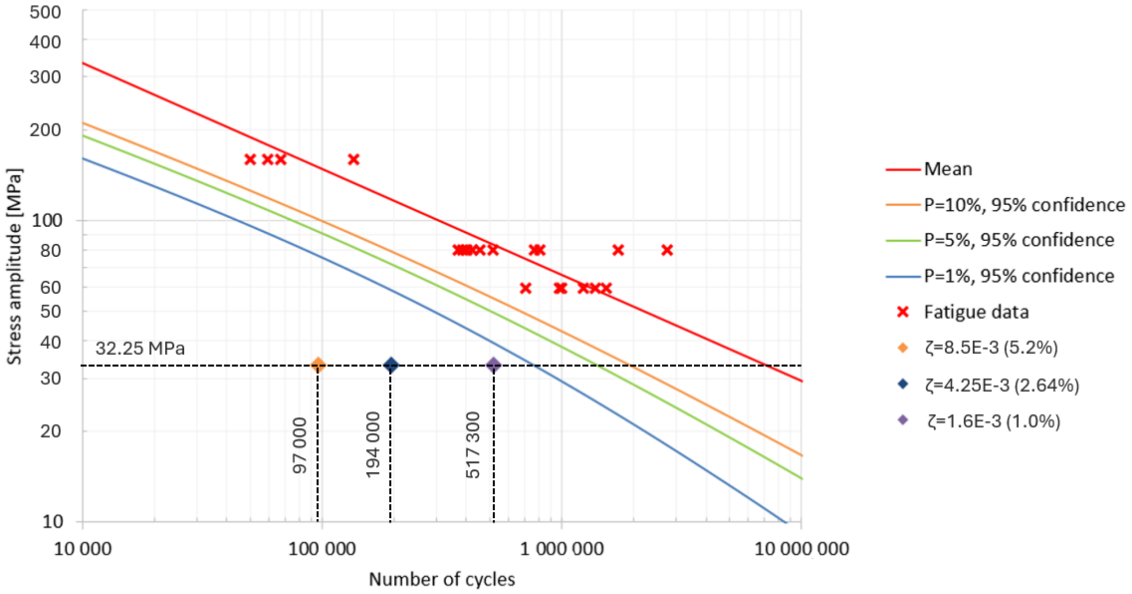

Figure 10.10: S-N curves computed from the raw fatigue data for probability-of-failure levels of 10%, 5%, and 1% at a 95% confidence interval. The HL operational conditions represent combined full- and half-intensity beam dumps with damping ratios of 5.2%, 2.64%, and 1% from Figure 10.8.

factors is provided in Appendix A.10. The square-root term accounts for the statistical uncertainty of the predicted fatigue life at the specified stress level, including both the scatter of the data and the uncertainty in the estimated regression parameters.

The resulting $\hat{Y}_{TL}$ represents the lower statistical tolerance limit of fatigue life. When transformed back to the original scale, it provides the fatigue life corresponding to the specified survival probability and confidence level required for design verification.

Table 10.1: Fatigue life of the electron beam welds given as safety factor in number of cycles based on the considered damping ratio and probability level of failure.

Probability level	Confidence level	$\zeta = 8.5 \times 10^{-3}$ (5.2% drop/cycle) Safety factor	$\zeta = 4.25 \times 10^{-3}$ (2.64% drop/cycle) Safety factor	$\zeta = 1.6 \times 10^{-3}$ (1% drop/cycle) Safety factor
1%	95%	8.3	4.1	1.6
5%		15.4	7.7	2.9
10%		21.1	10.6	4.0

Table 10.1 summarises the fatigue safety factors in terms of allowable number of cycles, evaluated for an equivalent stress amplitude of 32.25 MPa for different damping ratios and probability levels. The safety factor corresponds to the horizontal distance in Figure 10.10 between the operational loading point and the corresponding S-N curve. All values are calculated with a confidence level of 95%.

Even under the most conservative assumptions, 1% probability of failure and a damping ratio of 1%, a safety factor of 1.6 is obtained. For a more realistic, yet still conservative, damping ratio of 2.64%, the safety factor increases to 4.1. If a probability of failure of 5% is assumed instead, the safety factor increases further to 7.7. Considering the conservative assumptions adopted for the load spectrum in terms of distribution of dumped intensities, the reduced damping ratio relative to the measured Run 3 value, and the strict statistical post-processing using a 1% probability of failure, a safety factor of 4.1 is considered sufficient. The fatigue performance of the electron beam welded joint is therefore validated for operation under HL conditions.

Chapter

11

Window Response and Bolted Connection Assessment

The beam dump is enclosed at the upstream and downstream ends by 10 mm-thick titanium Grade 5 windows, which maintain an inert argon atmosphere and thereby prevent oxidation of the carbon-based core materials at elevated temperatures during beam impact. The windows are mounted to an integrated vessel flange using 72 M14 bolts. This chapter evaluates the window response to beam impact in Section 11.1 and the integrity of the bolted window–flange connection in Section 11.3.

When the beam impacts the dump, the first material encountered is the 10 mm-thick upstream window. At this location the particle shower has not yet developed, and FLUKA simulations [32] indicate an energy deposition of less than 0.01 MJ. In contrast, the downstream window is exposed to the fully developed particle shower and receives approximately 0.7 MJ, about 70 times more than the upstream window. The window response can be characterised on three distinct time scales:

- **Dynamic response during beam impact (0-100 μs):** The moving beam sweep as function of time generates rapidly varying energy deposition, producing transient stress waves in the window.
- **Post-impact thermal field ($\gtrsim$ 300 ms):** After the beam impact, the resulting thermal field induces mainly compressive stresses due to local temperature rise at the sweep location.
- **Vibration-induced response (3-500 ms:** Vibrations of the vessel flex the window due to its inertia and low geometrical stiffness, producing cyclic tensile and compressive stresses near the window centre during each oscillation.

11.1 Response to Dynamic Beam Sweep

To evaluate the stress waves generated in the windows during beam impact, a dedicated FE model was developed that is decoupled from the global model. This approach is necessary because accurate representation of the energy deposition requires a very fine mesh resolution. Incorporating such refinement in the global model would significantly increase the number of elements and reduce the time step, making the simulation computationally impractical. Furthermore, the stress-wave response occurs within the first 100 µs, before the global vibration of the vessel develops, which allows the phenomena to be analysed independently.

Figure 11.1 shows the FE model used to simulate the dynamic beam sweep response. The model consists only of the window which is unconstrained. Geometric partitions around the beam sweep were introduced to enable a very fine mesh in the sweep region, where accurate interpolation of the FLUKA energy deposition is required. All bodies share topology, and a gradual mesh growth rate ensures a smooth transition from the refined sweep region to the surrounding coarser mesh. The mesh contains approximately 870k solid linear hexahedral elements, with seven elements through the thickness to accurately capture through-thickness energy deposition gradients.

In contrast to the global model, the energy deposition is applied here as a time-dependent load within the beam impact duration. Each element therefore receives not only a spatially interpolated energy value from the FLUKA data, but also an associated load curve describing when this energy is deposited during the 86 µs beam impact. As illustrated in the figure, this results in individual load curves for each element in the model.

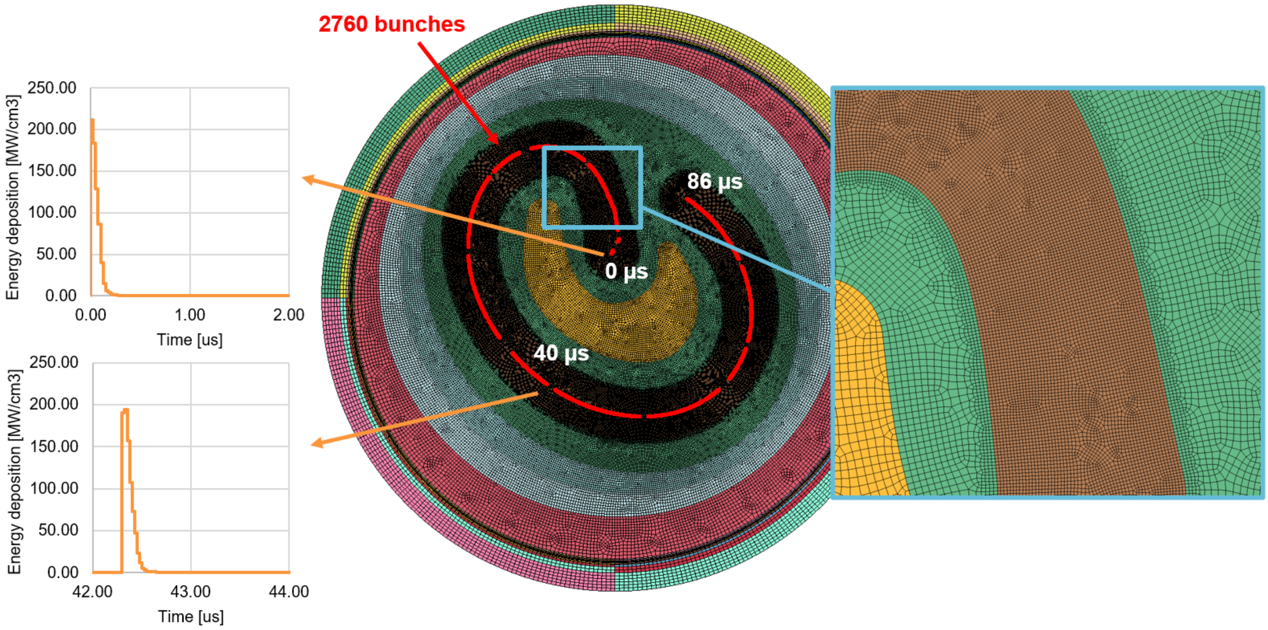

Figure 11.1: Finite element model used to assess dynamic sweep on the windows using spatial and temporal energy deposition and a high-density mesh in the beam sweep region.

Consequently, both spatial and temporal interpolation are required to reconstruct the full beam sweep. A Python script was developed to perform this interpolation. The script imports the developed mesh from the input file and uses the impact coordinates of all bunches as function of time to perform the interpolation. For elements affected by multiple bunches the energy deposition is accumulated. The script is provided in Appendix A.11 together with a visualizing script to validate the energy interpolation.

The simulation was run for an end time of 150 µs, corresponding to approximately twice the duration of the beam impact. The model was fully coupled thermo-mechanically, with the thermal solver constrained to a maximum temperature change of 1 °C between successive iterations to ensure accurate coupling. Separate simulations were performed for the upstream and downstream windows under both the nominal sweep pattern and the accidental case with two missing horizontal dilution kickers (2MKBH).

Figure 11.2 shows the temperature field (top) and von Mises stress field (bottom) at 23 µs (left), 60 µs (centre), and 89 µs (right) for the downstream window under nominal sweep conditions at full HL intensity. The temperature field develops progressively as successive bunches impact the window. The resulting thermal shock generates stress waves that propagate within the plane of the window, approximately at 45° relative to the beam sweep direction. The maximum temperature reaches 198 °C, occurring in the upper-left region of the sweep where the bunch spacing is smallest. The corresponding maximum von Mises stress across all time steps is 205 MPa. The stress state within the sweep region is dominated by in-plane compressive stresses, as confirmed by evaluation of the principal stress components. Although stress waves propagating away from the sweep region are visible, their magnitude remains small and they are likewise compressive. Consequently, the dynamic sweep loading does not produce stress states that are critical from a fatigue perspective.

Table 11.1 summarises the results for both upstream and downstream windows for the nominal sweep and the 2MKBH dilution failure scenario. The following observations can be made:

- **Upstream windows:** The upstream window experiences very limited energy deposition because the particle shower is undeveloped at this location. The window is therefore nearly transparent to the beam. As a result, temperatures remain below 80°C and stresses remain below 50 MPa for both sweep patterns. This leads to large safety factors to yield strength, approximately 27 for the nominal sweep and 19 for

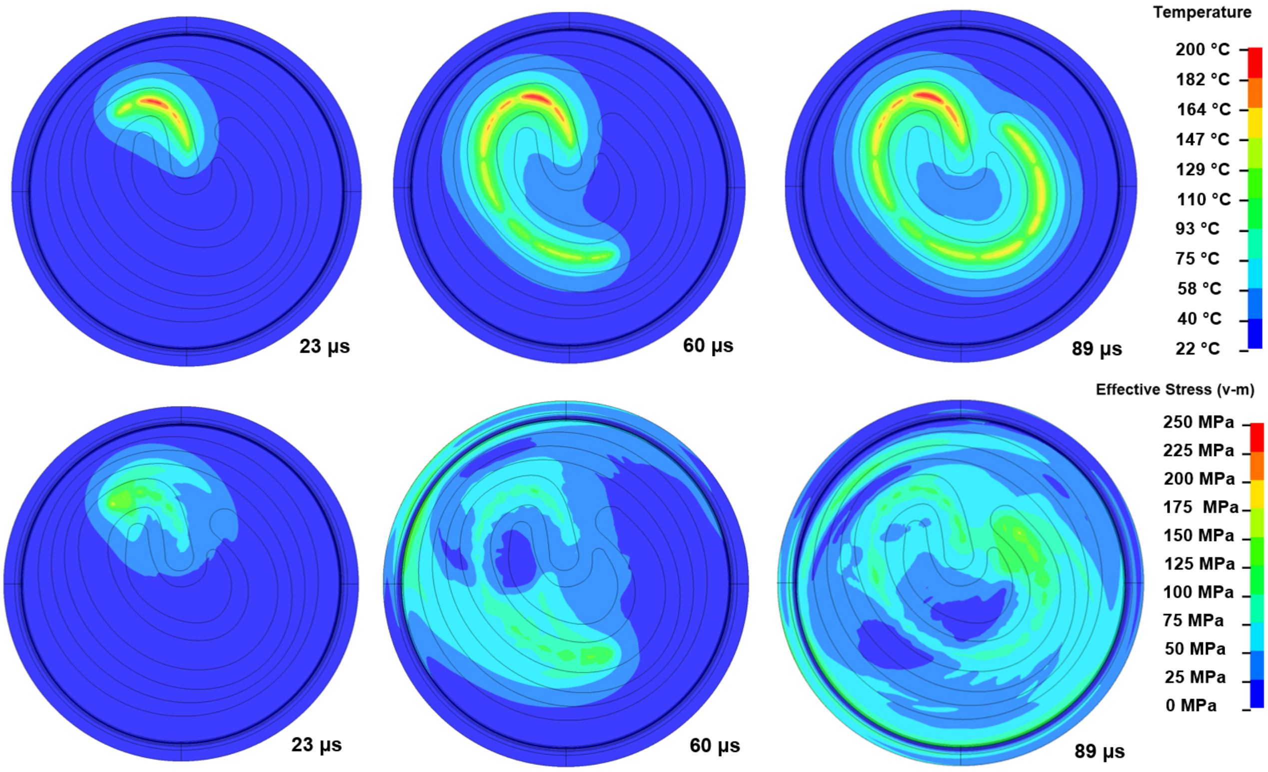

Figure 11.2: Simualted beam impact on the downstream window for nominal sweep and full-intensity HL beam parameters. Temperature (top) and von Mises stress (bottom) at 23 µs (left), 60 µs (centre), and 89 µs (right) during the beam sweep. Stress waves propagate in the plane of the window due to the rapid thermal loading.

the 2MKBH diluation failure case.

- **Dilution failure (2MKBH):** The 2MKBH scenario is more critical than the nominal sweep because the bunches have less horizontal spacing, producing higher local energy deposition. This results in higher peak temperatures and magnitude of compressive stresses. The downstream window reaches a peak temperature of about 276°C and a maximum stress of 241 MPa, corresponding to a safety factor to yield of approximately 2.7. Although this is the most severe case, its probability of occurrence is very low. It should therefore be treated as a single-impact verification rather than a fatigue case.

Table 11.1: Maximum temperature, von Mises stress, linear interpolated yield strength from Table 8.2, and safety factor to yield for different window positions and beam sweeps.

Position and beam sweep	Temperature [°C]	Max Stress [MPa]	Yield Strength [MPa]	Safety Factor
Upstream nominal	54	34	866	26.7
Downstream nominal	198	205	707	3.5
Upstream 2MKBH	76	45	838	18.6
Downstream 2MKBH	276	241	648	2.7

In summary, the stress waves generated by the dynamic sweep pattern do not pose a structural risk to the windows.

11.2 Response to Induced Vessel Vibrations

The stress waves generated in the window during beam impact are expected to decay rapidly and converge to a quasi-static state dominated by compressive stresses caused by the thermal field from the beam impact. These compressive stresses will gradually dissipate as the window cools down. Owing to the large surface area, small wall thickness, and direct exposure to ambient air on the outer surface of the beam dump, the thermal field is expected to decay relatively quickly. Nevertheless, the cooling process occurs over a time scale of seconds to minutes and can therefore be considered constant over the 50 ms simulation duration used to evaluate the vessel dynamics. Consequently, the window is subjected to a nearly constant compressive stress state originating from the thermal field while simultaneously experiencing the cyclic stresses induced by the vibration of the vessel. The resulting stress state is therefore the superposition of the thermal stresses and the vibration-induced stresses.

To evaluate this combined effect, the global model was updated with a refined mesh in the downstream window. The objective of the refinement was to capture the total energy deposition accurately, rather than resolving the local peak energy density, which had already been assessed in the dedicated dynamic sweep analysis in the previous section. Verification of the energy interpolation showed that the total deposited energy in the downstream window was reproduced within a deviation of 4%. The updated model was then used to assess the superimposed effect of energy deposition and vessel-induced vibrations.

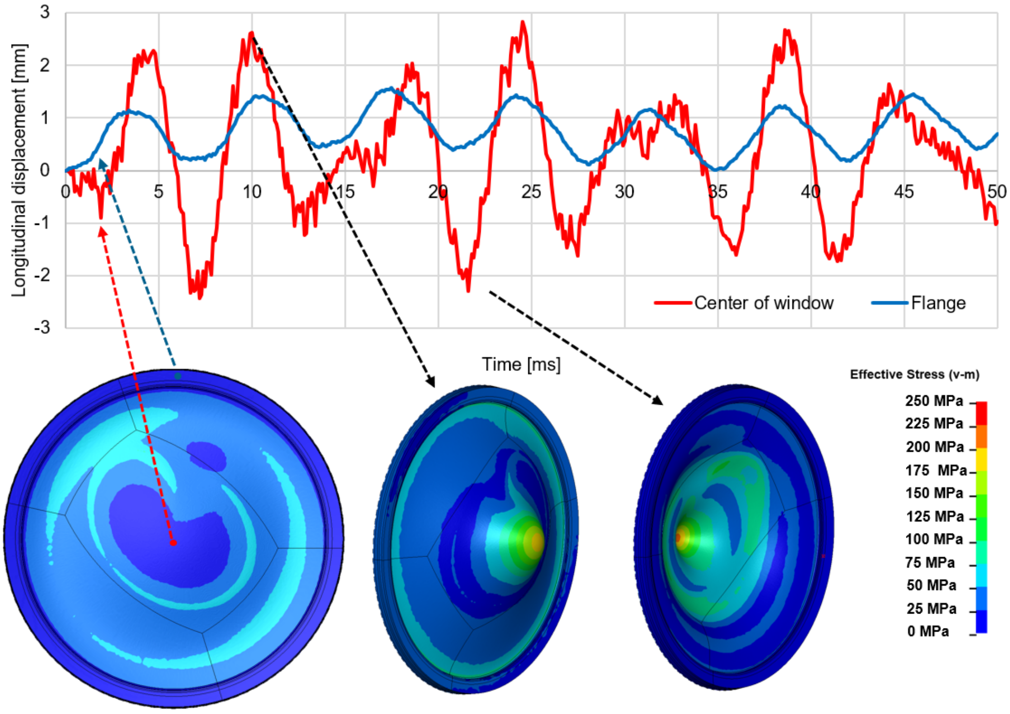

Figure 11.3: Longitudinal displacement of the downstream window centre and flange as a function of time, together with the corresponding von Mises stress fields during different stages of the vibration response. Deformations scaled with factor of 50.

Figure 11.3 presents the longitudinal displacement of the downstream window as a function of time, measured at the window centre and at the flange. The lower-left image shows the quasi static von Mises stress field at 1 ms, after the beam impact but before the vessel-induced vibrations dominate. The two images on the right illustrate different vibration states of the window, corresponding to outward and inward deformation.

The window deformation is primarily driven by the motion of the flange. As the vessel expands and contracts, the flange moves accordingly, while the inertia of the window causes a slight delay in the response of the window centre. This behaviour can be observed in the displacement history, where the centre of the window follows the flange motion with a small phase lag. After the second oscillation, some deviations appear in the displacement curve, which are attributed to the excitation of a secondary vibration mode of the window. The flange motion induces cyclic deformation of the window, producing alternating tensile and compressive stresses at the centre. The peak von Mises stress reaches approximately 225 MPa. Importantly, the maximum stress occurs at the window centre, whereas the beam-induced thermal stresses are located along the beam sweep. As a result, no significant superposition of these two stress fields occurs.

In addition, a modal analysis of the window was performed to verify that its natural frequencies do not coincide with the longitudinal vibration mode of the vessel, which could otherwise lead to resonance and amplification of the stress response. The natural frequency associated with the window deformation mode observed in Figure 11.3 was found to be 225 Hz, whereas the dominant longitudinal vibration mode of the vessel occurs at 158 Hz. Since these frequencies are well separated, dynamic resonance between the vessel vibration and the window deformation mode is not expected, and no significant amplification of window stresses is anticipated.

To evaluate the fatigue performance of the window, the same methodology as described in the fatigue assessment of the vessel in Chapter 10 was applied. The stress histories from the LS-Dyna simulations were exported and processed in nCode DesignLife. For each element, the peak stress amplitude evaluated from rainflow counting and its corresponding mean stress were extracted and used for the fatigue evaluation.

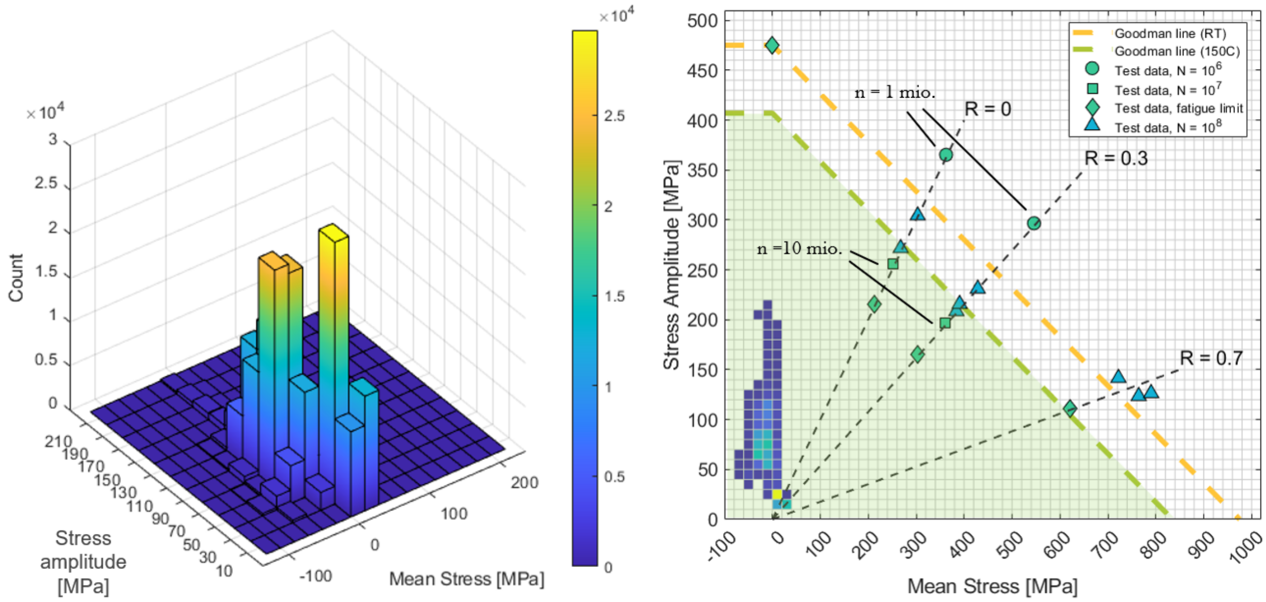

(a) Histogram showing the distribution of mean stress and stress amplitude obtained from the fatigue analysis.

(b) Goodman diagram of the stress histogram including reference fatigue limits from literature [30, 54].

Figure 11.4: Fatigue assessment of the beam dump windows based on vibration-induced stresses.

Figure 11.4a shows the resulting histogram of stress amplitude versus mean stress, while Figure 11.4b presents the same data plotted in a Goodman diagram using the same reference fatigue limits as previously [30, 54].

Most elements exhibit mean stresses close to 0 MPa, with a slight bias towards compressive values due to the compressive stresses induced by energy deposition. The maximum stress amplitude reaches approximately 220 MPa, occurring in elements near the centre of the window where the vibration-induced deformation is largest. When compared with the Goodman limit, all stress states remain well below the allowable fatigue boundary. Even when evaluated at an elevated temperature of 150 °C, which is conservative for the window centre, as this region is only weakly affected by energy deposition, the maximum stress amplitude of 220 MPa remains significantly below the corresponding Goodman limit of approximately 410 MPa. For reference, the infinite-life fatigue strength of Titanium Grade 5 at room temperature is indicated by the diamond marker at 475 MPa.

These results demonstrate that the window stresses during full-intensity beam impacts remain comfortably within the fatigue limits of the material, indicating that fatigue is not a limiting factor for the window design.

11.3 Assessment of the Bolted Window-Flange Connection

The bolted connection between the window and the vessel flange is a critical feature of the beam dump, as it represents the most likely location for potential leakage under dynamic loading. This joint is located at the extremities of the beam dump, where the accelerations induced by the vessel vibration are highest, reaching up to approximately 100 g. The connection must therefore maintain both structural integrity and leak tightness under the dynamic conditions, while still allowing the windows to be removed for inspection of the core. The ability to disassemble the joint is also the primary reason the windows are not welded to the vessel, as they provide the only practical access to the interior of the dump.

Figure 11.5 illustrates the design of the bolted window-flange connection. The joint is based on a ConFlat-type vacuum seal, where knife edges on both the flange and the window plastically deform a copper gasket to produce a metal-to-metal vacuum seal [48]. The geometry is designed such that, once the gasket is sufficiently compressed, the flange and window surfaces come into contact, thereby limiting further gasket compression and preventing excessive deformation. Three small screws on the inner side of the window retain the gasket during assembly. The screw heads are captured in recesses in the flange, preventing loosening during operation. When the window is removed, these screws retain the gasket so that the window and gasket can be extracted as a single unit, which is advantageous given the activated environment in which the component operates.

The joint preload is applied by 72 M14 titanium Grade 5 bolts distributed uniformly around the flange circumference to ensure uniform gasket pressure. Titanium was selected to avoid differential thermal expansion between the bolts and the window. Although CFD simulations by Notarianni indicate that the maximum temperature in this region is approximately 60 °C [19], using the same material eliminates differential expansion during cyclic thermo-mechanical loading, preventing gradual bolt wear and reducing the risk of preload loss over time. Stainless steel Helicoil inserts are installed in the flange threads to prevent galling and cold welding, which are known risks in titanium threaded connections due to the higher friction coefficient of the naturally occurring oxide layer and the potential exposure of highly weldable base metal when this thin oxide film is removed during tightening [44]. This ensures reliable tightening and repeatability in the assembly preload [4]. Nord-Lock washers are used to prevent bolt loosening caused by vibration from vessel dynamics [58]. Three alignment guides are placed around the circumference to centre the window relative to the flange during installation. These guides also enable robotic handling, which may be required due to the radioactive environment.

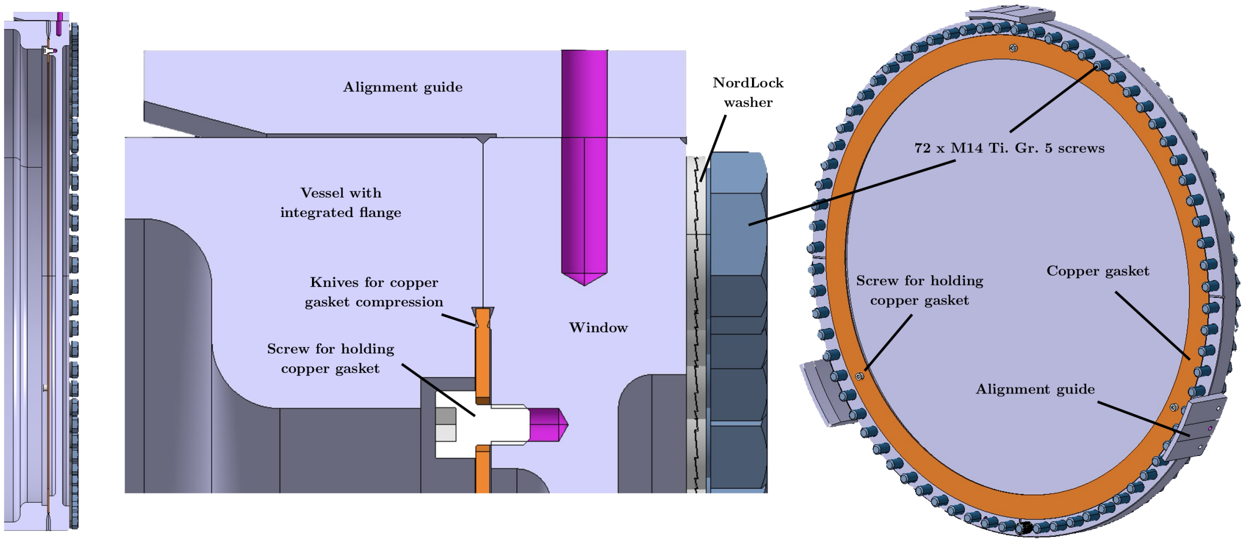

Figure 11.5: Overview of the window–vessel bolted connection. Left: full cross-sectional view. Centre: detailed view of the connection design. Right: window assembly showing bolts, alignment guides, and gasket.

11.3.1 Dimensioning of the Bolts

The bolted joint was dimensioned using a combined analytical and numerical approach. Preliminary dimensioning was carried out according to VDI 2230 [12] to determine the required preload and verify the load capacity of the screws. The calculations can be found in Appendix A.13 The analytical results were subsequently validated using finite element simulations of the flange assembly.

Because the loading of the connection is predominantly axial and the window mass is only approximately 30 kg, transverse loads due to gravity were neglected. The 72 bolts are uniformly distributed along the circumference and therefore assumed to experience identical loading. The loads acting on each bolt are summarised in Table 11.2. The dominant load is the force required to plastically deform the copper gasket and establish the vacuum seal. This force was determined using a 2D axisymmetric finite element model in which the flange knives were displaced into the copper gasket using an elastic–plastic material model. The reaction force corresponding to a total closure of 0.68 mm, equivalent to a knife penetration of 0.34 mm on each wide which is typical for ultra-high-vacuum ConFlat joints [48], was used to determine the required compression force. The internal pressure load results from the 0.2 bar relative argon overpressure inside the vessel acting on the window area. The inertial load originates from the dynamic response of the beam dump. It was calculated using Newton's second law with the longitudinal acceleration of approximately 100 g, obtained from the displacement response of the flange together with the mass of the window. A detailed description of the calculations can be found in Appendix A.13.

Table 11.2: Classification of external loads acting on each screw.

Symbol	Source of load	Description	Load per screw [kN]
F_{GC}	Gasket compression	Force required to plastically deform the copper gasket	7.00
F_{IP}	Internal pressure	Load from internal vessel pressure acting on the window	0.09
F_A	Dump vibration	Inertial load due to vessel acceleration	±0.42
Total		Range of load during operation	6.68 - 7.51

The bolt diameter was determined using the VDI 2230 load classification diagram provided in Appendix A.13. This procedure considers the maximum service load of 7.5 kN and applies additional factors accounting for eccentric loading, dynamic effects, and the tightening method. The resulting required bolt load capacity is approximately 40 kN, corresponding to a standard M14 bolt of strength class 8.8. VDI 2230 specifies a proof strength of 640 MPa for class 8.8 bolts and approximately 820 MPa for titanium alloys. An equivalent titanium bolt diameter providing the same load capacity can therefore be estimated to 12.4 mm. Selecting an M14 titanium bolt is therefore conservative and provides additional safety margin.

11.3.2 Torque-Tension Experimental Assessment

The relationship between applied torque and obtained tension in the screw is highly dependent on the friction in the threads and under the bolt head. To obtain reliable friction coefficients, a dedicated test campaign was conducted in collaboration with an external laboratory. The campaign measured the friction behaviour of M14 titanium Grade 5 bolts in threads of the same material with and without stainless-steel Helicoil inserts. This relation can be experimentally verified using a torque–tensile test bench, shown in Appendix A.12.1, in accordance with ISO 16047 [5] and is expressed as:

$$M_A = F_M \left(\frac{P}{2\pi} + \frac{d_2}{2}\mu_G + \frac{D_0 + d_h}{4}\mu_K \right) \tag{11.1}$$

Here, M_A denotes the tightening torque applied during assembly and F_M the resulting bolt tensile preload. The term P represents the thread pitch, while d_2 is the pitch diameter. μ_G and μ_K is the friction coefficient in the thread and under the bolt head respectively. The parameters D_0 and d_h denote the outer bearing diameter under the bolt head and the inner diameter of the bearing surface, respectively.

The torque-tension test bench measures the axial clamping force using a load cell, the applied tightening torque on the screw, and the torque transmitted to a reaction plate that simulates the bearing surface under the screw head. Friction between the bolt head and the reaction plate generates a measurable torque corresponding to the bearing friction μ_K. Since the total tightening torque and the bearing friction torque are measured directly, the thread friction μ_G can be determined as the only unknown in Equation 11.1.

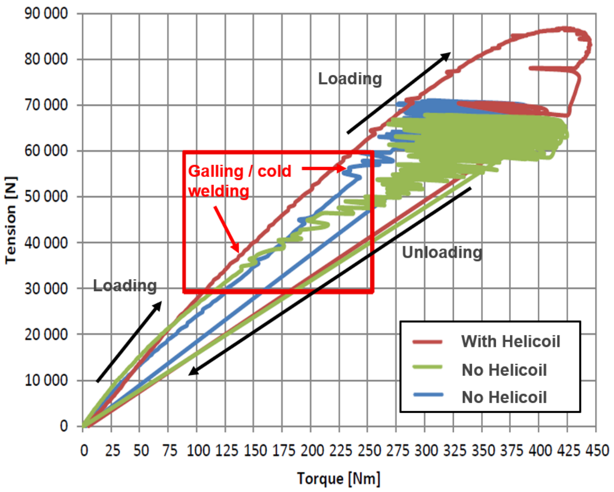

(a) Torque-tension relationship for configurations with and without the use of Helicoil thread inserts.

Spec.	Tension [N]	Torque [Nm]	Friction μ_G	Friction μ_K	Friction μ_{glob}
1	305	14.7	0.27	0.19	0.23
2	599	16.4	0.26	0.19	0.23
3	208	15.6	0.30	0.17	0.24
4	210	18.3	0.24	0.20	0.22
5	300	9.3	0.25	0.11	0.19
6	206	13.4	0.23	0.16	0.20
7	104	7.8	0.16	0.13	0.15
8	101	4.2	0.16	0.11	0.14
9	202	9.3	0.22	0.23	0.22
Min.	101	4.2	0.16	0.11	0.14
Max.	599	18.3	0.30	0.23	0.24
Avg.	248	12.1	0.23	0.16	0.20
σ [%]	-	-	20.4%	25.0%	18.9%
2σ [%]	-	-	40.7%	50.0%	37.8%

(b) Statistical post-processing of torque-tension relationship for configuration with Helicoil inserts.

Figure 11.6: Experimental results of torque-tension relationship for titanium Grade 5 screws in titanium Grade 5 threads with and without Helicoil inserts.

Figure 11.6a presents three torque–tension tests. One test was performed with a Helicoil insert (red), while the other two were conducted without Helicoils, meaning the titanium screw was tightened directly into a titanium thread. The loading and unloading directions are indicated in the figure. For the tests without Helicoils, galling occurs and appears as abrupt horizontal scatter in the torque signal due to rapidly increasing friction and local cold welding between the titanium surfaces. Galling starts at approximately 130 Nm in the green test and around 225 Nm in the blue test. Beyond this point, the torque–tension relationship becomes highly irregular, indicating that the applied torque no longer reliably corresponds to bolt preload. In contrast, the test with a Helicoil insert shows a smooth and stable torque–tension curve up to the elastic limit of the screw at approximately 400 N m, corresponding to a preload of about 87 kN. This demonstrates that the Helicoil insert effectively prevents galling and enables a predictable torque–preload relationship.

Table 11.6b summarises nine torque–tension tests performed with Helicoil inserts. While individual curves remain stable, noticeable scatter in preload is observed between screws, which is typical due to manufacturing imperfections. The table also provides the average friction coefficients and their standard deviations, illustrating the inherent variability. In addition the torque wrench intended for installation of the window bolts was calibrated using a torque cell, and the accuracy was found to be within ±10% from the measured torque on the wrench to the actually applied torque.

Assuming a target preload of 55 kN, corresponding to an installation torque of 200 Nm, Figure 11.7a and Table 11.7b show the combined effect of torque wrench deviation and the statistical variation of the friction coefficients. The tightening torque is assumed to vary by ±10%, while the friction coefficients are taken as the mean value and $\pm\sigma$ obtained from the torque-tension tests. The lowest achievable preload occurs for the conservative combination of reduced installation torque and high friction coefficients, while the highest preload occurs for increased installation torque and low friction coefficients. These two combinations define the lower and upper bounds of the achievable preload. As shown in Table 11.7b, the resulting preload can be as low as 41.6 kN, corresponding to 63% of the bolt yield strength, or as high as 76.5 kN, corresponding to 98% of the yield strength. Although the simultaneous occurrence of these extreme combinations of torque and friction is unlikely, they are adopted as bounding cases in the validation to ensure a conservative assessment of the joint. The range remains higher than the minimum requirement from VDI 2230 of 40 kN and lower than the yield strength of the bolt.

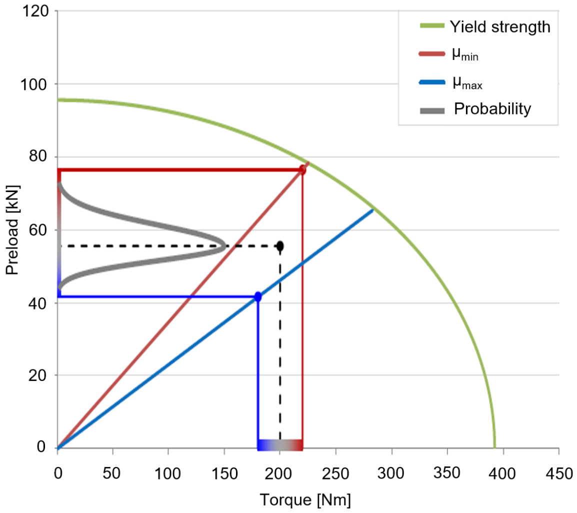

(a) Plot view showing statistical scatter and yield surface.

Friction	Minimal $-\sigma$	Nominal Mean	Maximal $+\sigma$
M_A [N m]	180	200	220
F_M [kN]	41.6	55.5	76.5
σ_t [MPa]	361	481	662
τ_{tor} [MPa]	218	243	268
σ_{eq} [MPa]	522	639	809
σ_{eq} [% $R_{p0.2b}$]	63%	77%	98%

(b) Corresponding tightening torque, preload, tensile stress, torsional stress, and equivalent von Mises stress.

Figure 11.7: Torque–tension relationship considering ± one standard deviation of the friction coefficients. The resulting preload states are compared with the yield surface of the Ti-6Al-4V (Grade 5) screw.

11.3.3 Validation of Joint Design

The joint design was validated using analytical calculations according to VDI 2230 (see Appendix A.13.1) together with a finite element submodel of the window-flange connection (see Appendix A.13.2). The VDI 2230 calculations verify the maximum expected loading of the bolt, while the finite element model is used to assess potential prying effects and gasket leakage at minimum expected loading of the bolt.

Figure 11.8 shows the finite element model developed. It represents a 5° sector of the joint containing a single bolt and was developed in Ansys static structural. Symmetry boundary conditions are applied on the radial faces. The model captures the state of the bolted connection under gasket compression from the knives, internal pressure, and inertia forces caused by vessel vibration. The bolt is modelled as a solid with a bonded contact to the flange hole, allowing the bolt preload to be applied directly. A detailed description of the model including mesh, contacts, boundary conditions and loads is provided in Appendix A.13.2.

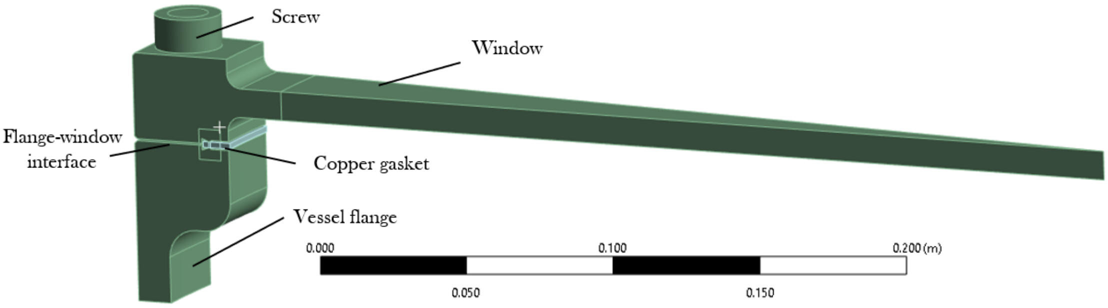

Figure 11.8: Finite element submodel used to assess the window–flange bolted connection.

The model was evaluated for the maximum (76.5 kN) and minimum (41.6 kN) bolt preloads derived from the experimental torque-friction study. Compared to this load range, the additional service loads are negligible: the inertia force from window vibrations contributes approximately 0.4 kN per bolt and the internal pressure approximately 0.09 kN per bolt. The joint can therefore be considered effectively static.

For the upper load limit it must be verified that the bolt remains below its allowable stress limit, while the lower load limit must ensure that full contact between the flange and window is maintained without prying that could reduce gasket compression and cause leakage.

For the upper load limit case, the VDI 2230 methodology was applied using inputs extracted from the finite element model. For the maximum bolt load of 76.5 kN, the bending moment obtained from the model was 7.6 N m. The bending moment is primarily introduced by the compression of the gasket, which acts eccentrically relative to the bolt axis. As a result, the reaction force generated by the gasket is offset from the bolt load path, producing a bending moment that is carried by the bolt. According to VDI 2230, bolts should primarily carry axial tension and not significant bending moments. In this case, the calculations show that the bending contribution accounts for only 6% of the total bolt equivalent stress, which is considered acceptable. The equivalent stress, including axial preload, bending, and torsional shear, is 762 MPa. This remains below the bolt yield strength of 820 MPa, corresponding to a safety factor of 1.075. This confirms that even the conservative maximum load case, representing the lowest measured friction values during installation, remains acceptable.

For the minimum preload of 41.6 kN, the main concern is potential prying between the window and flange and insufficient compression of the gasket. The contact pressure at the flange-window interface is shown in Figure 11.9a. Due to the proximity of the gasket, the contact pressure decreases toward the inner diameter. Nevertheless, a minimum pressure of 16 MPa is maintained at the location furthest from the bolt and closest to the gasket, while the average pressure over the interface is 57 MPa. This confirms that the entire interface remains in compression and that prying is not expected during operation. The contact pressure between the copper gasket and the knife edge is shown in Figure 11.9b. The pressure distribution is uniform around the circumference, indicating that the 5° bolt spacing provides sufficient load transfer to maintain consistent gasket compression even at the minimum preload of 41.6 kN.

Since the maximum temperature in this region was determined to be 60 °C [19], thermal effects were neglected. Titanium Grade 5 is used for the window, flange, and bolts, eliminating differential thermal expansion and limiting temperature-induced stresses.

Additional analytical verifications, including fatigue assessment, validation of thread engagement length, and contact pressure under the bolt head, are presented in Annex A.13.1.

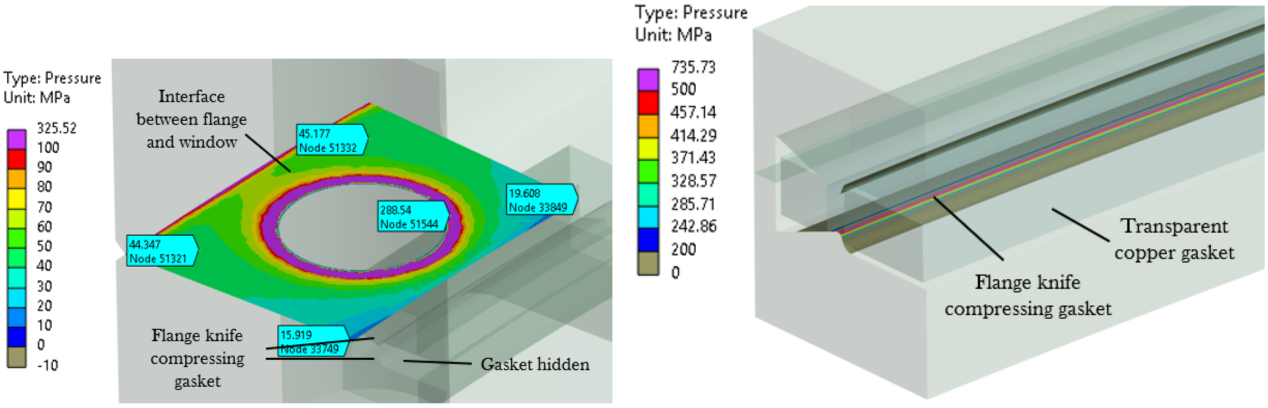

(a) Contact pressure distribution at the flange-window interface.

(b) Contact pressure distribution between the copper gasket and the knife edge showing uniform gasket compression.

Figure 11.9: Finite element results used to verify the minimum preload condition of 41.6 kN of the window-flange bolted connection.

This validation demonstrates that applying an installation torque of 200 Nm results in a bolt preload range of 41.6 kN to 76.5 kN. This range accounts for both torque application deviations from the wrench and variations in the friction coefficients of individual screws, and remains acceptable for operation.

CHAPTER

12

Conclusion

The development of the High-Luminosity beam dumps for the Large Hadron Collider addressed the challenges introduced by the significantly increased beam energy and the resulting dynamic loading of the system. The objective of this work was to develop a beam dump vessel capable of operating reliably under the High-Luminosity conditions while respecting the strict geometric and integration constraints of the existing installation. The central research question guiding the work was:

How can the HL-LHC beam dump vessel be designed to withstand 10 years of cyclic dynamic behaviour under HL operational conditions, without compromising structural integrity or violating existing integration constraints?

Finite element analyses of the currently operational Run 3 beam dumps showed that the existing design would experience excessive dynamic stresses caused by rapid energy deposition and the excitation of structural vibration modes, preventing safe operation under HL-LHC conditions. Several mitigation strategies were investigated, and changing the vessel material to titanium Grade 5 was identified as the only solution providing a substantial reduction in stress. Its low density, low coefficient of thermal expansion, low stiffness, and high strength significantly reduce the stress response during beam impact. The resulting design achieves a safety factor of approximately 3 against yielding under the full intensity High-Luminosity beam impact.

The cyclic loading of the vessel was assessed through a fatigue analysis using conservative assumptions for the load spectrum and damping behaviour. The vessel operates below its fatigue limit, resulting in effectively infinite fatigue life. The electron beam welds joining the vessel segments were also evaluated. Their asymmetric geometry introduces a stress concentration at the weld root, and their fatigue performance was verified experimentally, demonstrating a safety margin of 4.1 in cycles assuming half the damping measured in Run 3 at a 1% probability of failure and 95% confidence level with 10 years of operation.

The beam windows were analysed using finite element simulations capturing their response following beam impact. A safety factor of approximately 3.5 against yielding was obtained under nominal operational conditions, and the vibrations transmitted from the vessel remain within the infinite fatigue life regime of the windows. The bolted flange connection was validated using VDI 2230 calculations supported by experimental investigations, confirming the appropriate dimensioning of the joint.

In conclusion, the proposed High-Luminosity beam dump design resolves the limitations of the Run 3 configuration while respecting the integration constraints. The vessel, welds, windows, and bolted connections were validated through numerical analyses and experimental investigations, demonstrating that the design can safely withstand the cyclic loading during HL-LHC operation from the increased 710 MJ stored beam energy.

With the detailed design finalized in this work, the project now enters the production phase. Three HL beam dumps of this design (two operational units and one spare) will be manufactured. The main procurement contract for the titanium Grade 5 vessel and the core materials has been awarded to industrial suppliers, with the pre-series components expected to arrive in late summer 2026.

Bibliography

[1] AISI 318 LN. URL: https://woite-edelstahl.com/aisi318lnen.html.

[2] ASME B4.1: Limits and Fits Standard. Technical report.

[3] FLUKA.CERN website.

[4] HELICOIL® - Thread inserts for metals - Böllhoff. URL: https://www.boellhoff.com/ch-en/products/special-fasteners/helicoil-thread-inserts-for-metals/.

[5] ISO 16047:2005 - Fasteners - Torque/clamp force testing. Technical report. URL: https://www.iso.org/standard/27788.html.

[6] Our Member States. URL: https://home.cern/about/who-we-are/our-governance/member-states.

[7] *LS-DYNA. Keyword User's Manual, Volume I.* 1992. URL: https://www.dynasupport.com/manuals/ls-dyna-manuals/ls-dyna_manual_volume_i_r13.pdf.

[8] *LS-DYNA. Keyword User's Manual, Volume II, Material Models.* 1992. URL: https://www.dynasupport.com/manuals/ls-dyna-manuals/ls-dyna_manual_volume_ii_r13.pdf.

[9] Superconducting magnets for the LHC. 2007. URL: https://www.cea.fr/multimedia/Documents/publications/clefs-cea/archives/en/Clefs%2056_p04_09_MeurisGB.pdf.

[10] ISO 12107:2012: Metallic materials - Fatigue testing - Statistical planning and analysis of data. Technical report, 2012.

[11] nCode DesignLife Theory Guide. Technical report, 2013. URL: https://oss.jishulink.com/caenet/forums/upload/2014/09/25/191/45541828203375.pdf.

[12] VDI 2230 Part 1: Systematic calculation of highly stressed bolted joints. Technical report, 2014.

[13] ASME B381-21 - Specification Specification for Titanium and Titanium Alloy Forgings. Technical report, 2021.

[14] Ahdida C., Bozzato D., Calzolari D., Cerutti F., Charitonidis N., Cimmino A., Coronetti A., D'Alessandro G. L., Donadon Servelle A., Esposito L. S., Froeschl R., García Alía R., Gerbershagen A., Gilardoni S., Horváth D., Hugo G., Infantino A., Kouskoura V., Lechner A., Lefebvre B., Lerner G., Magistris M., Manousos A., Moryc G., Ogallar Ruiz F., Pozzi F., Prelipcean D., Roesler S., Rossi R., Sabaté Gilarte M., Salvat Pujol F., Schoofs P., Stránský V., Theis C., Tsinganis A., Versaci R., Vlachoudis V., Waets A., and Widorski M. New Capabilities of the FLUKA Multi-Purpose Code. *Frontiers in Physics 9, 788253 (2022)*, 2022.

[15] Roger Bailey and Paul Collier. Standard Filling Schemes for Various LHC Operation Modes (Revised). Technical report.

[16] Simon Baird. Accelerators for Pedestrians. Technical report, 2007. URL: https://cds.cern.ch/record/1017689/files/ab-note-2007-014.pdf.

[17] Giuseppe Battistoni, Till Boehlen, Francesco Cerutti, Pik Wai Chin, Luigi Salvatore Esposito, Alberto Fassò, Alfredo Ferrari, Anton Lechner, Anton Empl, Andrea Mairani, Alessio Mereghetti, Pablo Garcia Ortega, Johannes Ranft, Stefan Roesler, Paola R. Sala, Vasilis Vlachoudis, and George Smirnov. Overview of the FLUKA code. *Annals of Nuclear Energy*, 82:10–18, 2015.

[18] Battistoni G., Boehlen T., Cerutti F., Chin P.W., Esposito L.S., Fassò A., Ferrari A., Lechner A., Empl A., Mairani A., Mereghetti A., Garcia Ortega P., Ranft J., Roesler S., Sala P.R., Vlachoudis V., and Smirnov G. Frontiers in Physics 9. *Annals of Nuclear Energy 82*, pages 10–18, 2015.

[19] Giovanni Binda Notarianni. Design and validation of an enhanced cooling system model for the High-Luminosity LHC beam dump at CERN. 2025. URL: https://repository.cern/records/xm7zn-0s688, `doi:10.17181/XM7ZN-0S688`.

[20] C Bracco, M J Barnes, W Bartmann, E Carlier, L Ducimetiere, M Fraser, B Goddard, A Lechner, N Magnin, M Meddahi, A Perillo Marcone, V Senaj, J Uythoven, and F M Velotti. LBDS and Injection System. Technical report.

[21] C Bracco, W Bartmann, J Borburgh, L S Ducimetiere, E Carlier, B Goddard, N Magnin, T D Mottram, V Senaj, and C Wiesner. LBDS performance in RUN 2. Technical report.

[22] Jules Braken. Thermomechanical Analysis of Beam Dump Devices for the Large Hadron Collider, 2024. URL: https://repository.cern/records/cefbj-whk52.

[23] M Calviani. Beam Intercepting Devices at CERN Types, Challenges, Design, R&D and Operation - Lecture 1. Technical report, 2021.

[24] M Calviani. Beam Intercepting Devices at CERN Types, Challenges, Design, R&D and Operation - Lecture 2. Technical report, 2021.

[25] R Courant, K Friedrichs, and H Lewy. On the Partial Difference Equations of Mathematical Physics. 1967. URL: https://web.stanford.edu/class/cme324/classics/courant-friedrichs-lewy.pdf.

[26] Dipartimento Di, Sistemi Elettrici, Ed Automazione, and Ing Roberto Filippini. Dependability Analysis of a Safety Critical System: The LHC Beam Dumping System at CERN. Technical report, 2005.

[27] J H Dieperink, J.-L Bretin, E Carlier, L Ducimetière, G H Schröder, E Vossenberg, and Geneva Switzerland. Design aspects related to the reliability of the LHC beam dump kicker systems. Technical report, 1997.

[28] L Evans and P Bryant. LHC Machine. Technical report. URL: http://www.iop.org/EJ/jinst/.

[29] R. Filippini, E. Carlier, L. Ducimetière, B. Goddard, and J. Uythoven. Reliability analysis of the LHC Beam Dumping System. In *Proceedings of the IEEE Particle Accelerator Conference*, volume 2005, pages 1201–1203, 2005. `doi:10.1109/PAC.2005.1590707`.

[30] Yoshiyuki Furuya and Etsuo Takeuchi. Gigacycle fatigue properties of Ti–6Al–4V alloy under tensile mean stress. *Materials Science and Engineering: A*, 598:135–140, 3 2014. URL: https://www.sciencedirect.com/science/article/abs/pii/S0921509314000409, `doi:10.1016/j.msea.2014.01.019`.

[31] F Gerigk. Cavity types. Technical report.

[32] H Guerin, A Lechner, N Solieri, G Banks, A Lund, C Bracco, and Y Dutheil. Beam-Induced Energy Deposition in Dump Block Assembly. 2025. URL: https://indico.cern.ch/event/1487681/contributions/6346906/attachments/3029141/5346992/HLTDE_designreview_11032025_v2.pdf.

[33] M Guinchard. Experimental modal analysis of LHC beam dump on the final support for run 3 - EDMS 2339554. Technical report, 2020. URL: https://edms.cern.ch/document/2339554/1.

[34] John O Hallquist. *LS-DYNA Theory manual*. 2006. URL: https://www.dynasupport.com/manuals/additional/ls-dyna-theory-manual-2005-beta.

[35] Armin Hermann, Lanfranco Belloni, Gerhard John; Krige, Ulrike Mersits, and Dominique Pestre. *Launching the European Organization for Nuclear Research*. 1987.

[36] William Kappelgaard. Structural Analyses of the Support Structure for the Large Hadron Collider Beam Dumps. Technical report, 2026.

[37] N Karastathis, M Barnes, H Bartosik, K Brodzinski, X Buffat, F Cerutti, S Fartoukh, B Goddard, G Iadarola, S Le Naour, A Lechner, J Maestre Heredia, A Mereghetti, E Metral, D Missiaen, N Mounet, F X Nuiry, S Papadopoulou, Y Papaphilippou, B Petersen, G Rumolo, B Salvant, C Schwick, M Solfaroli Camillocci, G Sterbini, H Timko, R Tomas Garcia, J Uythoven, J Wenninger, and Geneva Switzerland. LHC Run 3 Configuration Working Group Report, in Proceedings of the 9 h Evian Workshop on LHC Beam Operation. Technical report, 2019.

[38] Michael Krause. *CERN how we found the Higgs boson*. World Scientific, 2014.

[39] Asbjørn L. Analysis of an Electron Beam Welded Ti6Al4V Joint with an Integrated Backing Ring: Using analytical, numerical and experimental methods. Technical report, 2025. URL: https://kbdk-aub.primo.exlibrisgroup.com/permalink/45KBDK_AUB/3b147k/alma9922350598105762.

[40] Lisa Latussek, Gabriel Banks, Nicola Solieri, Asbjoern Lund, and Francesco Velotti. Testing core materials for the HL-LHC Beam Dumps: The HRMT65-HLTDE Experiment . Technical report, 2025.

[41] Christoph Leyens and Manfred Peters. *Titanium and Titanium Alloys*. WILEY-VCH, 2003.

[42] Asbjoern Lassen Lund. Design of High-Irradiation Experimental Targets for Advanced Materials Testing. Technical report, Aalborg University, 2025. URL: https://kbdk-aub.primo.exlibrisgroup.com/permalink/45KBDK_AUB/a7me0f/alma9922100231005762.

[43] Asbjørn Lund. Design and Development of a New Beam Intercepting Device for the Super Proton Syncrotron at CERN. Technical report, CERN, Aalborg Univeristy, 2024. URL: https://kbdk-aub.primo.exlibrisgroup.com/permalink/45KBDK_AUB/a7me0f/alma9921966585605762.

[44] J U Ly. Technical PEM® - Thread Galling. Technical report, 2019. URL: https://www.pemnet.com/wp-content/uploads/sites/2/2022/06/Thread_Galling.pdf.

[45] J. Maestre, C. Torregrosa, K. Kershaw, C. Bracco, T. Coiffet, M. Ferrari, R. Franqueira Ximenes, S. Gilardoni, D. Grenier, A. Lechner, V. Maire, J. M. Martin Ruiz, E. Matheson, N. Solieri, A. Perillo-Marcone, T. Polzin, V. Rizzoglio, D. Senajova, C. Sharp, M. Timmins, and M. Calviani. Design and behaviour of the Large Hadron Collider external beam dumps capable of receiving 539 MJ/dump. *Journal of Instrumentation*, 16(11), 11 2021. `doi:10.1088/1748-0221/16/11/P11019`.

[46] J Maestre, C Torregrosa, J M Martin, K D Kershaw Senajova, M Ferrari, M Calviani, V Rizzoglio, A Lechner EN-STI, C Bracco TE-ABT, P Serrano EN-SMM, P D Pastuszak, J Guardia EN-MME, O Aberle, A Bertarelli, G Bregliozzi, A Dallocchio, M Di Castro, J-f Fuchs, M Guinchard, A Lechner, S Marsh, F-x Nuiry, A Perillo-Marcone, I Ruehl, S Sgobba, H Vincke, F Carra, K Kershaw, M Timmins, C Bertone, C Garion, W M Bartmann Calviani S Gilardoni M Lamont, P Collier, M Jimenez, B Goddard, R Folch, and R Losito. CERN LHC External Beam Dumps (TDE) Modifications and upgrades during LS2 for Run3. Technical report, 2021.

[47] A. Majorell, S. Srivatsa, and R. C. Picu. Mechanical behavior of Ti-6Al-4V at high and moderate temperatures-Part I: Experimental results. *Materials Science and Engineering: A*, 326(2):297–305, 3 2002. URL: https://www.researchgate.net/publication/248470361_Mechanical_behavior_of_Ti-6Al-4V_at_high_and_moderate_temperatures_-_Part_I_Experimental_results, `doi:10.1016/S0921-5093(01)01507-6`.

[48] Ana Miarnau Marin, Gerhard Schneider, and Raymond Veness. Development and test of a rectangular CERN ConFlat-type flange. 2015. URL: https://cds.cern.ch/record/2054374/files/1-s2.0-S0042207X15300440-main.pdf?subformat=pdfa, `doi:10.1016/j.vacuum.2015.08.018`.

[49] Maurizio Vretenar. Linac 4 design report. `doi:10.23731/CYRM-2020-006`.

[50] Y Papaphilippou, G Arduini, T Argyropoulos, W Bartmann, H Bartosik, T Bohl, C Bracco, S Cettour-Cave, K Cornelis, L Drosdal, J Esteban Muller, B Goddard, A Guerrero, W Höfle, V Kain, G Rumolo, B Salvant, E Shaposhnikova, H Timko, D Valuch, G Vanbavinckhove, J Wenninger, and Geneva Switzerland. Operational Performance of the LHC Proton Beams With The SPS Low Transition Energy Optics. Technical report, 2013. URL: https://inspirehep.net/files/8e660782f33bceaac37fc759a2f32232.

[51] Thomas Sven Pettersson and P Lefèvre. *The Large Hadron Collider conceptual design*. 1995. URL: https://cds.cern.ch/record/291782.

[52] Rune Bergmann. Simulations of Beam Dump Upgrades for HL Run of the LHC at CERN. Technical report, 2025.

[53] Francesc Salvat Pujol. An introduction to the Monte Carlo method for the simulation of radiation transport. 2022. URL: https://indico.cern.ch/event/1088394/.

[54] Michael D. Sangid. Ensuring fatigue performance via location-specific lifing in aerospace components made of titanium alloys and nickel-base superalloys. 2020. URL: https://www.researchgate.net/publication/352888321_Ensuring_fatigue_performance_via_location-specific_lifing_in_aerospace_components_made_of_titanium_alloys_and_nickel-base_superalloys, `doi:10.25394/PGS.12168018.v1`.

[55] N. Solieri, A. Lund, A. P. Bernardes, L. R. Buonocore, A. Cherif, S. De Man, M. Di Castro, S. Di Giovannantonio, G. Dumont, S. El-Idrissi, E. Farina, D. Grenier, E. Grenier-Boley, M. Himmerlich, A. Infantino, A. Lechner, R. Mouret, D. Pazem, A. T. Perez-Fontenla, E. Romagnoli, S. Sgobba, C. Tromel, C. Veiga Almagro, and M. Calviani. Decommissioning and post-irradiation examination of the LHC beam dumps. *Physical Review Accelerators and Beams*, 28(8):083001, 8 2025. URL: https://journals.aps.org/prab/abstract/10.1103/kv1h-3t6z, `doi:10.1103/kv1h-3t6z`.

[56] Marilena Streit-Bianchi. *CERN science bridging cultures*. 2018.

[57] K. Tokaji. High cycle fatigue behaviour of Ti-6Al-4V alloy at elevated temperatures. *Scripta Materialia*, 54(12):2143–2148, 6 2006. `doi:10.1016/j.scriptamat.2006.02.043`.

[58] Manual User. Nord-Lock Original Washers. Technical report. URL: https://www.nord-lock.com/globalassets/mediavalet/web-assets/downloads/user-manual/nord-lock-original-washers--user-manual-en2.pdf.

[59] J M Zazula and S Péraire. LHC Beam Dump Design Study, Part 1: Simulation of energy deposition by particle cascades: Implications for the dump core and beam sweeping system. Technical report, 1996.

[60] Jan M Zazula and Serge Péraire. LHC Beam Dump Design Study, Part 2: Thermal analysis; implication for abort repetition and cooling system, Tech. Rep. Technical report, 1997.

A

Appendix

A.1 Overview of Complete Set of Beam Dilution Failures

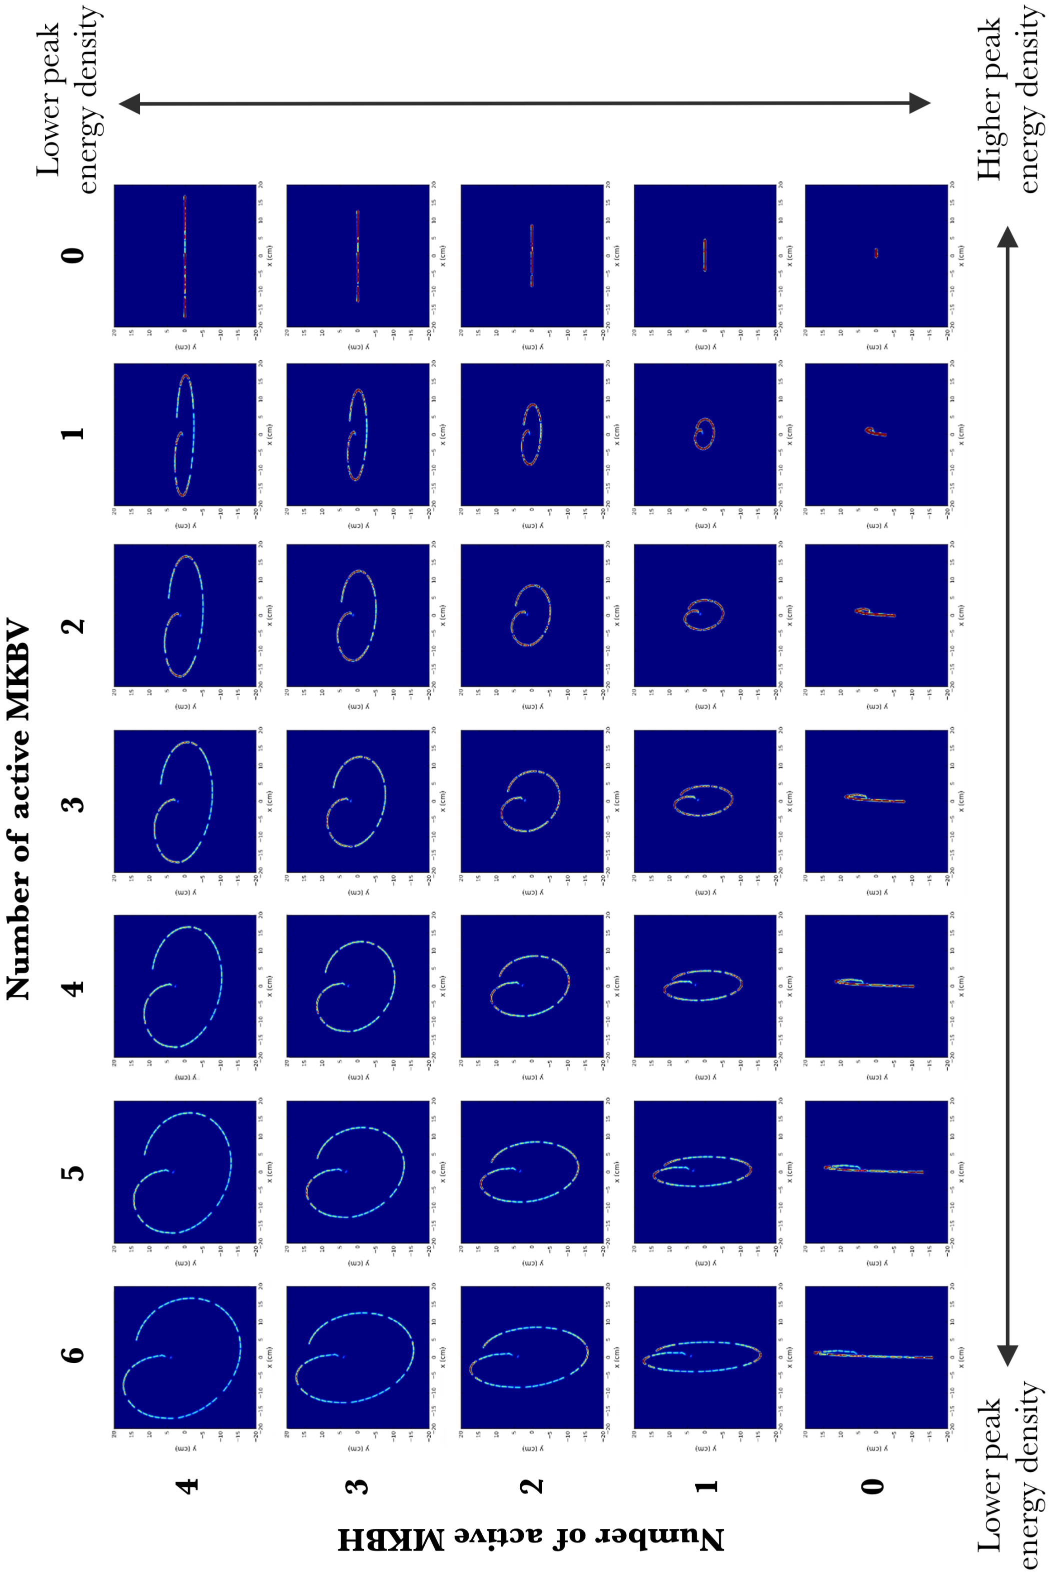

Figure A.1: Overview of all dilution failures based on the number of active horizontal and vertical kickers during the beam dump [21].

A.2 FLUKA Monte-Carlo Simulation for Energy Deposition Studies

Figure A.2 presents an example of a FLUKA Monte Carlo simulation in which 150 MeV protons impact a water target. The images on the left illustrate representative particle tracks generated during the simulation, while the plots on the right show the corresponding statistical uncertainty of the mean energy deposition as a function of depth in the target.

A comparison of Figures A.2b to A.2f demonstrates that increasing the number of simulated primary protons reduces the statistical uncertainty in the calculated energy deposition. This behaviour reflects the statistical convergence inherent to Monte Carlo methods, where larger sample sizes improve the accuracy of the estimated mean response. Simulating a sufficiently large number of primary particles is therefore essential to obtain a reliable approximation of the energy deposition profile. This requirement is well justified as the HL beam contains on the order of 6×10^{14} protons.

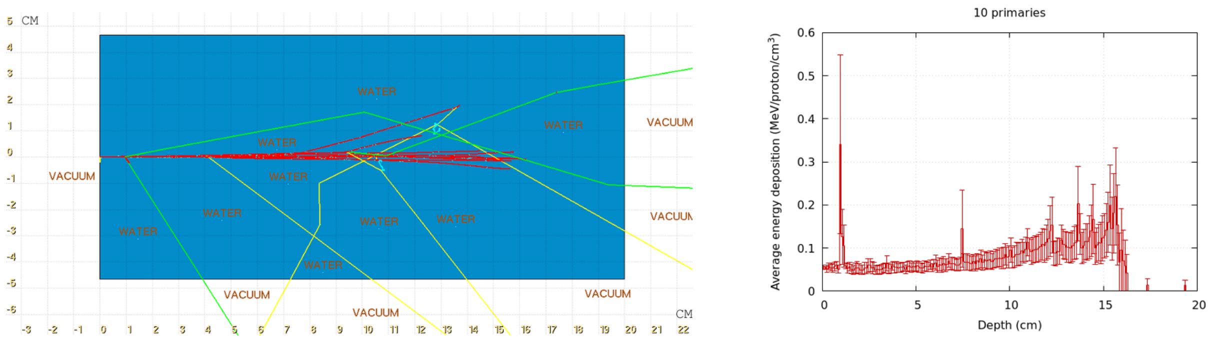

(a) 10 primary 150 MeV protons in H_2O.

(b) Average energy dep. as function of depth.

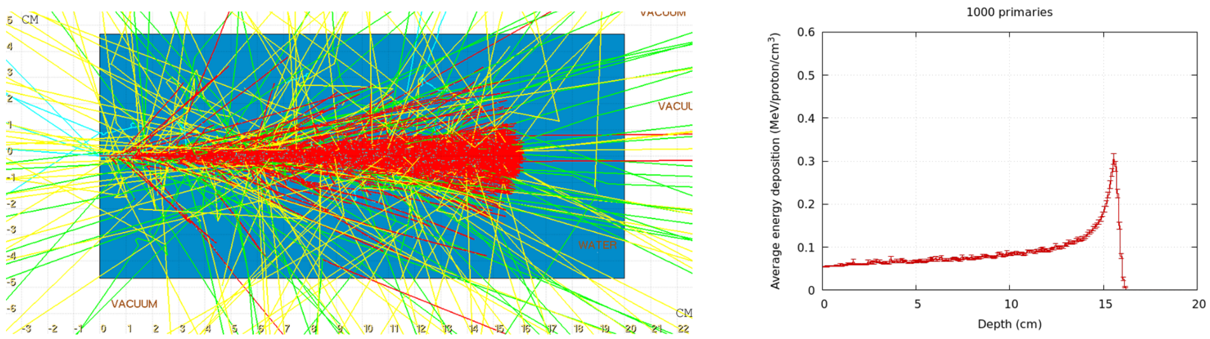

(c) 1 000 primary 150 MeV protons in H_2O.

(d) Average energy dep. as function of depth.

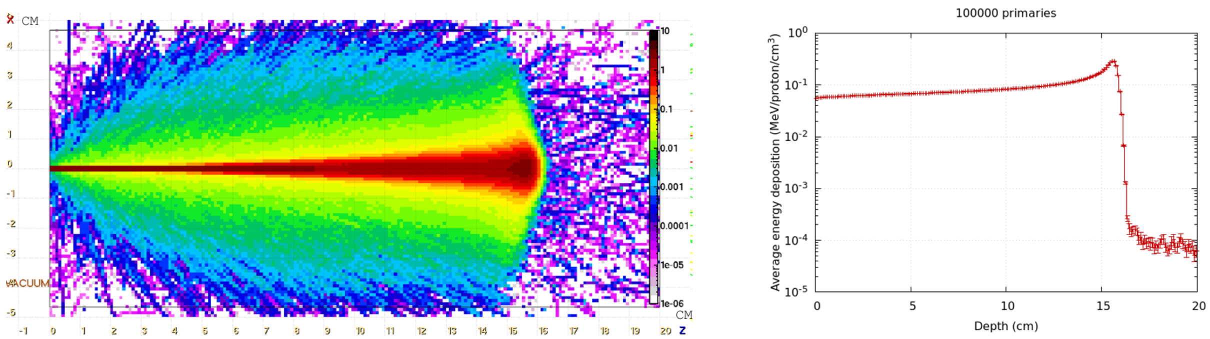

(e) 100 000 primary 150 MeV protons in H_2O. Proton tracks replaced by distribution.

(f) Average energy dep. (logarithmic scale) as function of depth.

Figure A.2: Illustration of a FLUKA Monte Carlo simulation of 150 MeV protons impacting a water target. In subfigures (a), (c), and (e), particle tracks are shown, where red lines represent protons, yellow lines photons, light blue lines electrons, and green lines neutrons. Subfigures (b), (d), and (f) show the statistical uncertainty on the mean energy deposition as function of depth [53, 43].

A.3 Global Model Mesh

This appendix provides additional information regarding the discretisation of the global finite element model.

The mesh consists of approximately 970k nodes and 870k solid hexahedral elements. The circumferential discretisation comprises 100 elements, while a longitudinal spacing of approximately 20 mm is used along the vessel axis. Four elements are introduced through the 12 mm vessel thickness. This through-thickness refinement is required to capture the variation in energy deposition and the resulting thermal gradients across the vessel wall. Considering the large diameter-to-thickness ratio of the vessel, four elements are deemed sufficient to accurately capture bending behaviour while maintaining computational efficiency.

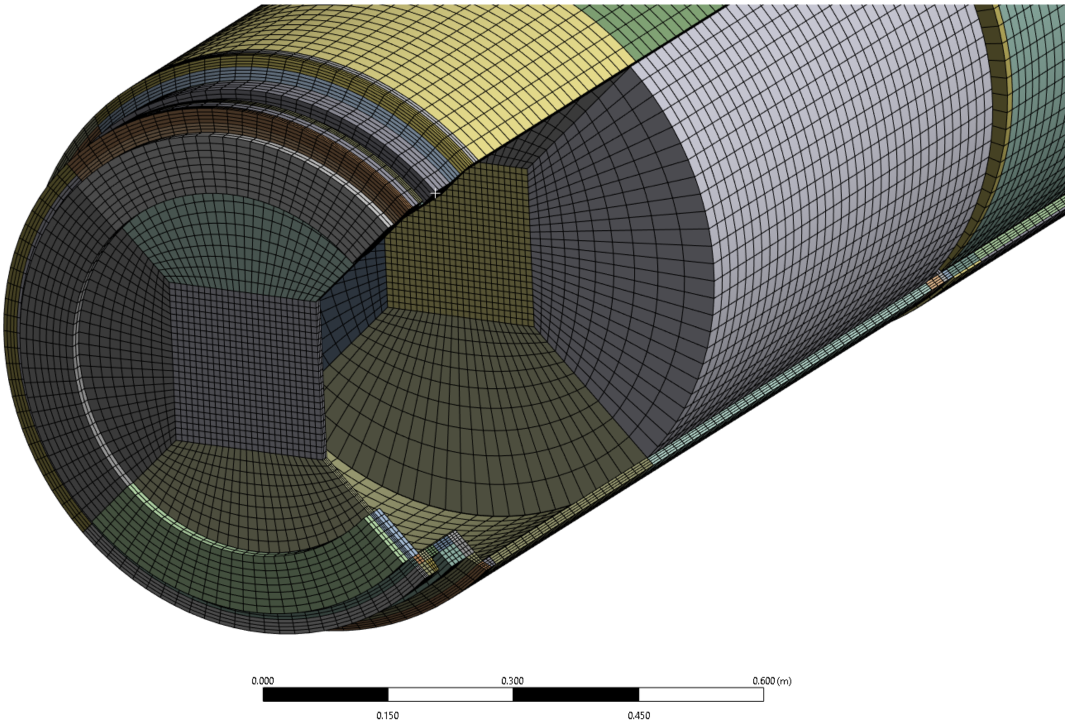

Figure A.3: Isometric view of the upstream region of the global model with the vessel partially removed to expose the internal components and nodal conformity between core and vessel.

Figure A.3 shows an isometric view of the upstream region of the model, with the right-hand side of the vessel removed to expose the internal components. The longitudinal and circumferential nodal alignment between the core and vessel can be observed, ensuring correct implementation of the shrink-fit interference.

Figure A.4 presents a cross-sectional view of the upstream region, highlighting the discretisation of the vessel wall. The four elements through the thickness are clearly visible.

Figure A.5 shows a histogram of the Ansys element quality metric defined as

$$\text{Quality} = c\,\frac{\text{volume}}{\sqrt{\left(\sum \text{edge length}^2\right)^3}}$$

where a value of 1 corresponds to a perfect cube and a value of 0 indicates a degenerate element with zero or negative volume. For hexahedral elements, the normalisation constant is $c = 41.57$. The distribution indicates that the majority of elements are within an acceptable quality range.

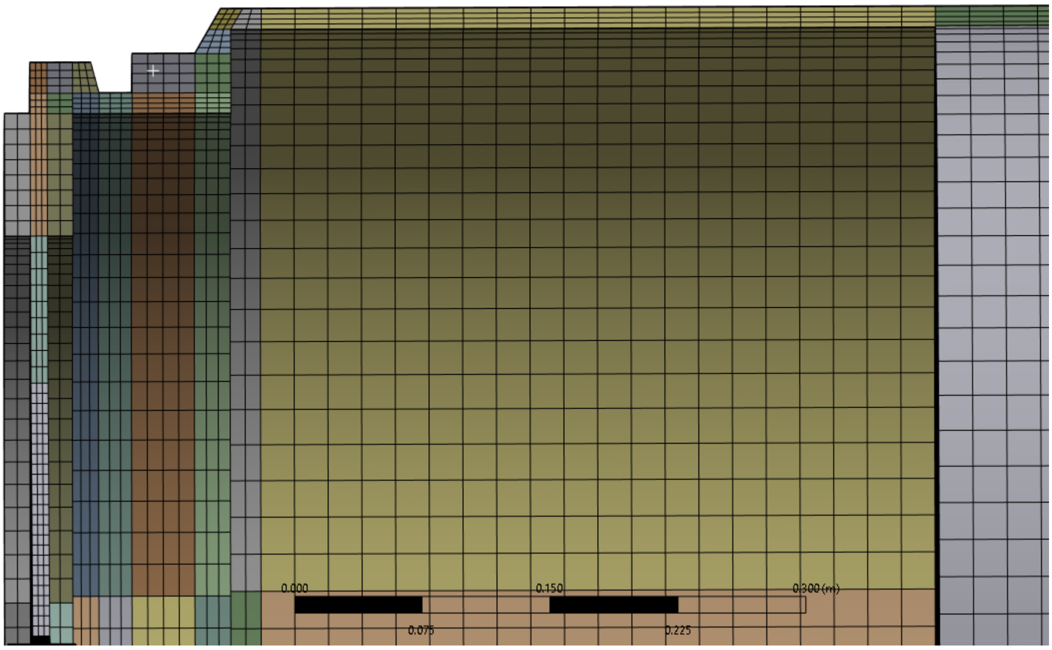

Figure A.4: Cross-sectional view of the upstream region showing the discretisation of the vessel wall with four elements through the thickness.

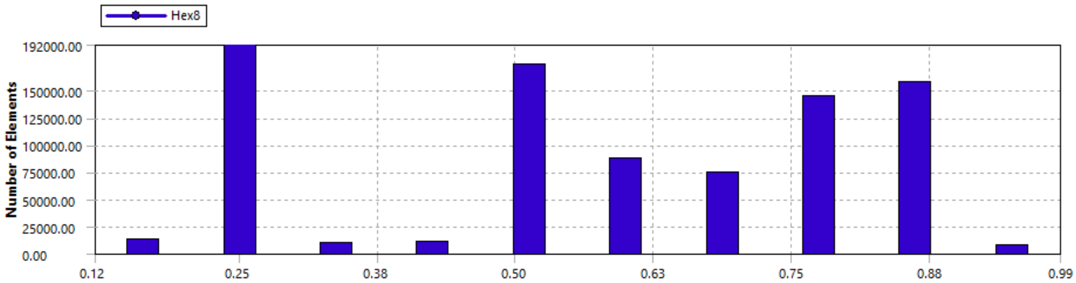

Figure A.5: Histogram of the Ansys element quality metric for the global model.

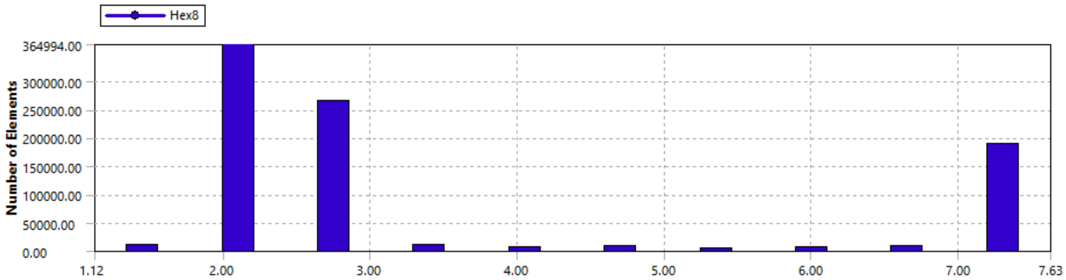

Figure A.6: Histogram of the element aspect ratio for the global model.

Figure A.6 shows a histogram of the element aspect ratio. Lower aspect ratios are observed primarily in the core region where elements are closer to cubic. The highest aspect ratios occur in vessel elements due to their relatively small through-thickness dimension compared to the in-plane element size. The average aspect ratio of the model is 3.46, with a maximum value of 7.6.

A.4 Material Models

This annex provides an overview of the materials used in this thesis. The caption provide the keyword name for those materials implemented in LS-Dyna. The material cards can be found in [8].

A.4.1 Titanium Grade 5 - Ti-6Al-4V (Annealed, AMS 4911)

An overview of the mechanical and thermal material properties is provided in Table A.1 and Table A.2 respectively.

Table A.1: Temperature-dependent mechanical properties of Ti-6Al-4V (`*MAT_ELASTIC_PLASTIC_THERMAL`).

T [°C]	E [GPa]	ν [-]	α [10^{-6}/K]	σ_y [MPa]
20	110.0	0.31	8.76	834
204	97.1	0.31	9.88	591
316	90.5	0.31	10.4	520
538	81.6	0.31	8.13	483
2000	81.6	0.31	8.13	483

Table A.2: Temperature-dependent thermal properties of Ti-6Al-4V (`*MAT_THERMAL_ISOTROPIC_TD_LC`).

T [°C]	c_p [J/kgK]	k [W/mK]
-1000	544	7.2
20	544	7.2
204	574	8.8
316	611	10.4
538	720	11.9
2000	720	11.9

A.4.2 SS318LN

An overview of the mechanical and thermal material properties is provided in Table A.3 and Table A.4 respectively.

Table A.3: Temperature-dependent mechanical properties of SS318LN (`*MAT_ELASTIC_PLASTIC_THERMAL`).

T [°C]	E [GPa]	ν [-]	α [10^{-6}/K]	σ_y [MPa]
20	195.2	0.31	12.44	470
50	194.5	0.31	12.79	421
100	190.8	0.31	13.32	355
150	183.6	0.31	13.79	306
200	174.5	0.31	14.22	276
250	163.1	0.31	14.58	265
2000	163.1	0.31	14.58	265

Table A.4: Temperature-dependent thermal properties of SS318LN (`*MAT_THERMAL_ISOTROPIC_TD_LC`).

T [°C]	c_p [J/kgK]	k [W/mK]
20	481	14.1
100	502	15.7
150	513	16.5
200	523	17.2
250	540	18.1
2000	540	18.1

A.4.3 Carbon Fibre Reinforced Carbon (EDMS 3064824)

The material model for the CFC was derived from testing performed at CERN and documented in EDMS 3064824. Owing to the composite architecture and layup configuration of the CFC, an orthotropic material formulation was adopted. The material is assumed to be transversely isotropic, exhibiting identical properties in the in-plane directions and distinct properties in the out-of-plane direction.

Tables A.5 and A.6 summarise the mechanical properties defined in the three principal material directions. Tables A.7 and A.8 present the corresponding thermal properties.

Table A.5: CFC mechanical model definition (`*MAT_TEMPERATURE_DEPENDENT_ORTHOTROPIC`).

Density ρ	1286 kg/m^3
Elastic model	Orthotropic, temperature dependent
Poisson's ratios	0.20 (constant)
Shear moduli	$G_{12} = 10.9$ GPa, $G_{23} = 4.55$ GPa, $G_{31} = 4.55$ GPa (constant)

Table A.6: Temperature-dependent orthotropic elastic properties and thermal expansion coefficients for CFC (`*MAT_TEMPERATURE_DEPENDENT_ORTHOTROPIC`).

T [°C]	E_1 [GPa]	E_2 [GPa]	E_3 [GPa]	α_1 [10^{-6}/K]	α_2 [10^{-6}/K]	α_3 [10^{-6}/K]
20	24.0	24.0	10.0	-2.17	-2.17	8.39
200	24.0	24.0	10.0	-0.50	-0.50	8.72
400	24.0	24.0	10.0	1.06	1.06	9.02
600	24.0	24.0	10.0	2.29	2.29	9.27
1200	24.0	24.0	10.0	4.14	4.14	9.68
3000	24.0	24.0	10.0	4.14	4.14	9.68

Table A.7: CFC thermal model definition (`*MAT_THERMAL_ORTHOTROPIC_TD`).

Thermal model	Orthotropic, temperature dependent
Density ρ	1286 kg/m^3
Specific heat $c_p(T)$	Temperature dependent (table below)
Conductivity $k_1(T), k_2(T), k_3(T)$	Orthotropic, temperature dependent (table below)
Axis definition	k_1, k_2, k_3 aligned with material axes 1–3

Table A.8: Temperature-dependent thermal properties for CFC (`*MAT_THERMAL_ORTHOTROPIC_TD`).

T [°C]	c_p [J/kgK]	k_1 [W/mK]	k_2 [W/mK]	k_3 [W/mK]
20	655	2.60	2.60	10.0
200	1180	3.05	3.05	10.0
400	1520	3.30	3.30	10.0
600	1730	3.45	3.45	10.0
1000	1900	3.55	3.55	10.0
1400	1970	3.56	3.56	10.0
1800	1995	3.59	3.59	10.0
3000	2000	3.60	3.60	10.0

A.4.4 Isostatic and Extruded Graphite

The material model for isostatic and extruded graphite was derived from testing performed at CERN and documented in EDMS 3064824 in combination with datasheets supplied by suppliers.

Tables A.9 and A.10 summarise the mechanical and thermal properties respectively.

Table A.9: Mechanical properties of isostatic and extruded graphite (`*MAT_ELASTIC`).

Material model	Linear elastic, isotropic
Density ρ	1730 kg/m^3
Young's modulus E	10.0 GPa
Poisson's ratio ν	0.10
Temperature dependence	Not included

Table A.10: Thermal properties of isostatic and extruded graphite (`*MAT_THERMAL_ISOTROPIC`).

Thermal model	Isotropic
Density ρ	1730 kg/m^3
Specific heat c_p	650 J/kgK
Thermal conductivity k	70 W/mK
Temperature dependence	Not included

A.4.5 Flexible Graphite

The material model for flexible graphite was derived from supplier datasheets and available technical documentation. Owing to its function in the assembly and the scope of the present study, the material was modelled as a linear elastic, isotropic solid without temperature dependence in the mechanical formulation.

Tables A.11 and A.12 summarise the mechanical and thermal properties, respectively.

Table A.11: Mechanical properties of flexible graphite (`*MAT_ELASTIC`).

Material model	Linear elastic, isotropic
Density ρ	1200 kg/m^3
Young's modulus E	0.465 GPa
Poisson's ratio ν	0.35
Temperature dependence	Not included (mechanical model)

Table A.12: Thermal properties of flexible graphite (`*MAT_THERMAL_ISOTROPIC`).

Thermal model	Isotropic
Density ρ	1200 kg/m^3
Specific heat c_p	700 J/kgK
Thermal conductivity k	150 W/mK
Temperature dependence	Not included

A.5 Contact Modelling

Table A.13 provides a detailed overview of the implemented contact formulations in the global model. Explicit reference to the keyword used is found in the Keyword column. Frictional values are stated for each of the frictional contacts. More informaiton on the contact formulations can be found in [7].

Table A.13: Summary of contact definitions implemented in the global model.

Interface	**Keyword**	**Type**	**Purpose / Remarks**
Core - Vessel (radial)	`*CONTACT_SURFACE_TO_SURFACE_INTERFERENCE`	Interference (frictional)	Models the shrink-fit configuration. Initial geometric penetration is allowed and contact stiffness is ramped using a load curve until penetration is eliminated. Friction coefficients: $\mu_s = 0.5$, $\mu_d = 0.35$.
Window - Vessel Flange (axial)	`*CONTACT_TIED_SURFACE_TO_SURFACE`	Bonded (bolted)	Represents the bolted connection. Implemented as tied contact to enable extraction of interface forces for sub-modelling of the window-vessel joint.
Vessel - Cable Guides (radial)	`*CONTACT_TIED_SURFACE_TO_SURFACE`	Bonded (welded)	Used to extract interface forces for validation of cable guide fillet welds (assessment outside the scope of this thesis).
Vessel - Retaining Ring (radial)	`*CONTACT_TIED_SURFACE_TO_SURFACE`	Bonded (fixed)	Allows force extraction for submodelling and structural dimensioning of the retaining ring (assessment outside the scope of this thesis).
Extruded graphite - LD sector (axial)	`*CONTACT_SURFACE_TO_SURFACE`	Frictional	Allows relative sliding between extruded graphite and LD sector. Friction coefficients: $\mu_s = 0.5$, $\mu_d = 0.35$.
Retaining ring - extruded graphite (axial)	`*CONTACT_SURFACE_TO_SURFACE`	Frictional	Allows relative sliding between graphite sector and retaining rings. Friction coefficients: $\mu_s = 0.5$, $\mu_d = 0.35$.

A.6 Energy Deposition Interpolation from FLUKA to Ansys

This appendix summarises the interpolation of volumetric energy deposition from the FLUKA model to the Ansys finite element model.

Figure A.7 shows an isometric view of the interpolated power density on the vessel and windows. The highest power density occurs on the top left side of the vessel, close to the longitudinal centre of the dump. This location corresponds to the region where the beam sweep approaches the vessel wall most closely. As a result, a larger fraction of the particle shower is absorbed by the vessel in this area.

The upstream window appears almost uniformly blue, indicating very low energy deposition. At this location the particle shower is not yet fully developed, and the window is nearly transparent to the primary protons.

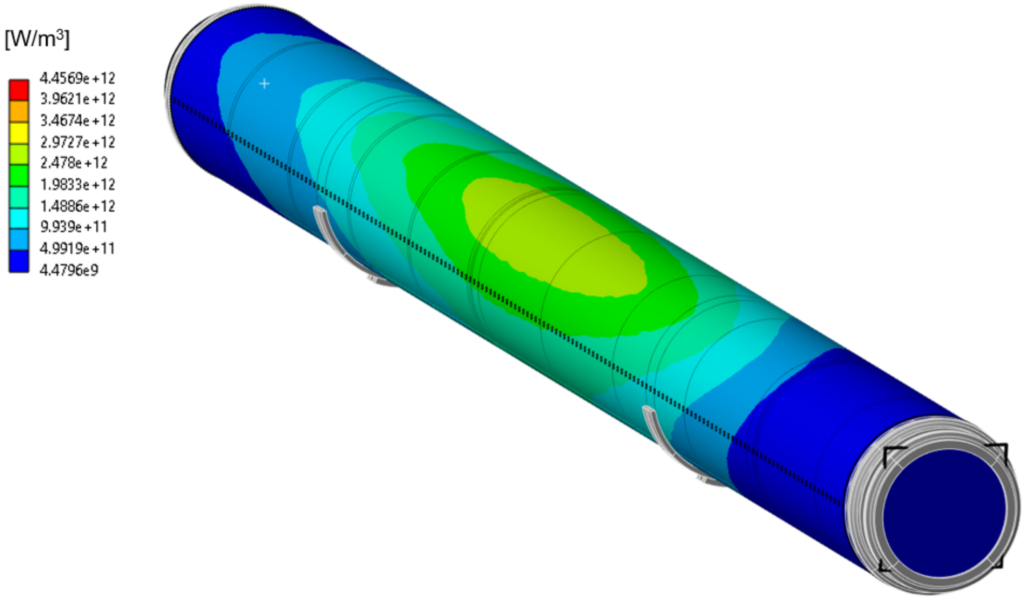

Figure A.7: Isometric view of interpolated volumetric power density from FLUKA mapped onto the Ansys model. The highest pwoer density occurs on the top left side of the vessel where the beam sweep approaches the vessel wall most closely.

Figure A.8 shows a cross-section at the axial position of peak vessel power density. The power density varies with angular position and through the vessel thickness. The maximum value is located near the inner surface, reflecting the radial development of the particle shower.

Figure A.9 compares the axial distribution of peak power density in the vessel and in the core. The vessel peak is located further downstream than the core peak. This shift is expected, as the particle shower develops both longitudinally and radially. Due to the larger vessel radius, the maximum power density in the vessel occurs further downstream.

Figure A.10 shows the interpolated power density on the upstream and downstream windows (Figures A.10a and A.10b). The colour contours differ by more than one order of magnitude.

For the upstream window, the energy distribution appears scattered due to the coarse mesh and the highly localised energy deposition. For the downstream window, the shower is fully developed and the energy deposition is more spatially distributed. The downstream window is shown from the backside, resulting in a horizontal mirroring of the sweep pattern.

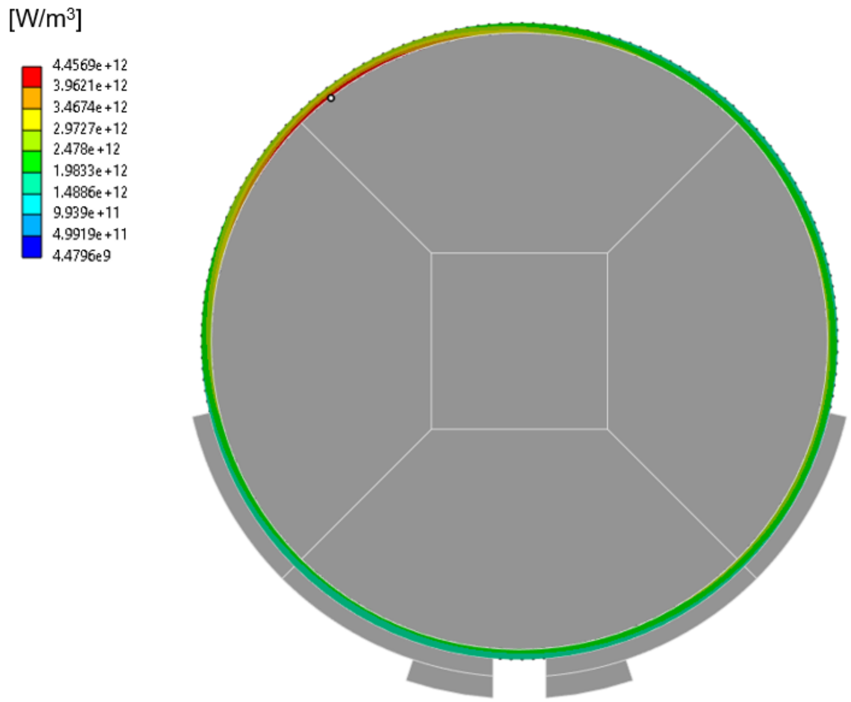

Figure A.8: Cross-sectional view at the axial location of peak vessel power density. The distribution varies with angular position and through the wall thickness, with the maximum located near the top left inner surface.

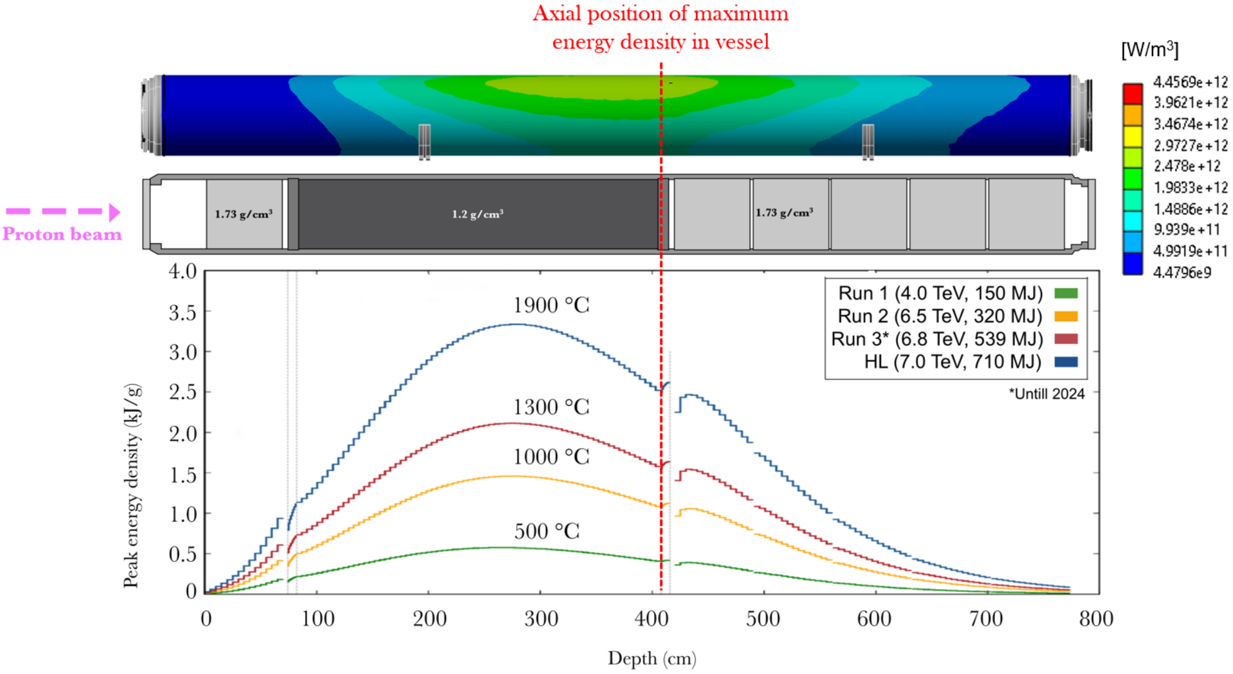

Figure A.9: Axial distribution of peak power density in the vessel and core after interpolation. The vessel peak is located further downstream due to longitudinal and radial shower development.

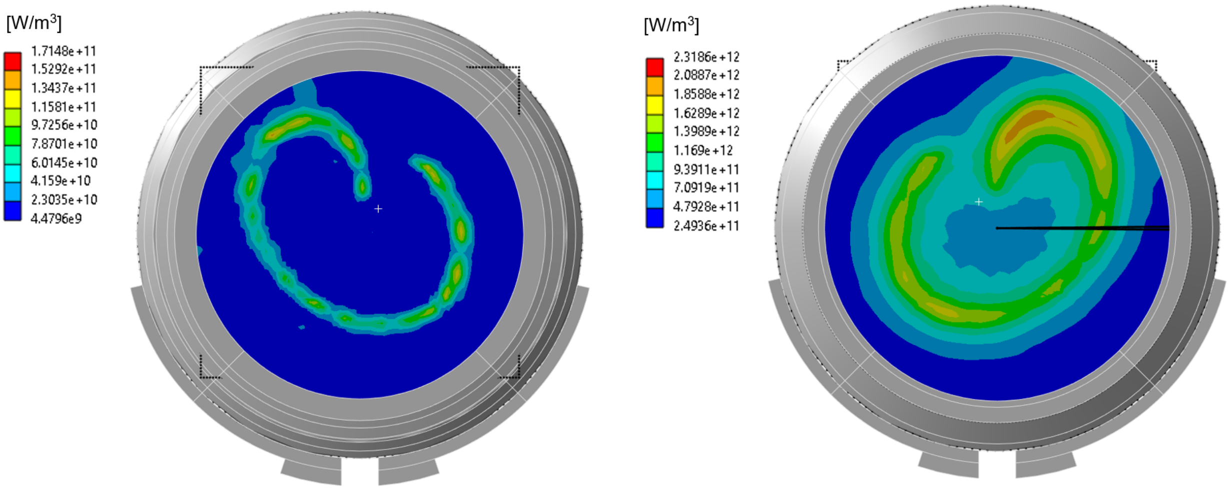

(a) Interpolated power density on the upstream window. Low deposition is observed as the particle shower is not yet developed.

(b) Interpolated power density on the downstream window. Higher and more distributed deposition is observed due to the developed particle shower.

Figure A.10: Comparison of interpolated volumetric power density on upstream and downstream windows. Note that the colour contours differ by more than one order of magnitude.

A.7 Calculations on Interference Fit According to ASME B4.1

This appendix presents the analytical calculations used to evaluate the stresses induced by the interference fit between the isostatic graphite core and the 318LN vessel for the Run 3 beam dumps.

A.7.1 Material Properties and Geometry

The material properties and geometric parameters used in the analysis are summarised in Table A.14.

Table A.14: Material and geometric parameters for the interference fit calculations.

Parameter	Value
Inner material Poisson's ratio, ν_i	0.10
Outer material Poisson's ratio, ν_o	0.30
Inner material elastic modulus, E_i	10 GPa
Outer material elastic modulus, E_o	195 GPa
Inner material inner diameter, $D_{i,i}$	0 mm
Inner material outer diameter, $D_{i,o}$	700 mm
Outer material inner diameter, $D_{o,i}$	698.5 mm
Outer material outer diameter, $D_{o,o}$	722.5 mm
Maximum interference, Δ_i	1.5 mm

The interference between the two components is defined as

$$\Delta_i = D_{i,o} - D_{o,i} \tag{A.1}$$

A.7.2 Interference Pressure

The interference contact pressure P_i generated at the interface due to the interference fit is given by:

$$P_i = \frac{\Delta_i}{\frac{D_{o,i}}{E_o}\left(\frac{D_{o,o}^2+D_{o,i}^2}{D_{o,o}^2-D_{o,i}^2}+\nu_o\right)+\frac{D_{i,o}}{E_i}\left(\frac{D_{i,o}^2+D_{i,i}^2}{D_{i,o}^2-D_{i,i}^2}-\nu_i\right)} \tag{A.2}$$

Substituting the values from Table A.14 yields:

$$P_i = 8.8\text{MPa} \tag{A.3}$$

A.7.3 Stress Components in the Outer Vessel

The circumferential (hoop) stress in the outer vessel at the interface due to the interference pressure is given by

$$\sigma_{\theta,P} = P_i\left(\frac{D_{o,o}^2 + D_{o,i}^2}{D_{o,o}^2 - D_{o,i}^2}\right). \tag{A.4}$$

Using the calculated interface pressure gives

$$\sigma_{\theta,P} = 2.611 \times 10^8 \text{ Pa.} \tag{A.5}$$

The remaining stress components at the interface are:

$$\sigma_r = -P_i, \tag{A.6}$$

$$\sigma_z = 0, \tag{A.7}$$

$$\tau = 0, \tag{A.8}$$

$$\sigma_{\theta,\text{cfg}} = 0, \tag{A.9}$$

where σ_r is the radial stress, σ_z is the axial stress, τ is the shear stress, and $\sigma_{\theta,\text{cfg}}$ represents any additional circumferential stress contributions, which are neglected in this analysis.

A.7.4 Von Mises Stress

The equivalent von Mises stress is calculated using:

$$\sigma_{\text{VM}} = \sqrt{\frac{(\sigma_r - \sigma_{\theta,P} - \sigma_{\theta,\text{cfg}})^2 + (\sigma_{\theta,P} + \sigma_{\theta,\text{cfg}} + \sigma_z)^2 + (\sigma_r - \sigma_z)^2}{2} + 3\tau^2} \tag{A.10}$$

Substituting the stress components yields:

$$\sigma_{\text{VM}} = 2.656 \times 10^8 \text{ Pa} = 265.6 \text{ MPa} \tag{A.11}$$

The calculated von Mises stress represents the equivalent stress induced in the vessel due to the interference fit with the isostatic graphite core.

A.8 Mechanical Drawing of Electron Beam Fatigue Specimen

Figure A.4.1 shows the mechanical drawings for the titanium Grade 5 electron beam welded fatigue specimen.

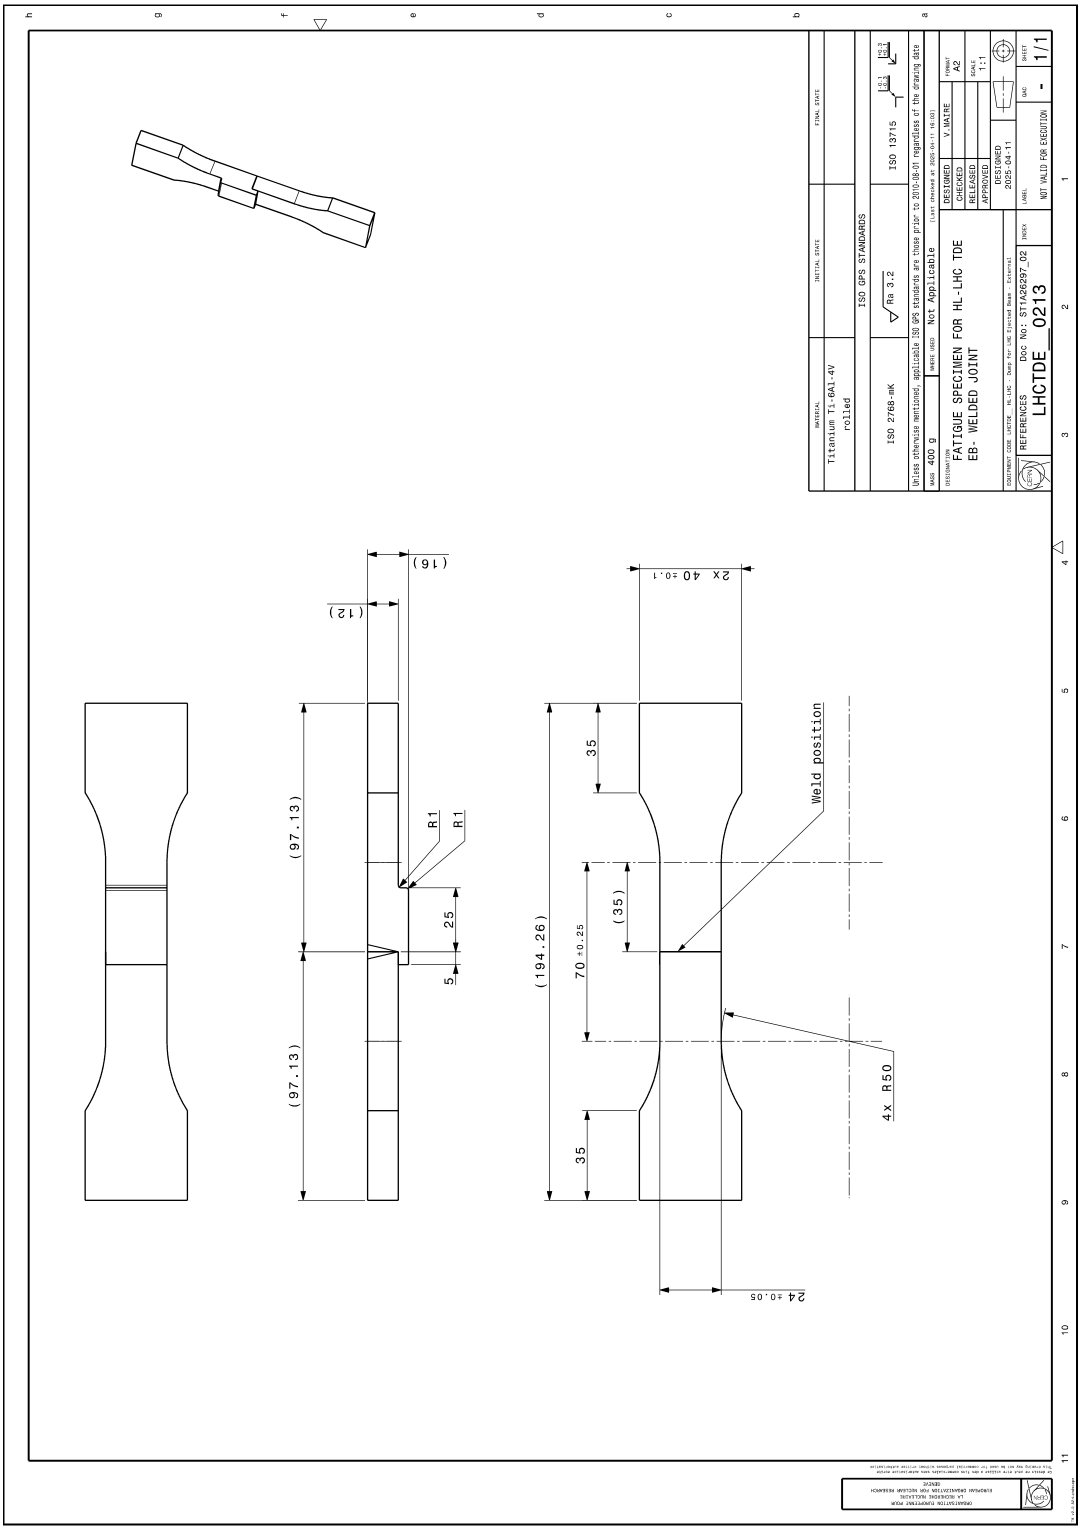

Figure A.11: Mechanical drawing for the titanium Grade 5 electron beam welded fatigue specimen

A.9 Computation of Equivalent Stress and Load Cycles from Damped Extrapolated Simulation Data

This annex describes how the simulation stress histories were processed to obtain an equivalent stress range and the corresponding number of cycles. The procedure combines direct use of the simulated time histories with an idealized damped oscillation model that reproduces the same initial peak stress amplitude.

The stress histories are imported from `.csv` files extracted from the LS-Dyna simulation. For each load case, the first column contains time and the second column contains stress. No extrapolation is applied to the imported simulation data, and no cycle filtering is applied to these simulation signals before rainflow counting.

In addition to the simulated stress history, the script defines a theoretical damped oscillation with the same initial maximum stress amplitude as the simulation. This is not a curve fit in the strict optimization sense. Instead, the response is idealized as a decaying harmonic signal with prescribed natural frequency, initial amplitude, and damping ratio. For both load cases, the natural frequency is taken as 158 Hz which gives the undamped circular frequency:

$$\omega_0 = 2\pi f_n. \tag{A.12}$$

The damped circular frequency is:

$$\omega_d = \omega_0 \sqrt{1-\zeta^2}, \tag{A.13}$$

where ζ is the damping ratio. The theoretical stress history is then written as:

$$\sigma(t) = -A\, e^{-\zeta\omega_0 t} \cos(\omega_d t), \tag{A.14}$$

where A is chosen to match the initial peak stress amplitude of the simulation. $A = 80$ MPa for full intensity beam dumps and $A = 25$ MPa for half intensity beam dumps. Three damping ratios are investigated of $\zeta = 0.0085, 0.00425, 0.0016$. For each damping ratio, the actual logarithmic decrement used in the plots is:

$$\delta = \frac{2\pi\zeta}{\sqrt{1-\zeta^2}}, \tag{A.15}$$

which gives the percentage drop per cycle as:

$$\%\text{drop} = \left(1 - e^{-\delta}\right) 100. \tag{A.16}$$

The fatigue-relevant stress cycles are extracted using MATLAB's `rainflow` function. For each cycle, the function returns the cycle count n_i, the stress range $\Delta\sigma_i$, and the the mean stress $\sigma_{m,i}$.

For the theoretical damped signals, rainflow counting is applied, but only cycles above a prescribed threshold of 10 MPa are retained, assuming that cycles with lower amplitude have no fatigue damage. The total number of counted cycles is obtained by summing the rainflow counts:

$$N = \sum_i n_i. \tag{A.17}$$

The script computes a single equivalent stress range from the set of variable-amplitude cycles using a power-law damage aggregation with exponent $m = 2.8469$ (obtained from fatigue testing campaign). The equivalent stress range is calculated as:

$$\Delta\sigma_{\text{eq}} = \left(\frac{\sum_i n_i \, (\Delta\sigma_i)^m}{\sum_i n_i} \right)^{1/m}. \tag{A.18}$$

This gives one constant-amplitude stress range that produces the same weighted effect as the original variable-amplitude spectrum within the analysed history.

The script is provided below:

```matlab
%% -------------------------------------------------------------------------
% Global time settings
% --------------------------------------------------------------------------
tEnd = 1.5;                 % Total duration [s]
dt   = 1e-4;                % Time step [s]
t    = 0:dt:tEnd;           % Common time vector [s]

%% -------------------------------------------------------------------------
% Rainflow filtering for theoretical signals only
% --------------------------------------------------------------------------
ampMin   = 10;              % Minimum stress amplitude [MPa]
rangeMin = 2 * ampMin;      % Minimum stress range [MPa]

% Keep only rainflow rows with range >= rangeMin
filterRainflow = @(rf) rf(rf(:,2) >= rangeMin, :);

%% -------------------------------------------------------------------------
% Material / fatigue parameters
% --------------------------------------------------------------------------
m         = 2.8469;         % Fatigue exponent
sigmaUlt  = 900;            % Ultimate stress for Goodman correction [MPa]

%% -------------------------------------------------------------------------
% Damping ratios
% --------------------------------------------------------------------------
zetaValues = [0.0085, 0.00425];

% Additional damping ratio corresponding to 1% amplitude drop per cycle
targetDropPct = 1;
deltaTarget   = -log(1 - targetDropPct / 100);
zeta1pct      = deltaTarget / sqrt((2*pi)^2 + deltaTarget^2);

zetaValues = [zetaValues, zeta1pct];

fprintf('Damping ratio for 1%% drop per cycle: %.6f\n', zeta1pct);

%% -------------------------------------------------------------------------
% Case definitions
% --------------------------------------------------------------------------
% Each case contains:
% - label: plot/output label
% - fn: natural frequency [Hz]
% - A: initial stress amplitude [MPa]
% - file: CSV file with simulation results
% - offset: stress offset removed after import [MPa]
% - plotColor: color used for simulation curve
```

```
% - xLimits: x-axis limit for comparison plot [ms]
cases(1).label     = 'Full Intensity';
cases(1).fn        = 140;
cases(1).A         = 80;
cases(1).file      = 'C:/DesignReviewSimulations/Full intensity with CFC geometry/Max_Z_stress.csv';
cases(1).offset    = 78.573312;
cases(1).plotColor = 'k';
cases(1).xLimits   = [0 250];

cases(2).label     = '52% Intensity';
cases(2).fn        = 140;
cases(2).A         = 25;
cases(2).file      = 'D:/LHC Target Dump Externals/High luminosity/Welds/Fatigue resistance of welds
↪  calculations/MaxZStress53%intensity.csv';
cases(2).offset    = 80.8;
cases(2).plotColor = 'g';
cases(2).xLimits   = [0 1500];

%% ---------------------------------------------------------------------
% Histogram settings
% ----------------------------------------------------------------------
maxRange = 200;
numBins  = 30;
binEdges = linspace(0, maxRange, numBins + 1);

%% ---------------------------------------------------------------------
% Main loop over load cases
% ----------------------------------------------------------------------
for c = 1:numel(cases)

    %% Case-specific parameters
    caseName = cases(c).label;
    fn       = cases(c).fn;
    omega0   = 2 * pi * fn;
    A        = cases(c).A;

    %% -----------------------------------------------------------------
    % Import simulation data
    % ------------------------------------------------------------------
    data = readmatrix(cases(c).file);

    timeExp   = data(:,1) * 1000;           % Convert time to [ms]
    stressExp = data(:,2) / 1e6 - cases(c).offset;   % Convert to [MPa] and remove offset

    %% -----------------------------------------------------------------
    % Plot simulation data together with theoretical damped oscillations
    % ------------------------------------------------------------------
    figure; hold on;

    for i = 1:numel(zetaValues)
        zeta   = zetaValues(i);
        omegaD = omega0 * sqrt(1 - zeta^2);

        % Idealized damped oscillation
        stressTh = -A * exp(-zeta * omega0 * t) .* cos(omegaD * t);

        % Logarithmic decrement and percentage amplitude drop per cycle
        delta       = (2 * pi * zeta) / sqrt(1 - zeta^2);
        dropPct     = (1 - exp(-delta)) * 100;

        plot(t * 1000, stressTh, '-', 'LineWidth', 1.2, ...
            'DisplayName', sprintf('\\zeta = %.4f (%.2f%% drop/cycle)', zeta, dropPct));
    end

    plot(timeExp, stressExp, cases(c).plotColor, 'LineWidth', 1.5, ...
        'DisplayName', ['Simulation ', caseName]);

    title(['Damped Oscillations vs. ', caseName, ' Simulation']);
    xlabel('Time [ms]');
    ylabel('Axial stress [MPa]');
    legend('show');
    grid on;
```

```
    xlim(cases(c).xLimits);
    ylim([-110 110]);
    set(gca, 'FontSize', 12);

    %% ------------------------------------------------------------------
    % Rainflow counting for simulation data (unfiltered)
    % -------------------------------------------------------------------
    rfExp    = rainflow(stressExp, timeExp);
    rangeExp = rfExp(:,2);

    figure;
    histogram(rangeExp, binEdges);
    xlabel('Stress Range [MPa]');
    ylabel('Cycle Count');
    title(['Rainflow Histogram - ', caseName, ' Simulation (Unfiltered)']);
    xlim([0 maxRange]);
    grid on;

    %% ------------------------------------------------------------------
    % Equivalent stress for simulation data
    % -------------------------------------------------------------------
    countsExp = rfExp(:,1);
    rangesExp = rfExp(:,2);
    meansExp  = rfExp(:,3);

    cyclesExp      = sum(countsExp);
    rangeEqExp     = (sum(countsExp .* rangesExp.^m) / sum(countsExp))^(1/m);
    meanEqExp      = sum(countsExp .* meansExp) / sum(countsExp);
    rangeEqExpAdj  = rangeEqExp / (1 - meanEqExp / sigmaUlt);

    fprintf('\n=== Equivalent Load Summary: %s Simulation (Unfiltered) ===\n', caseName);
    fprintf('Total number of cycles      : %10.0f\n', cyclesExp);
    fprintf('Equivalent stress range     : %10.2f MPa\n', rangeEqExp);
    fprintf('Equivalent mean stress      : %10.2f MPa\n', meanEqExp);
    fprintf('Goodman-adjusted stress     : %10.2f MPa\n', rangeEqExpAdj);

    %% ------------------------------------------------------------------
    % Theoretical damped oscillations: rainflow + filtering + equivalent load
    % -------------------------------------------------------------------
    for i = 1:numel(zetaValues)
        zeta   = zetaValues(i);
        omegaD = omega0 * sqrt(1 - zeta^2);

        % Idealized damped oscillation
        stressTh = -A * exp(-zeta * omega0 * t) .* cos(omegaD * t);

        % Rainflow counting on theoretical response
        rfThRaw = rainflow(stressTh, t * 1000);

        % Apply minimum range filter only to theoretical signal
        rfTh = filterRainflow(rfThRaw);

        countsTh = rfTh(:,1);
        rangesTh = rfTh(:,2);
        meansTh  = rfTh(:,3);

        delta   = (2 * pi * zeta) / sqrt(1 - zeta^2);
        dropPct = (1 - exp(-delta)) * 100;

        % Histogram of filtered theoretical cycles
        figure;
        histogram(rangesTh, binEdges);
        xlabel('Stress Range [MPa]');
        ylabel('Cycle Count');
        title(sprintf('Rainflow Histogram - %s Theoretical (Amp \\geq %g MPa), \\zeta = %.4f', ...
            caseName, ampMin, zeta));
        xlim([rangeMin maxRange]);
        grid on;

        % Equivalent stress calculation
        cyclesTh = sum(countsTh);

        if cyclesTh > 0
```

```
            rangeEqTh    = (sum(countsTh .* rangesTh.^m) / sum(countsTh))^(1/m);
            meanEqTh     = sum(countsTh .* meansTh) / sum(countsTh);
            rangeEqThAdj = rangeEqTh / (1 - meanEqTh / sigmaUlt);
        else
            rangeEqTh    = NaN;
            meanEqTh     = NaN;
            rangeEqThAdj = NaN;
        end

        fprintf('\n--- Theoretical %s (1.5 s, Amp >= %g MPa), zeta = %.4f (%.2f%% drop/cycle) ---\n', ...
            caseName, ampMin, zeta, dropPct);
        fprintf('Total number of cycles      : %10.0f\n', cyclesTh);
        fprintf('Equivalent stress range     : %10.2f MPa\n', rangeEqTh);
        fprintf('Equivalent mean stress      : %10.2f MPa\n', meanEqTh);
        fprintf('Goodman-adjusted stress     : %10.2f MPa\n', rangeEqThAdj);
    end
end
```

A.10 Statistical Post-Processing of Fatigue Results According to ISO 12107:2012

This appendix provides a detailed description of the post-processing of the fatigue test results.

Figure A.12 shows the results of the fatigue campaign in number of cycles reached at different load levels. Fractures of the specimens are indicated by red crosses. Note that the load level corresponds to the stress amplitude for fully reversed cycles (e.g. the 80 MPa load level corresponds to an oscillation with a stress range of 160 MPa, from -80 to 80 MPa).

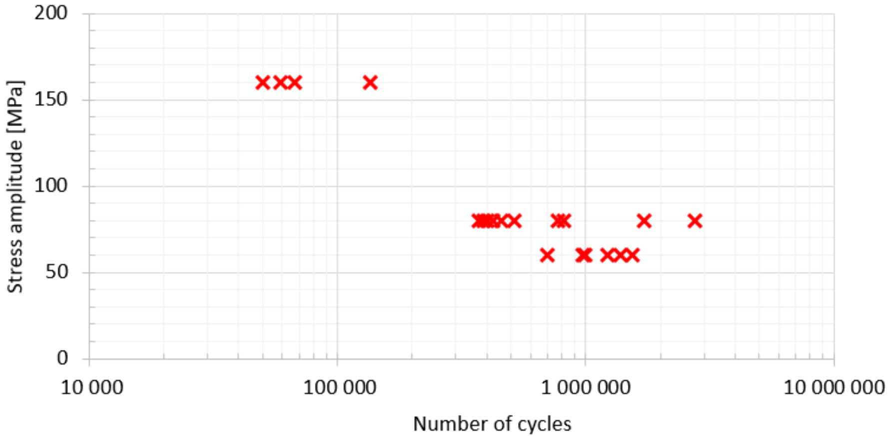

Figure A.12: Fatigue test results for notched titanium Grade 5 specimens welded using electron beam technology, showing number of cycles at given applied stress amplitudes.

The fatigue data were post-processed in accordance with ISO 12107:2012 [10]. It is well established that fatigue life values at a given stress level are approximately normally distributed when expressed on a logarithmic scale. In other words, the logarithms of the fatigue lives follow a normal distribution [10]. Under this assumption, the data points in logarithmic space should align along a straight line, enabling the use of linear regression to derive the S-N curve:

$$\hat{Y}_i = b_0 + b_1 X_i \tag{A.19}$$

where $\hat{Y}_i$ is the predicted number of cycles according to the linear model and:

$$X_i = \log_{10}(S_i), \qquad Y_i = \log_{10}(N_i) \tag{A.20}$$

Here S is the stress level and N the number of cycles.

When the fatigue data are plotted in logarithmic space as shown in Figure A.13, the fitted linear regression yields parameters $b_0 = 11.178$ and $b_1 = -2.8469$. The coefficient of determination $R^2 = 0.7711$ indicates that the model explains approximately 77% of the observed variation in the data. Evaluation of the quality of fit is undertaken by evaluating plots of residuals and plots of residual versus cumulative normal probability. A residual is defined as:

$$e_i = Y_i - \hat{Y}_i \tag{A.21}$$

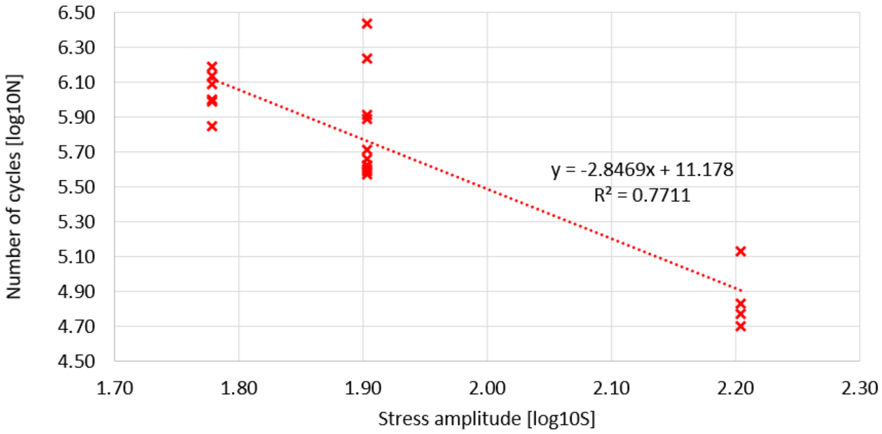

Figure A.13: Fatigue data plotted in logarithmic metrics with the axis swapped to perform linear regression with stress as the independent variable and number of cycles the dependent variable.

This is essentially the difference between the number of cycles estimated by the linear regression model and the fatigue data at respective load levels. One property of the residuals is that they should sum to zero, hence a plot of the residuals versus the corresponding predicted $\hat{Y}_i$ values should uniformly populate the plot without biases or trends.

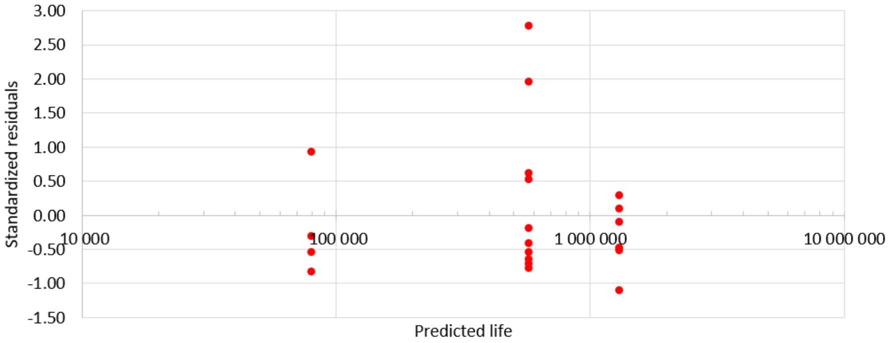

Figure A.14: Residual plot of fatigue data with standardized residuals (residuals scaled by the standard deviation) versus the predicted life.

Figure A.14 shows standardized residuals (residuals normalized by the standard deviation) plotted against the predicted fatigue life from the linear regression model. A few outliers can be observed in the upper region of the second column, corresponding to specimens with significantly higher fatigue life than predicted. Apart from these outliers, the residuals appear to be evenly distributed around zero without any pronounced trends or systematic biases, suggesting that the linear model provides an adequate fit for most of the data.

Plots of the standardized residuals versus cumulative normal probability are also useful in evaluation of the fit of the model to the data. The residuals are assumed to be log-normally distributed, and this assumption can be evaluated by determining if the residuals from the analysis plot reasonably as a straight line. As seen in Figure A.15, this seems to be the case for the data.

Therefore, in conclusion, the data satisfies the assumption of being log-normal distributed and the data processing can be conducted under these assumptions for modifications to the S-N curve when considering confidence and probability of failure levels.

The tolerance coefficient $k_{(P,1-\alpha,v)}$ is found in Figure A.16.

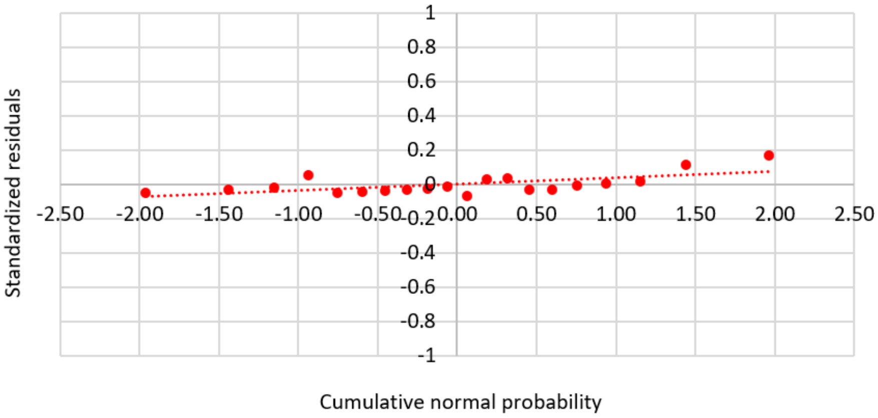

Figure A.15: Cumulative normal probability plot displaying acceptable conformance to normality.

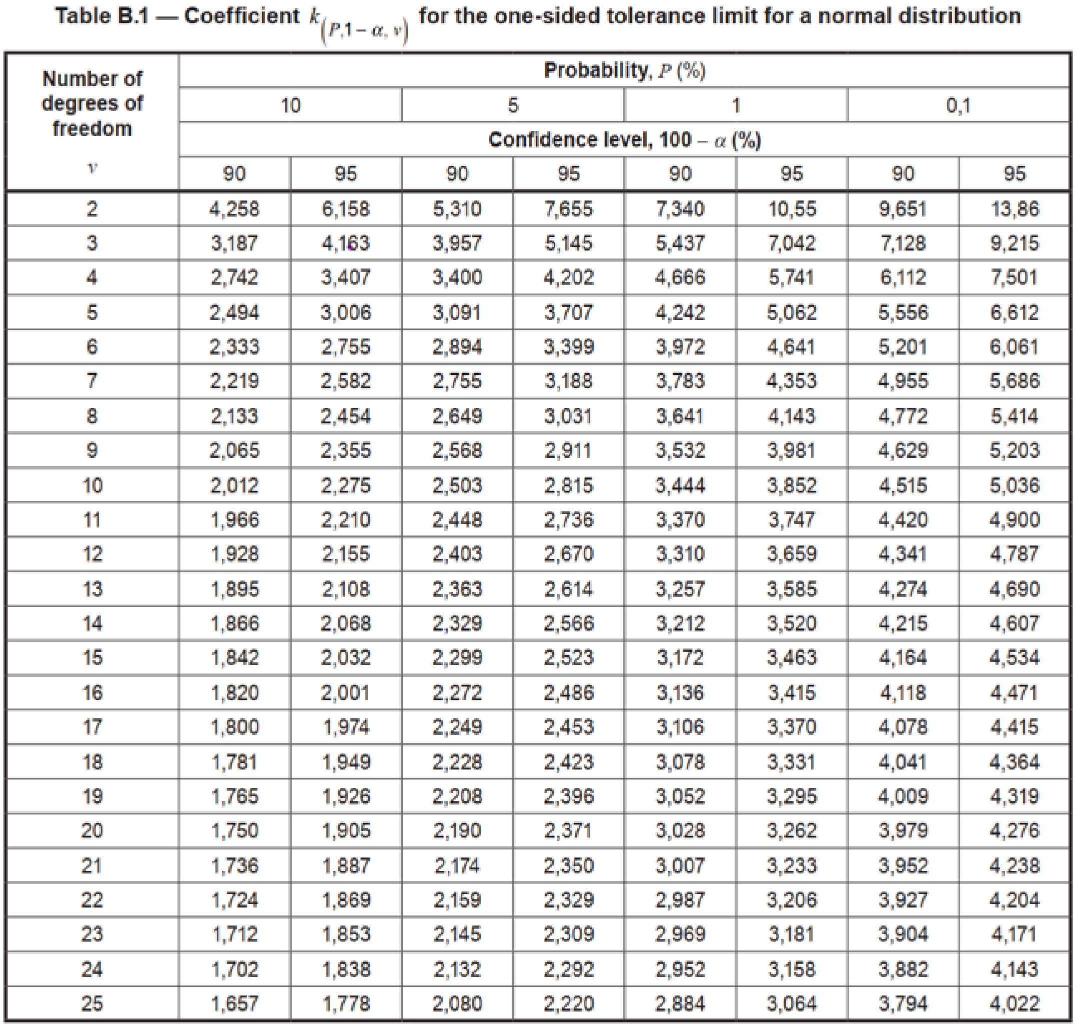

Table B.1 — Coefficient $k_{(P,1-\alpha,\nu)}$ for the one-sided tolerance limit for a normal distribution

Number of degrees of freedom ν	Probability, P (%)							
	10		5		1		0,1	
	Confidence level, 100 – α (%)							
	90	95	90	95	90	95	90	95
2	4,258	6,158	5,310	7,655	7,340	10,55	9,651	13,86
3	3,187	4,163	3,957	5,145	5,437	7,042	7,128	9,215
4	2,742	3,407	3,400	4,202	4,666	5,741	6,112	7,501
5	2,494	3,006	3,091	3,707	4,242	5,062	5,556	6,612
6	2,333	2,755	2,894	3,399	3,972	4,641	5,201	6,061
7	2,219	2,582	2,755	3,188	3,783	4,353	4,955	5,686
8	2,133	2,454	2,649	3,031	3,641	4,143	4,772	5,414
9	2,065	2,355	2,568	2,911	3,532	3,981	4,629	5,203
10	2,012	2,275	2,503	2,815	3,444	3,852	4,515	5,036
11	1,966	2,210	2,448	2,736	3,370	3,747	4,420	4,900
12	1,928	2,155	2,403	2,670	3,310	3,659	4,341	4,787
13	1,895	2,108	2,363	2,614	3,257	3,585	4,274	4,690
14	1,866	2,068	2,329	2,566	3,212	3,520	4,215	4,607
15	1,842	2,032	2,299	2,523	3,172	3,463	4,164	4,534
16	1,820	2,001	2,272	2,486	3,136	3,415	4,118	4,471
17	1,800	1,974	2,249	2,453	3,106	3,370	4,078	4,415
18	1,781	1,949	2,228	2,423	3,078	3,331	4,041	4,364
19	1,765	1,926	2,208	2,396	3,052	3,295	4,009	4,319
20	1,750	1,905	2,190	2,371	3,028	3,262	3,979	4,276
21	1,736	1,887	2,174	2,350	3,007	3,233	3,952	4,238
22	1,724	1,869	2,159	2,329	2,987	3,206	3,927	4,204
23	1,712	1,853	2,145	2,309	2,969	3,181	3,904	4,171
24	1,702	1,838	2,132	2,292	2,952	3,158	3,882	4,143
25	1,657	1,778	2,080	2,220	2,884	3,064	3,794	4,022

Figure A.16: Tolerance factors for different confidence intervals and levels of probability.

A.11 Window Energy Deposition Interpolation Scripts

To perform the spatial and temporal energy deposition interpolation for the window dynamic sweep models three python scripts were used. This appendix aims at providing the scripts and explaining how they work.

A.11.1 Parameter File

The parameter file sets all of the settings for the FLUKA interpolation script. It inputs the FLUKA file, sets beam parameters, imports the mesh from the FE-model and sets the output directory of the file containing the energy deposition load curves that should be submitted to the solver together with the input file. In the following comments are incorporated in the script.

```
# FLUKA file contains the averaged energy deposition in 10 mm thick titanium grade 5 which must be scaled to the
↪  number of protons in each bunch. Bunch length of 25 ns is the time duration for the energy deposition to be
↪  deposited in the material for the given bunch.
# Scaling of Energy Deposition from GeV/cm^3/p+ to J/m^3:
numberOfProtons        = 2.3E11
protonCharge           = 1E6*1E9*1.602176487e-19
bunchLength            = 25E-9
scaleFactor            = numberOfProtons*protonCharge/bunchLength

# Definition of Directories:


# Input and Output for LS-Dyna files:
KFileName              = "DownstreamNominal/DownstreamNominal.k" # LS Dyna input file containing the mesh.
outputDirectory        = 'DownstreamNominal/' # Output directory for the generated load curves.

# Directory of FLUKA file:
flukaDirectory         = 'DownstreamNominal/'
flukaFileName          = flukaDirectory + "Downstream.txt"

# Directory of Sweep patterns:
patternDirectory       = 'SweepPatternsProcessed/'
coordinateFileName     = patternDirectory + "Nominal.txt" # Provides the bunch impact locations as function of
↪  time.

# Settings for Interpolation:
zOffet                 = 0 # If there is any difference betweent he origin of the FLUKA file in the z-direction
↪  and the developed mesh
# partNumber           = [61, 82, 88, 92, 93, 94, 95, 96, 97, 98, 99, 100, 101, 102, 103, 104, 105, 106] # All of
↪  the part numbers from LS-Dyna in which energy deposition interpolation should be performed.
ElementType            = 'SOLID'
```

A.11.2 FLUKA Interpolation

This is the main script in which the data from the parameter file is being processed, the energy interpolation performed and the outputs generated. AI was used as a coding support to develop the script. An overview and description of the main functions of the script is provided in A.15.

Table A.15: Overview of the Python functions used for FLUKA-to-LS-DYNA energy deposition interpolation and load generation.

Function	**Description**
`main()`	Main execution routine. Calls the load generation routine
`generateLoadApplicationFile(...)`	Controls the overall workflow. Opens the required files, reads the FLUKA energy deposition data, generates the time vector, extracts mesh information from the LS-DYNA `k`-file, reads the bunch sweep coordinates, interpolates the energy deposition onto the mesh, writes the load definitions, and closes all files.
`generateTimeVector(...)`	Builds the time vector used for the load curves from the bunch timing information in the coordinate file.
`getDataFromFlukaFile(...)`	Reads the FLUKA input file, generates the FLUKA coordinate axes, applies the scaling factor, and returns the processed FLUKA data.
`processHeaderData(...)`	Extracts header information from the FLUKA file and converts it into a usable format containing the coordinate limits, number of bins, and bin spacing for each direction.
`makeFlukaCoordinateAxes(...)`	Generates the spatial coordinate axes of the FLUKA bins from the header data.
`getMeshInformationArray(...)`	Extracts element numbers and centroid coordinates from the LS-DYNA mesh.
`getElementCoordinates(...)`	Reads the `*NODE` and `*ELEMENT` sections of the LS-DYNA `k`-file, identifies elements belonging to the selected part, and computes their centroid coordinates from the nodal positions.
`makeLineSolidElemenIterator(...)`	Reads solid element definitions from the LS-DYNA `k`-file and returns the element connectivity data. The function combines the two lines used to define each solid element.
`makeLineIterator(...)`	Converts each line of a file block into typed entries. It is defined for generic parsing, although it is not central to the main workflow.
`chunks(...)`	Splits a string into fixed-width fields. This is required because LS-DYNA keyword files use fixed-column formatting.
`giveLineOfBlock(...)`	Yields lines from the current file block until a termination marker is encountered, e.g. a new keyword or comment line.
`getSweepCoordinates(...)`	Reads the bunch impact coordinates from the CSV file. For 3D analyses, the prescribed out-of-plane offset is appended to each coordinate set.
`interpolateOnMeshAndWriteFile(...)`	Core interpolation routine. Creates the `RegularGridInterpolator`, evaluates the FLUKA energy deposition at each element centroid for all bunch positions, writes the corresponding LS-DYNA load curves and heat generation cards.
`iterateOverFiniteElements(...)`	Evaluates the interpolated energy deposition field for all finite elements at each bunch position and yields the time history of deposited energy for each element.
`writeLoadCurve(...)`	Writes the LS-DYNA `*DEFINE_CURVE` and `*LOAD_HEAT_GENERATION` entries for one element. It also integrates the time-dependent heat generation curve to compute the total deposited energy in that element.

```python
import numpy as np
import time
from scipy.interpolate import RegularGridInterpolator
from Parameterfile import *
import PostProcessCoords

def main():
    PostProcessCoords.main(coordinateFileName)
    csvCoordinateFileName = coordinateFileName.split('\\')[-1].split('.')[0] + '.csv'
```

```
    print(csvCoordinateFileName)
    fileNames=[csvCoordinateFileName, flukaFileName, KFileName, csvCoordinateFileName, outputKFileName]
    IOTypes=['r','r','r','r','w']
    fileAccessInfo = (i for i in zip(fileNames,IOTypes))
    maximumDepositedEnergy =
    ↪  generateLoadApplicationFile(fileAccessInfo,scaleFactor,thresholdFactor,zOffet,ElementType,partNumber)
    printMaximumEnergyDeposition(maximumDepositedEnergy)
    printDuration(int(time.process_time()))

def generateLoadApplicationFile(fileAccessInfo,scaleFactor,thresholdFactor,zOffet,ElementType,partNumber):
    maximumDepositedEnergy=[0,0]
    fileList = [open(fileName,IOType) for (fileName,IOType) in fileAccessInfo]
    print('reading FLUKA ...')
    flukaCoordinateAxes, flukaDataArray, dimension =
    ↪  getDataFromFlukaFile(fileList[1],scaleFactor,thresholdFactor)
    timeVector = generateTimeVector(fileList[0],endtime = 1)
    print('reading KFile ...')
    elementNumbers, meshInformationArray = getMeshInformationArray(fileList[2],partNumber,dimension)
    coordinates = getSweepCoordinates(fileList[3],zOffet,dimension)
    print('interpolate ...')
    maximumEnergyDepostion = interpolateOnMeshAndWriteFile(fileList[4],flukaCoordinateAxes,flukaDataArray,eleme⌋
    ↪  ntNumbers,meshInformationArray,timeVector,coordinates)
    for file in fileList:
        file.close()
    return maximumEnergyDepostion


def generateTimeVector(coordinateFile,endtime):
    next(coordinateFile)
    times = (float(line.strip().split(',')[1]) for line in coordinateFile)
    firstEntry = next(times)
    rampTime = firstEntry/10000
    offsets = [0,rampTime]
    timeVector = [0,rampTime,firstEntry,firstEntry+rampTime]
    return timeVector + [time+offset for time in times for offset in offsets] + [endtime]

def getDataFromFlukaFile(flukaFile,scaleFactor,thresholdFactor):
    headerRaw = (next(flukaFile).strip().split() for x in range(8))
    headerData = processHeaderData(headerRaw)
    xmax=int(headerData[0][2])
    ymax=int(headerData[1][2])
    zmax=int(headerData[2][2])

    lines = [line.strip().split() for line in flukaFile]
    energyDepositionDensities=(float(energyDepositionDensityString) for line in lines[:lines.index([])] for
    ↪  energyDepositionDensityString in line)

    flukaDataArray=np.zeros((xmax,ymax,zmax),dtype=np.float64)
    for z in range(zmax):
         for y in range(ymax):
            for x in range(xmax):
                flukaDataArray[x,y,z]=next(energyDepositionDensities)
    threshold = float(np.amax(flukaDataArray)) * thresholdFactor
    flukaDataArray[flukaDataArray < threshold] = 0
    dimension = '3D'
    if flukaDataArray.shape[2]==1:
        dimension = '2D'
        flukaDataArray=np.squeeze(flukaDataArray,axis=(2,))
    flukaCoordinateAxes = makeFlukaCoordinateAxes(headerData,dimension)
    flukaDataArray = scaleFactor * flukaDataArray
    return flukaCoordinateAxes, flukaDataArray, dimension

def processHeaderData(headerRaw):
    tempHead=[]
    for line in headerRaw:
        tempLine=[]
        for value in line:
            try:
                tempLine.append(float(value))
            except Exception:
                continue
        tempHead.append(tempLine)
```

```python
    headerData=tempHead[2:5]
    for index,line in enumerate(headerData):
        headerData[index]=[entry/100. for entry in line]
        headerData[index][2]=int(line[2])
    return headerData


def makeFlukaCoordinateAxes(headerData,dimension):
    flukaCoordinateAxes=[]
    for i in range(3):
        flukaCoordinateAxes.append(np.linspace(headerData[i][0]+0.5*headerData[i][3],headerData[i][1]-0.5*heade
        ↪  rData[i][3],num=headerData[i][2]))
    if dimension == '2D':
        filteredFlukaCoordinateAxes=(flukaCoordinateAxes[0], flukaCoordinateAxes[1])
    else:
        filteredFlukaCoordinateAxes=(flukaCoordinateAxes[0], flukaCoordinateAxes[1], flukaCoordinateAxes[2])
    return filteredFlukaCoordinateAxes


def getMeshInformationArray(KFile,partNumber,dimension):
    elementNumbers,meshInformationArray = getElementCoordinates(KFile,partNumber)
    if dimension == '2D':
        meshInformationArray=meshInformationArray[:,:2]
    return elementNumbers,meshInformationArray


def getElementCoordinates(KFile,partNumber):
    locstr = str
    locfloat = float
    locint = int
    goTo(KFile,"*NODE")
    nodeDict=makeLineDict(KFile,16,0,3,locfloat)
    goTo(KFile,"*ELEMENT")
    elementNodesIterator = makeLineSolidElemenIterator(KFile,8,0,-1,locint)
    elementCentroids = []
    elementNumbers = []
    for finiteElement in elementNodesIterator:
        if finiteElement[1] in partNumber:
            elementNumbers.append(finiteElement[0])
            elementCentroids.append(np.mean(np.array([nodeDict[locstr(locint(i))] for i in
            ↪  finiteElement[2:]]),axis=0))
    return elementNumbers,np.array(elementCentroids)

def goTo(file,signal):
    for line in giveLine(file):
        if signal in line:
            return

def giveLine(file):
    for line in file:
        yield line

def makeLineDict(KFile,chunkSize,start,end,variableType):
    dictionary={}
    for line in giveLineOfBlock(KFile,['$','*']):
        dictionary[line[:8].strip()]=[element for element in chunks(line[8:],chunkSize)]
    for key,value in dictionary.items():
        dictionary[key]=[variableType(element) for element in value[start:end]]
    return dictionary

def makeLineSolidElemenIterator(KFile,chunkSize,start,end,variableType):
    ElementVector=[]
    gen=giveLineOfBlock(KFile,['$','*'])
    for line in gen:
        ElementType=[]
        for subelemnt in ([element for element in chunks(line,chunkSize)][:-1]):
            ElementType.append(variableType(subelemnt))
            if (ElementType[0]==1):
                aa=1
        line=next(gen)
        for subelemnt in ([element for element in chunks(line,chunkSize)][:-1]):
            ElementType.append(variableType(subelemnt))
```

```python
        ElementVector.append(ElementType)
    return ElementVector

def makeLineIterator(KFile,chunkSize,start,end,variableType):
    return (([variableType(entry) for entry in [element for element in chunks(line,chunkSize)][:-1]]) for line in giveLineOfBlock(KFile,['$','*']))

def chunks(s, n):
    """Produce `n`-character chunks from `s`."""
    for start in range(0, len(s), n):
        yield s[start:start+n]

def giveLineOfBlock(file,signal):
    for line in file:
        if signal[0] in line:
            return
        elif signal[1] in line:
            return
        yield line

def getSweepCoordinates(coordinateFile,zOffet,dimension):
    next(coordinateFile)
    coordinates=(line.strip().split(',')[2:] for line in coordinateFile)
    if dimension == '3D':
        return [np.array([float(value) for value in line]+[zOffet]) for line in coordinates]
    else:
        return [np.array([float(value) for value in line]) for line in coordinates]

def interpolateOnMeshAndWriteFile(outputKFile,flukaCoordinateAxes,flukaDataArray,elementNumbers,meshInformationArray,timeVector,coordinates):
    maximumDepositedEnergy = [0,0]
    outputKFile.write('$\n$\n$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$\n')
    outputKFile.write('$                               LOAD DEFINITIONS                               $\n')
    outputKFile.write('$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$$\n$\n')

    interpolationFunction = RegularGridInterpolator(flukaCoordinateAxes, flukaDataArray, bounds_error=False, fill_value=0.)

    previousElementNumber = 0
    sections = int(meshInformationArray.shape[0] / 300000)
    if sections == 0: sections = 1
    for i,array in enumerate(iter(np.array_split(meshInformationArray, sections, axis=0))):
        #print(i)
        for localFiniteElementNumber, heatGenerationCurve in enumerate(iterateOverFiniteElements(array, coordinates, interpolationFunction), start=0):
            finiteElementNumber = elementNumbers[localFiniteElementNumber+previousElementNumber]
            heatGenerationCurve = np.array(heatGenerationCurve)
            heatGenerationCurve[heatGenerationCurve < 0] = 0
            if np.amax(heatGenerationCurve) > 0:
                depositedEnergy = writeLoadCurve(outputKFile, [ElementType, heatGenerationCurve, finiteElementNumber, timeVector])
                if depositedEnergy > maximumDepositedEnergy[1]:
                    maximumDepositedEnergy = [finiteElementNumber,depositedEnergy]
        previousElementNumber = previousElementNumber + array.shape[0]
    return maximumDepositedEnergy

def iterateOverFiniteElements(meshInformationArray, coordinates, interpolationFunction):
    meshDataIterator = (np.nditer(interpolationFunction(np.subtract(meshInformationArray,coordinateSystem))) for coordinateSystem in coordinates)
    for heatGenerationCurve in zip(*meshDataIterator):
        yield heatGenerationCurve

def writeLoadCurve(outputKFile, parameters):
    ElementType, heatGenerationCurve, finiteElementNumber, timeVector = parameters
    heatGenerationCurve = [0] + [element for element in heatGenerationCurve for _ in range(2)] + [0, 0]
    outputKFile.write('*DEFINE_CURVE\n% 10.0f% 10.0f% 10.0f% 10.0f% 10.0f% 10.0f' % (finiteElementNumber,0,1,1,0,0))
    energyDeposition = []
    for BunchNumber,time in enumerate(timeVector):
        try:
            cond1 = heatGenerationCurve[BunchNumber-1]>0
```

```
            cond2 = heatGenerationCurve[BunchNumber+1]>0
            cond3 = time == 0
            if cond1 or cond2 or cond3:
                outputKFile.write('\n%20.7E%20.6E' % (time, heatGenerationCurve[BunchNumber]))
                energyDeposition.append([time, heatGenerationCurve[BunchNumber]])
        except IndexError:
            outputKFile.write('\n%20.7E%20.6E' % (time, heatGenerationCurve[BunchNumber]))
            energyDeposition.append([time, heatGenerationCurve[BunchNumber]])
    outputKFile.write('\n*LOAD_HEAT_GENERATION_'+ElementType+'\n% 10.0f% 10.0f% 10.0f\n' % (finiteElementNumber,
    ↪   finiteElementNumber,1))
    energy = 0
    for index,line in enumerate(energyDeposition[1:],start = 1):
        energy = energy+((line[1]+energyDeposition[index-1][1])/2)*(line[0]-energyDeposition[index-1][0])
    return energy

def printMaximumEnergyDeposition(maximumDepositedEnergy):
    print("Maximum deposited Energy in Element "+ str(maximumDepositedEnergy[0]) +":\n" +
    ↪   str(maximumDepositedEnergy[1]/1000000) + " J/cm3.")
    with open('EnergyDeposition.txt','w') as file:
        file.write("\nMaximum deposited Energy in Element "+ str(maximumDepositedEnergy[0]) +":\n" +
        ↪   str(maximumDepositedEnergy[1]/1000000) + " J/cm3.")

def printDuration(durationInSeconds, StringFormat = '\nIt took %i:%i:%i to finish.'):
    RestTime = durationInSeconds
    Hours = int(durationInSeconds / 3600)
    RestTime = durationInSeconds - Hours * 3600
    Minutes = int(RestTime / 60)
    RestTime = RestTime - Minutes * 60
    Seconds = int(RestTime)
    print(StringFormat % (Hours , Minutes , Seconds))

if __name__ == '__main__':
    main()
```

A.11.3 Plotter

The plotter script is used to visualise the interpolated energy deposition and verify that the resulting energy densities are consistent with the FLUKA predictions. This verification step ensures that the spatial and temporal interpolation has been performed correctly before submitting the LS-Dyna simulation to the computing cluster. Figure A.17 shows the beam sweep pattern obtained from the script.

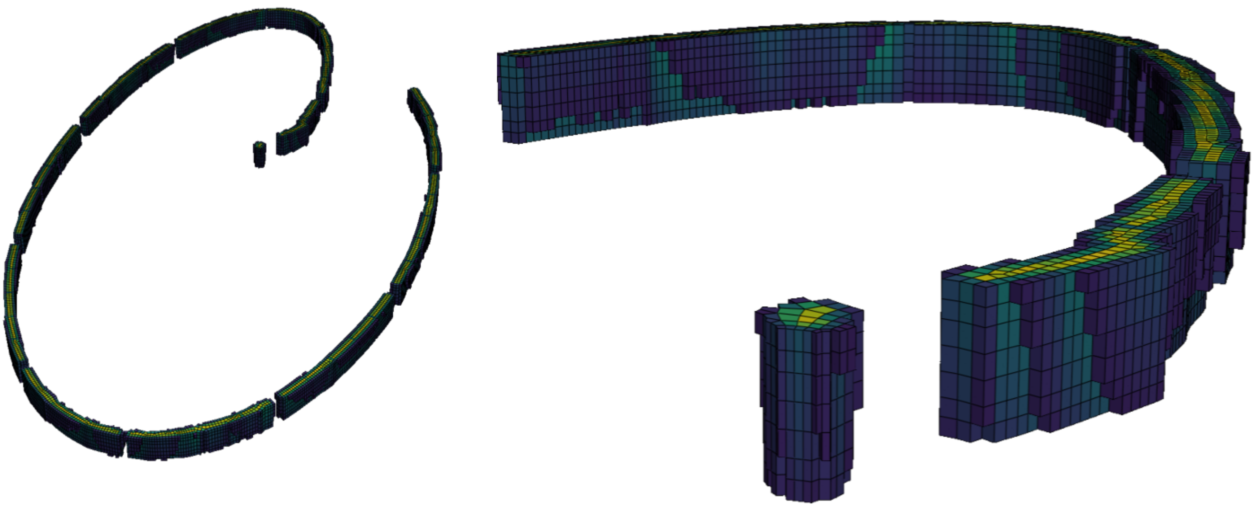

Figure A.17: Visualisation of accumulated energy deposition in the beam sweep for the 2MKBH dilution failure scenario, generated with the plotter script to verify the FLUKA energy deposition interpolation.

```
import os
import re
```

```python
import time
from pathlib import Path
import pickle

from matplotlib import pyplot as plt
import plotly.express as px
import plotly.io as pio
import numpy as np
import pandas as pd
from tqdm import tqdm
from scipy import integrate
import pyvista as pv
from pyvista import examples, CellType

pio.renderers.default = "browser"

file_path = "C:/V9Windows/UpstreamNominal/"
mesh_path = (file_path + "UpstreamNominal.k")
load_curve_path = (file_path + "UpstreamNominal_loaded.k")

def time_convert(seconds):
    minutes = seconds // 60
    seconds = seconds % 60
    hours = minutes // 60
    minutes = minutes % 60
    print(f"Processing time = {int(hours)}:{int(minutes)}:{seconds:.0f} \n")

def line_count(file_path):
    with open(file_path, "r") as f:
        return sum(1 for _ in f)

def read_mesh(path):
    geometry_start_time = time.time()
    with open(path, encoding="utf8") as file:
        lines = file.readlines()
        start_element = 0

        print("Reading input file lines...")
        print("Identifying data for extraction...")

        start_node_found = False
        start_element_found = False
        end_node_found = False
        end_element_found = False

        for line_number, line in enumerate(lines):
            if "*NODE" in line:
                start_node = line_number + 1
                start_node_found = True
                continue

            if not re.match("^[0-9 .-]+$", line.strip()) and start_node_found and not end_node_found:
                end_node = line_number  # maybe subtract 1
                end_node_found = True

            if "*ELEMENT_SOLID" in line:
                start_element = line_number + 1
                start_element_found = True
                continue

            if not re.match("^[0-9 ]+$", line.strip()) and start_element_found and not end_element_found:
                end_element = line_number
                end_element_found = True

            if end_element_found and end_node_found:
                break

        print("Extracting elements, parts and coordinates...")

        ids = []
        xs = []
        ys = []
        zs = []
```

```python
        for line_number, line in enumerate(lines[start_node:end_node]):
            split_line = line.split()
            ids.append(int(split_line[0]))
            xs.append(float(split_line[1]))
            ys.append(float(split_line[2]))
            zs.append(float(split_line[3]))

        element_ids = []
        part_ids = []
        node_ids = []
        for line_number, line in enumerate(lines[start_element:end_element]):
            if line_number % 2 == 0:
                split_line = line.split()
                element_ids.append(int(split_line[0]))
                part_ids.append(int(split_line[1]))
            else:
                split_line = line.split()
                node_ids.append([int(value) for value in split_line])

    xs, ys, zs, ids = np.array(xs), np.array(ys), np.array(zs), np.array(ids, dtype=int)

    element_ids, part_ids, node_ids = np.array(element_ids, dtype=int), np.array(part_ids, dtype=int), np.array(
        node_ids, dtype=int)
    return xs, ys, zs, ids, element_ids, part_ids, node_ids

def find_centroids(node_ids, element_ids, part_ids, xs, ys, zs):
    centroid_coordinates = np.empty(shape=(node_ids.shape[0], 5))

    length = len(element_ids)
    with tqdm(total=length) as pbar:
        for i, (element_id, part_id, node_id) in enumerate(zip(element_ids, part_ids, node_ids)):
            coordinate_indexes = np.searchsorted(ids, node_id)
            x = np.mean(xs[coordinate_indexes])
            y = np.mean(ys[coordinate_indexes])
            z = np.mean(zs[coordinate_indexes])
            centroid_coordinates[i, :] = [element_id, part_id, x, y, z]
            pbar.update(1)

    print("Building dataframe...")

    df_centroid_coordinates = pd.DataFrame({"element_id": centroid_coordinates[:, 0],
                                            "part_id": centroid_coordinates[:, 1],
                                            "x": centroid_coordinates[:, 2],
                                            "y": centroid_coordinates[:, 3],
                                            "z": centroid_coordinates[:, 4]})

    print("Geometry done...")

    geometry_end_time = time.time()
    geometry_time_lapsed = geometry_end_time - geometry_start_time
    time_convert(geometry_time_lapsed)

    return df_centroid_coordinates

def read_load_curves(path):
    load_start_time = time.time()

    print("Creating load curve dictionary...")

    load_curve_dict = {}
    # with tqdm(total=os.path.getsize(path)) as pbar:
    with open(path, "r", encoding="utf8") as file:
        load_curve_time = []
        load_curve_value = []
        reading_curve = False
        read_element_id = False
        n = 0

        print("Populating dictionary with load curves...")

        for line_no, line in enumerate(file):
            if reading_curve:
                if line_no == n + 1:
```

```
                    continue
                if read_element_id:
                    element_id = line.split()[0]
                    load_curve_dict[int(element_id)] = [np.array(load_curve_time),
                                                        np.array(load_curve_value)]
                    load_curve_time = []
                    load_curve_value = []
                    reading_curve = False
                    read_element_id = False
                if "*LOAD_HEAT_GENERATION_SOLID" in line:
                    read_element_id = True
                elif len(line.split()) < 3:
                    split_line = line.strip().split()
                    t, value = split_line
                    load_curve_time.append(float(t))
                    load_curve_value.append(float(value))
            if "*DEFINE_CURVE" in line:
                reading_curve = True
                n = line_no

                # pbar.update(len(line))

    print("Heat generation done...")

    load_end_time = time.time()
    load_time_lapsed = load_end_time - load_start_time
    time_convert(load_time_lapsed)

    return load_curve_dict

def find_max_load(load_curve_dict: dict):
    max_load = 0.
    time_max_load = 0.
    element_id_max_load = 0.
    value_max_load = 0.
    for key, item in load_curve_dict.items():
        load_curve_time = item[0]
        load_curve_values = item[1]
        key_max_load = np.max(load_curve_values)
        if key_max_load < max_load:
            continue

        max_load = key_max_load
        max_load_index = np.argmax(load_curve_values)
        value_max_load = load_curve_values[max_load_index]
        time_max_load = load_curve_time[max_load_index]
        element_id_max_load = key

    print(f"{max_load = }, {time_max_load = }, {element_id_max_load = }")
    return time_max_load, value_max_load

def find_elementwise_max_loads(load_curve_dict: dict):
    max_loads = {}
    for key, item in load_curve_dict.items():
        load_curve_values = item[1]
        max_loads[key] = np.max(load_curve_values)
    max_loads = pd.Series(max_loads)
    return max_loads

def interpolate_load_to_time(load_curve_dict, time):
    load_interpolated = {}
    for key, item in load_curve_dict.items():
        load_curve_time = item[0]
        load_curve_values = item[1]
        load_interpolated[int(key)] = np.interp(time, load_curve_time, load_curve_values)
    load_interpolated = pd.Series(load_interpolated)
    return load_interpolated

def integrate_load(load_curve_dict):
    load_integrated = {}
    for key, item in load_curve_dict.items():
        load_curve_time = item[0]
        load_curve_values = item[1]
```

```
        load_integrated[int(key)] = integrate.trapezoid(load_curve_values, load_curve_time)
    load_integrated = pd.Series(load_integrated)
    return load_integrated

class LoadCurveTimeSlider:
    def __init__(self, mesh, load_curve_dict, scalars):
        self.output = mesh  # Expected PyVista mesh type
        self.load_curve_dict = load_curve_dict
        self.scalars = scalars
        # default parameters
        self.kwargs = {
            'time': 0.
        }

    def __call__(self, param, value):
        self.kwargs[param] = value
        self.update()

    def update(self):
        # This is where you call your simulation
        interpolated_load = interpolate_load_to_time(self.load_curve_dict, **self.kwargs)
        self.scalars["load"] = interpolated_load
        self.scalars = self.scalars.fillna(0)
        self.output["load"] = self.scalars["load"]
        return

print("\nOpening LS-Dyna input file...")
xs, ys, zs, ids, element_ids, part_ids, node_ids = read_mesh(mesh_path)

print("Creating grid in PyVista")

points = np.stack([xs, ys, zs]).T
node_ids_indexes = np.searchsorted(ids, node_ids)
cells = np.insert(node_ids_indexes, 0, 8, axis=1).ravel()

celltypes = np.empty(node_ids_indexes.shape[0], dtype=np.uint8)
celltypes[:] = CellType.HEXAHEDRON

grid = pv.UnstructuredGrid(cells, celltypes, points)

print("Opening file with heat generation load curves...")
if Path(file_path + 'load_curve_dictionary.pkl').is_file():
    with open(file_path + 'load_curve_dictionary.pkl', 'rb') as f:
        load_curve_dict = pickle.load(f)
else:
    load_curve_dict = read_load_curves(load_curve_path)
    with open(file_path + 'load_curve_dictionary.pkl', 'wb') as f:
        pickle.dump(load_curve_dict, f)

print("Identifying peak energy deposition...")
time_max_load, value_max_load = find_max_load(load_curve_dict)

print("Interpolating load curves to create instantaneous plot at peak...")
load_interpolated_at_max_load_time = interpolate_load_to_time(load_curve_dict, time_max_load)

print("Finding max load for each curves to create plot showing elementwise peak load...")
elementwise_max_loads = find_elementwise_max_loads(load_curve_dict)

print("Computing cell volumes...")
grid = grid.compute_cell_sizes(length=False, area=False, volume=True)
cell_volumes = grid.cell_data["Volume"]

print("Integrating load curves...")
elementwise_total_energy_per_m3 = integrate_load(load_curve_dict)

print("Creating dataframe with cell load scalars")
scalars = pd.DataFrame({"element_id": element_ids})
scalars = scalars.set_index("element_id")
scalars["load_at_T_load_max"] = load_interpolated_at_max_load_time
scalars["elementwise_max_load"] = elementwise_max_loads
scalars["elementwise_total_energy_per_m3"] = elementwise_total_energy_per_m3
scalars = scalars.fillna(0)

```

```
print("Computing elementwise total energy:")
elementwise_total_energy = scalars["elementwise_total_energy_per_m3"].to_numpy()*cell_volumes
total_integrated_energy = np.sum(elementwise_total_energy)
scalars["elementwise_total_energy"] = elementwise_total_energy
print(f"Finished computing elementwise total energy. Combined total energy: {total_integrated_energy:.2f} J")

print("Populating grid with cell load data")
for key in scalars.keys():
    grid.cell_data.set_array(scalars[key].to_numpy(), key)

print("Plotting grid")

plotter = pv.Plotter()
_ = plotter.add_mesh_threshold(grid, show_edges=True, scalars="elementwise_max_load")
plotter.show()
```

A.12 Experimental Tests for the Window-Flange Bolted Connection

A.12.1 Torque-Friction Test Bench

The measurements for torque-friction relationship were performed using a torque–tensile testing bench compliant with ISO 16047 [5], shown in Figure A.18. The bench measures the axial clamping force, the applied tightening torque, and the torque acting on a reaction plate. Friction between the bolt head and the plate produces a measurable torque, which directly gives the bearing friction. The remaining torque corresponds to the thread and pitch contributions, allowing the thread friction coefficient to be determined.

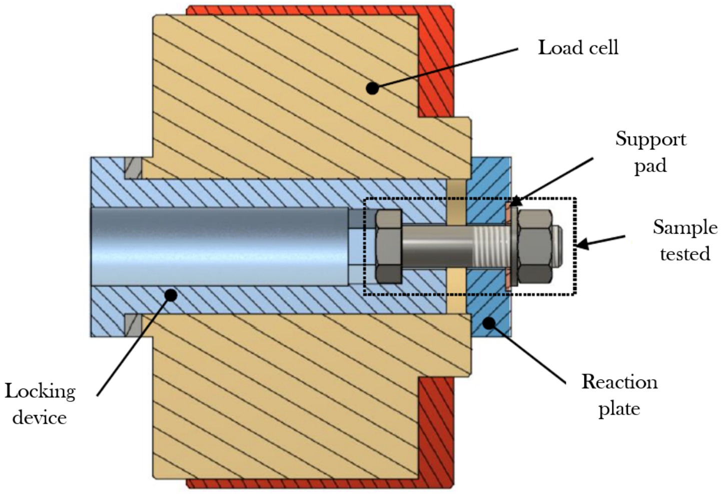

Figure A.18: Torque–tensile testing setup used to determine friction coefficients for the M14 titanium bolt configuration in accordance with ISO 16047.

A.13 Window-Vessel Joint Dimensioning According to VDI 2230

A.13.1 Analytical Calculations

The analytical calculations for the window-flange bolted connection were performed in accordance with VDI 2230 and are described in the following MathCad document.

S1 - Workload

Number of bolts $i := 72$

Internal pressure

Window inside diameter $D_{Ziflange} := 722.8\ mm - 2 \cdot 35\ mm = 0.653\ m$

Internal pressure $F_{IP} := \dfrac{\pi}{4 \cdot i} \cdot D_{Ziflange}^{\ 2} \cdot 0.2\ bar = 0.093\ kN$

Dump vibration

Mass of window $m_{win} := 30\ kg$

Force from acceleration $F_{VI} := \dfrac{1000\ \dfrac{m}{s^2} \cdot m_{win}}{i} = 0.417\ kN$

Gasket compression

Gasket compression $F_{GC} := 7000\ N$

Total load

Maximum load acting on bolt during service $F_{A.max} := F_{GC} + F_{VI} + F_{IP} = 7.51\ kN$

Minimum load acting on bolt during service $F_{A.min} := F_{GC} - F_{VI} + F_{IP} = 6.676\ kN$

S2 - Nominal bolt diameter

Following the bolt size estimation guide in VDI 2230 Table A7

Step	Evaluation	Result
Step A	$F_{A.max} = 7.51\ kN$ < $10\ kN$	$10\ kN$
Step B	+ 2 steps for dynamic and eccentrically applied load	$25\ kN$
Step C	+ 1 step for using a torque wrench	$40\ kN$
Step D	Strength grade 8.8 give nominal diameter:	$14\ mm$

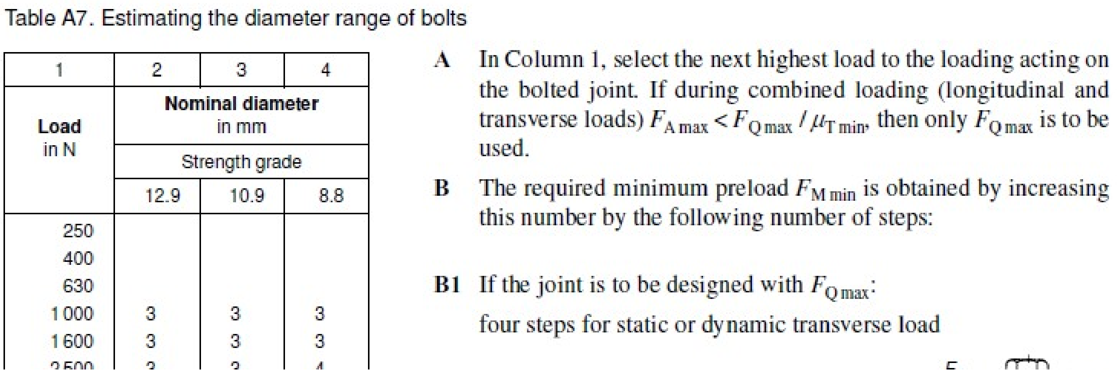

Table A7. Estimating the diameter range of bolts

1	2	3	4
Load in N	**Nominal diameter** in mm		
	Strength grade		
	12.9	10.9	8.8
250			
400			
630			
1 000	3	3	3
1 600	3	3	3

A In Column 1, select the next highest load to the loading acting on the bolted joint. If during combined loading (longitudinal and transverse loads) $F_{A\,max} < F_{Q\,max} / \mu_{T\,min}$, then only $F_{Q\,max}$ is to be used.

B The required minimum preload $F_{M\,min}$ is obtained by increasing this number by the following number of steps:

B1 If the joint is to be designed with $F_{Q\,max}$:

four steps for static or dynamic transverse load

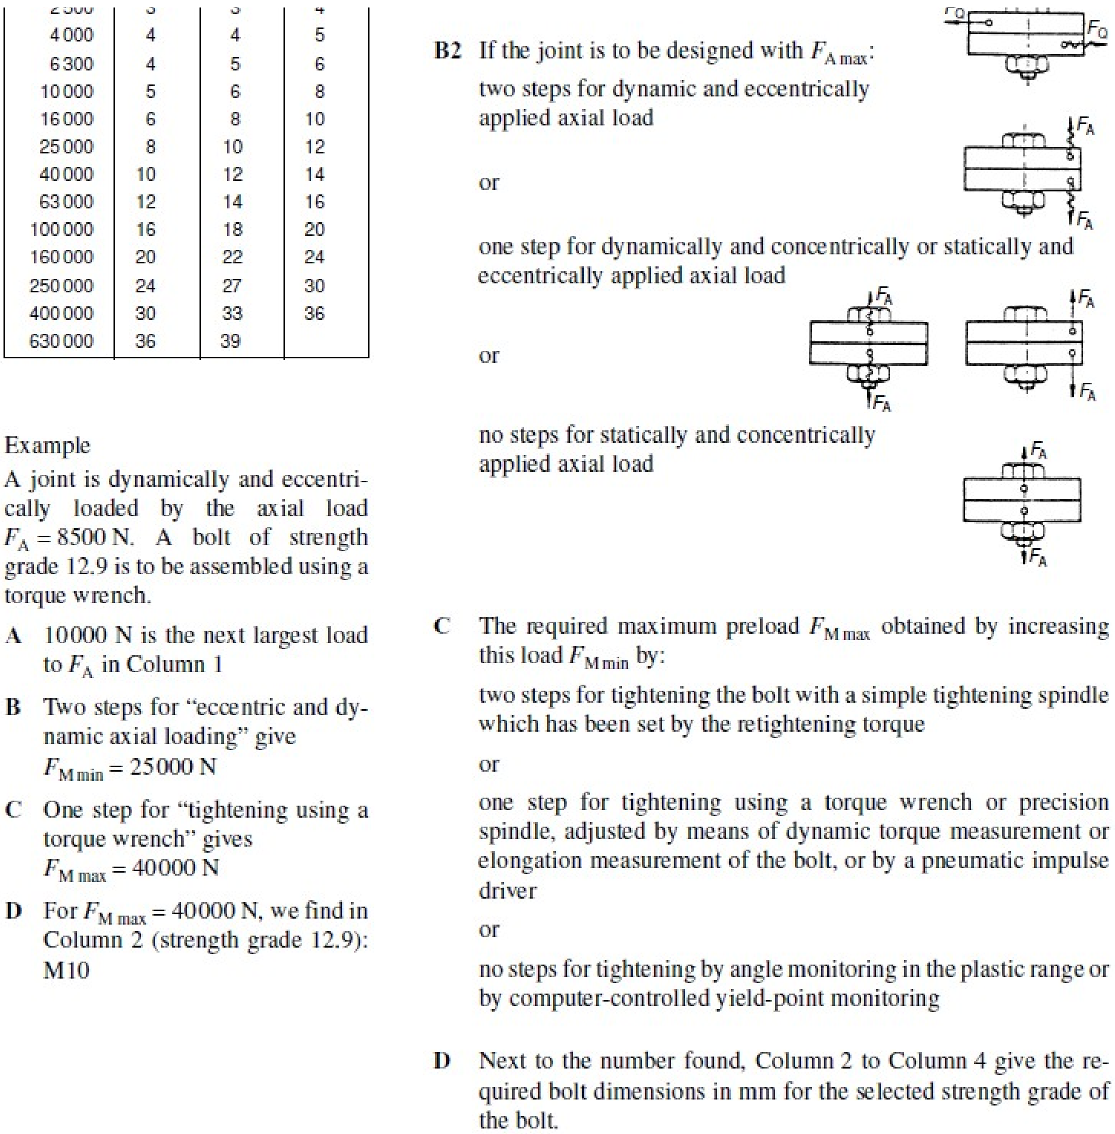

2 500	3	3	4
4 000	4	4	5
6 300	4	5	6
10 000	5	6	8
16 000	6	8	10
25 000	8	10	12
40 000	10	12	14
63 000	12	14	16
100 000	16	18	20
160 000	20	22	24
250 000	24	27	30
400 000	30	33	36
630 000	36	39	

Example

A joint is dynamically and eccentrically loaded by the axial load $F_A = 8500$ N. A bolt of strength grade 12.9 is to be assembled using a torque wrench.

A 10000 N is the next largest load to F_A in Column 1

B Two steps for "eccentric and dynamic axial loading" give $F_{M\,min} = 25000$ N

C One step for "tightening using a torque wrench" gives $F_{M\,max} = 40000$ N

D For $F_{M\,max} = 40000$ N, we find in Column 2 (strength grade 12.9): M10

B2 If the joint is to be designed with $F_{A\,max}$:

two steps for dynamic and eccentrically applied axial load

or

one step for dynamically and concentrically or statically and eccentrically applied axial load

or

no steps for statically and concentrically applied axial load

C The required maximum preload $F_{M\,max}$ obtained by increasing this load $F_{M\,min}$ by:

two steps for tightening the bolt with a simple tightening spindle which has been set by the retightening torque

or

one step for tightening using a torque wrench or precision spindle, adjusted by means of dynamic torque measurement or elongation measurement of the bolt, or by a pneumatic impulse driver

or

no steps for tightening by angle monitoring in the plastic range or by computer-controlled yield-point monitoring

D Next to the number found, Column 2 to Column 4 give the required bolt dimensions in mm for the selected strength grade of the bolt.

Converting bolt properties from steel to Titanium Grade 5.

$$d_{ti} := \sqrt{\frac{640\ MPa}{820\ MPa} \cdot (14\ mm)^2} = 12.368\ mm$$

<u>Taking M14 titanium is conservative</u>

Properties of selected bolt (Table A9 and A11 VDI 2230)

Nominal diameter $d := 14\ mm$

Pitch $P := 2\ mm$

Yield strength of the bolt $R_{p0.2min} := 820\ MPa$

Shear strength $\tau_{Bmin} := 600\ MPa$

Limiting surface pressure $p_G := 1340\ MPa$

Friction coeff. in the threads (experimental) $\mu_{Gmin} := 0.164$

Friction coeff. under head (experimental) $\mu_{Kmin} := 0.231$

Effective bolt head diameter $d_W := 20.67\ mm$

Diameter of clearance hole $d_h := 15.5\ mm$

Calculated parameters

Pitch diameter $d_2 := d - \frac{3 \cdot \sqrt{3}}{8} \cdot P = 12.701\ mm$

Minor diameter of the bolt $d_3 := d - P \cdot \left(\frac{17 \cdot \sqrt{3}}{24}\right) = 11.546\ mm$

S3 - Tightening factor

Using VDI 2230 Table A8 for friction coefficient class B, knowing a precision torque wrench will be used and choosing a high value in the allowed range give the tightening factor:

Tightening factor $\alpha_A := 1.6$

Table A8. Guide values for the tightening factor α_A (continued)

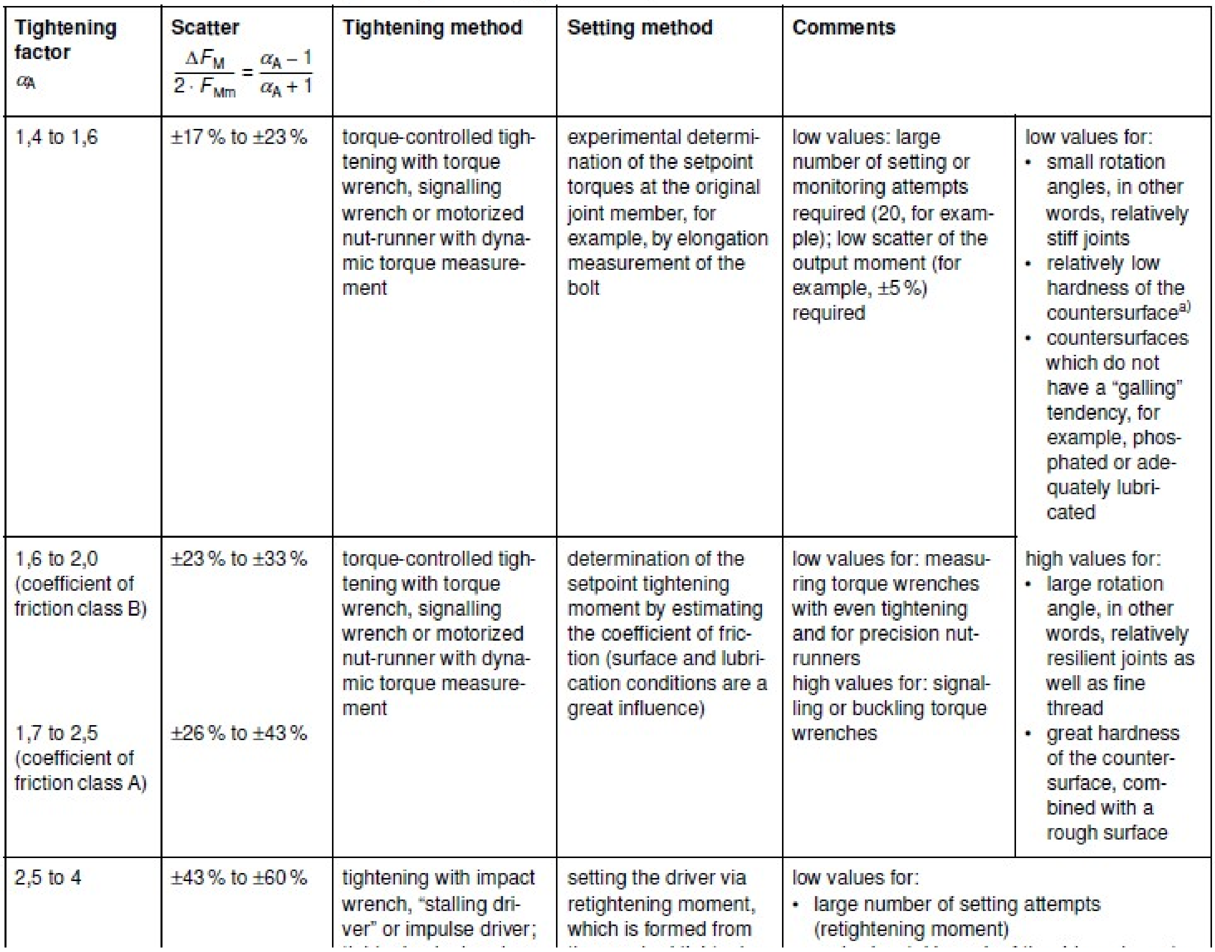

Tightening factor α_A	Scatter $\frac{\Delta F_M}{2 \cdot F_{Mm}} = \frac{\alpha_A - 1}{\alpha_A + 1}$	Tightening method	Setting method	Comments	
1,4 to 1,6	±17 % to ±23 %	torque-controlled tightening with torque wrench, signalling wrench or motorized nut-runner with dynamic torque measurement	experimental determination of the setpoint torques at the original joint member, for example, by elongation measurement of the bolt	low values: large number of setting or monitoring attempts required (20, for example); low scatter of the output moment (for example, ±5 %) required	low values for: • small rotation angles, in other words, relatively stiff joints • relatively low hardness of the countersurface$^{a)}$ • countersurfaces which do not have a "galling" tendency, for example, phosphated or adequately lubricated
1,6 to 2,0 (coefficient of friction class B)	±23 % to ±33 %	torque-controlled tightening with torque wrench, signalling wrench or motorized nut-runner with dynamic torque measurement	determination of the setpoint tightening moment by estimating the coefficient of friction (surface and lubrication conditions are a great influence)	low values for: measuring torque wrenches with even tightening and for precision nut-runners high values for: signalling or buckling torque wrenches	high values for: • large rotation angle, in other words, relatively resilient joints as well as fine thread • great hardness of the countersurface, combined with a rough surface
1,7 to 2,5 (coefficient of friction class A)	±26 % to ±43 %				
2,5 to 4	±43 % to ±60 %	tightening with impact wrench, "stalling driver" or impulse driver;	setting the driver via retightening moment, which is formed from	low values for: • large number of setting attempts (retightening moment)	

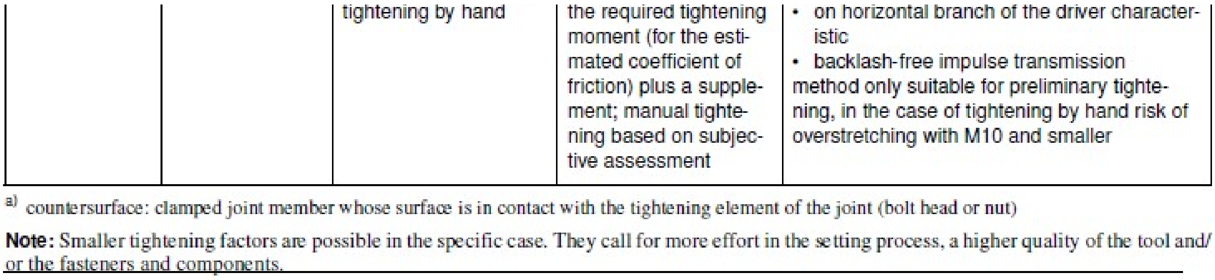

		tightening by hand	the required tightening moment (for the estimated coefficient of friction) plus a supplement; manual tightening based on subjective assessment	• on horizontal branch of the driver characteristic • backlash-free impulse transmission method only suitable for preliminary tightening, in the case of tightening by hand risk of overstretching with M10 and smaller

[a] countersurface: clamped joint member whose surface is in contact with the tightening element of the joint (bolt head or nut)

Note: Smaller tightening factors are possible in the specific case. They call for more effort in the setting process, a higher quality of the tool and/or the fasteners and components.

S4 - Maximum assembly preload

Utilization factor determines how much of the capacity of the bolt that should be used during assembly.

Calculation quantities

Stress diameter $d_S := 0.5 \cdot \left(d_2 + d_3\right) = 12.124\ mm$

Stress area $A_S := \frac{\pi}{4} \cdot {d_S}^2 = 115.439\ mm^2$

Parameters

Bolt capacity utilization factor $v := 0.9$

Calculations

Allowed assembly stress $\sigma_{Mzul} := v \cdot R_{p0.2min} = 738\ MPa$

Maximum permitted assembly load

$$F_{Mzul} := A_S \cdot \frac{v \cdot R_{p0.2min}}{\sqrt{1 + 3 \cdot \left(\frac{3}{2} \cdot \frac{d_2}{d_S} \cdot \left(\frac{P}{\pi \cdot d_2} + 1.155 \cdot \mu_{Gmin}\right)\right)^2}} = 71.366\ kN$$

S5 - Minimum assembly preload

Minimum assembly preload $F_{Mzul.min} := \frac{F_{Mzul}}{\alpha_A} = 44.603\ kN$

The VDI 2230 calculation predicts an expected preload range between 44.6 kN and 71.4 kN. However, experimental tightening tests provide a more realistic estimate of the scatter associated with the torque wrench.

$\boxed{F_{Mzul}} := 76.5\ kN$

Based on these measurements, a minimum preload of 41.6 kN and a maximum preload of 76.5kN are used in the validation.

$\boxed{F_{Mzul.min}} := 41.6\ kN$

S6 - Results and analysis

Finite element model developed with the maximum and minimum loads applied. For the maximum load the integrity of the bolt must be validated. For the minimum load, prying effects must be evaluated.

Max assembly load

Max bolt load $F_{Smax} := 76.5\ kN$

Max bolt bending moment (extracted from FE-model) $M_{Sbo} := 7.6\ N \cdot m$

Min assembly load

Min bolt load $F_{Vmin} := 41.6\ kN$

S7 - Work stress

Parameters

Residual torque factor. Remaining torsional stress in the bolt after tightening (assumed value from VDI 2230). $k_t := 0.5$

Calculations

Bending modulus for bolt cross section

$$W_S := \frac{\pi}{32} \cdot {d_S}^3 = 174.943\ mm^3$$

Thread torque caused by tightening

$$M_G := F_{Mzul} \cdot \frac{d_2}{2} \cdot \left(\frac{P}{\pi \cdot d_2} + 1.155 \cdot \mu_{Gmin} \right) = 116.373\ N \cdot m$$

Polar moment of inertia to compute shear stress from torque

$$W_P := \frac{\pi}{16} \cdot {d_S}^3 = 349.885\ mm^3$$

Combined axial and bending normal stress

$$\sigma_{zb.max} := \frac{F_{Smax}}{A_S} + \frac{M_{Sbo}}{W_S} = 706.128\ MPa$$

Bending stress (isolated)

$$\sigma_{Sbo} := \frac{M_{Sbo}}{W_S} = 43.443\ MPa$$

Torsional shear stress

$$\tau_{max} := \frac{M_G}{W_P} = 332.604\ MPa$$

Working (von mises) stress

$$\sigma_{red.B} := \sqrt{{\sigma_{zb.max}}^2 + 3 \cdot \left(k_t \cdot \tau_{max} \right)^2} = 762.618\ MPa$$

Bending stress contribution to work stress

$$\Delta\sigma_{Sbo.p} := \frac{\sigma_{Sbo}}{\sigma_{red.B}} = 0.057 \qquad 5.7\%$$

Criteria

$$\sigma_{red.B} = 762.618\ MPa \ < \ R_{p0.2min} = 820\ MPa$$

Safety factor

$$S_F := \frac{R_{p0.2min}}{\sigma_{red.B}} = 1.075$$

At the maximum considered bolt load the equivalent stress is 763 MPa which is below the yield strength of 820 MPa with a factor of 1.075. The maximum bolt load should be withstood

S8 - Clamping requirements

Pressure plot at minimum load of 41.6 kN

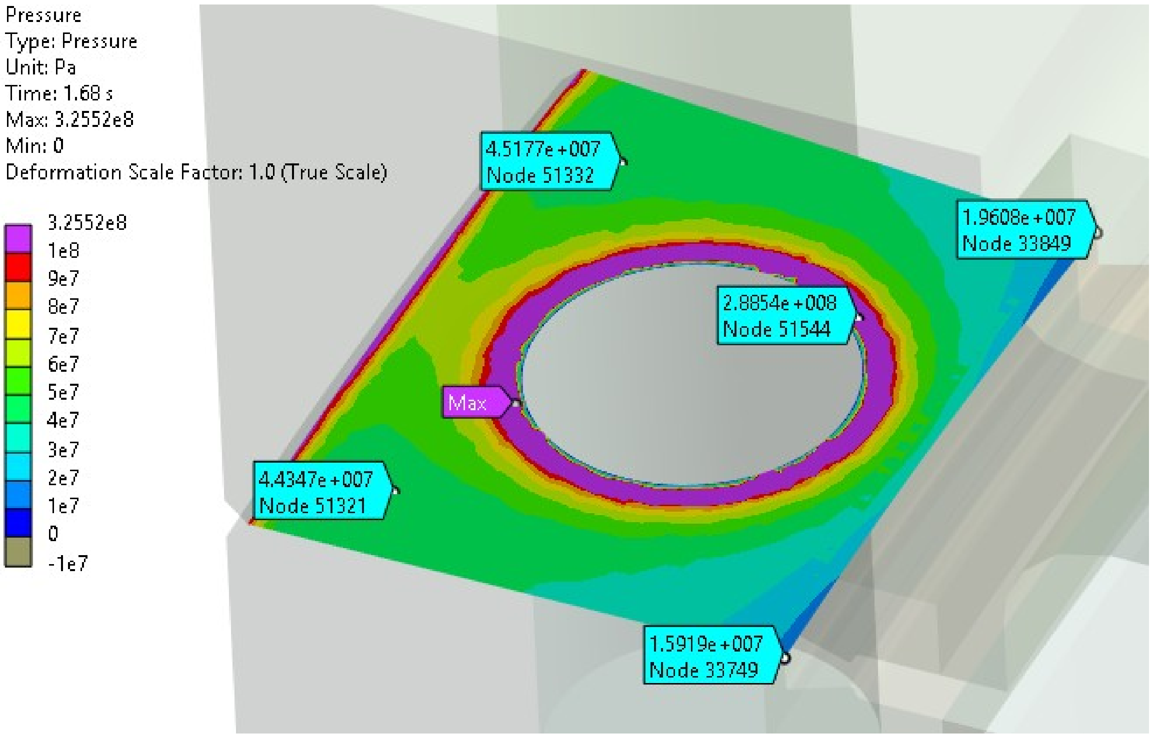

Contact pressure is present on the complete surface with minimum values close to the gasket knives of 15 MPa. The average contact pressure is 57 MPa.

S9 - Tightening torque

Average head friction diameter $$D_{Km} := \frac{d_W + d_h}{2} = 18.085\ mm$$

Average load between minimum and maximum $$F_{avg} := \frac{F_{Mzul} + F_{Mzul.min}}{2}$$

Tightening torque for average load $$M_A := F_{avg} \cdot \left(0.16 \cdot P + 0.58 \cdot d_2 \cdot \mu_{Gmin} + \frac{D_{Km}}{2} \cdot \mu_{Kmin}\right) = 213.58\ N \cdot m$$

S10 - Alternating stress

Titanium threads on the bolt are rolled after heat treatment

Permissible continuous alternating stress for bolts rolled before heat treatment (VDI 2230 page 157)

$$\sigma_{ASV} := 0.85 \cdot \left(\frac{150}{12} + 45\right) \mathit{MPa} = 48.875 \ \mathit{MPa}$$

Worst case pre-load

$$F_{Mzul} - \left(F_{VI} + F_{IP}\right) = 75.99 \ \mathit{kN}$$

Maximum axial additional load due to external load

$$F_{SAO} := F_{VI} + F_{IP} = 0.51 \ \mathit{kN}$$

Minimum axial additional load due to external load

$$F_{SAU} := -F_{VI} + F_{IP} = -0.324 \ \mathit{kN}$$

Steady mean load on the bolt

$$F_{sm} := \frac{F_{SAO} + F_{SAU}}{2} + F_{Mzul} = 76.593 \ \mathit{kN}$$

Permissible continuous alternating stress for bolts rolled after heat treatment

$$\sigma_{ASG} := \left(2 - \frac{F_{sm}}{R_{p0.2min} \cdot A_S}\right) \cdot \sigma_{ASV} = 58.204 \ \mathit{MPa}$$

Due to the high preload, the joint is almost static, meaning that the alternating stress is far below the permissible one.

S11 - Surface pressure (under bolt head)

Bolt head minimum pressure area

$$A_{p.min} := \frac{\pi}{4} \cdot \left(d_W{}^2 - d_h{}^2\right) = 146.869 \ \mathit{mm}^2$$

Assembled state

$$p_{M.max} := \frac{F_{Mzul}}{A_{p.min}} = 520.874 \ \mathit{MPa}$$

Achieved surface pressure is below the allowed.

$$p_{M.max} = 520.874 \ \mathit{MPa} < \ p_G = \left(1.34 \cdot 10^3\right) \ \mathit{MPa}$$

S12 - Length of engagement

Pitch diameter

$$D_2 := d_2 = 12.701 \ \mathit{mm}$$

Minor diameter of the bolt

$$D_1 := d_3 = 11.546 \ \mathit{mm}$$

Proof strength of the bolt

$$R_{S.p0.2.min} := R_{p0.2min}$$

$$R_{mS} := R_{S.p0.2.min} = 820 \ \mathit{MPa}$$

Ultimate tensile strength of the threaded material (Table A9)

$$R_{mM} := 890 \ \mathit{MPa}$$

Shear strength ratio from Table 6

$$\tau_B := 0.6 \cdot R_{mM} = 534 \ \mathit{MPa}$$

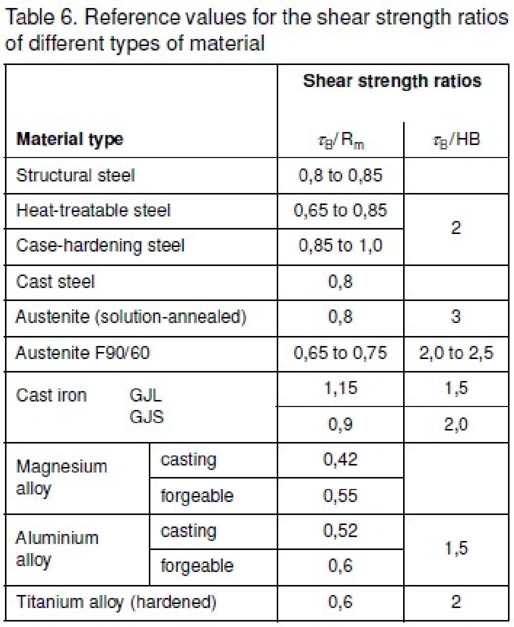

Table 6. Reference values for the shear strength ratios of different types of material

Material type		Shear strength ratios τ_B/R_m	τ_B/HB
Structural steel		0,8 to 0,85	
Heat-treatable steel		0,65 to 0,85	2
Case-hardening steel		0,85 to 1,0	
Cast steel		0,8	
Austenite (solution-annealed)		0,8	3
Austenite F90/60		0,65 to 0,75	2,0 to 2,5
Cast iron	GJL	1,15	1,5
	GJS	0,9	2,0
Magnesium alloy	casting	0,42	
	forgeable	0,55	
Aluminium alloy	casting	0,52	1,5
	forgeable	0,6	
Titanium alloy (hardened)		0,6	2

Assessment of whether the thread in the hole or the thread on the bolt will fail first:

$$R_S := \frac{d \cdot \left(\frac{P}{2} + (d - D_2) \cdot \tan(30\ deg)\right)}{d \cdot \left(\frac{P}{2} + (d_2 - D_1) \cdot \tan(30\ deg)\right)} \cdot \frac{R_{mM}}{R_{mS}} = 1.14$$

> 1 : bolt fails first

< 1 : thread fails first

It is expected that the bolt will fail first.

Outer diameter of component with internal thread

$$s := 10\ mm \cdot 2 = 20\ mm$$

$$\frac{s}{d} = 1.429$$

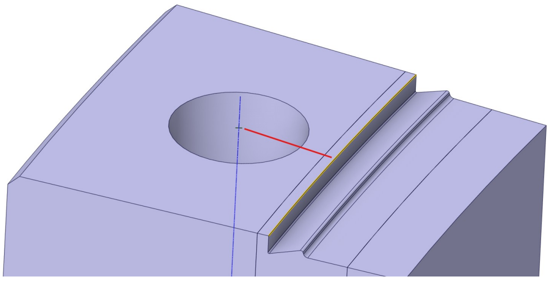

Lowest possible thread distance to the wall. Probably not representative so lowest C1 is used of 1.4

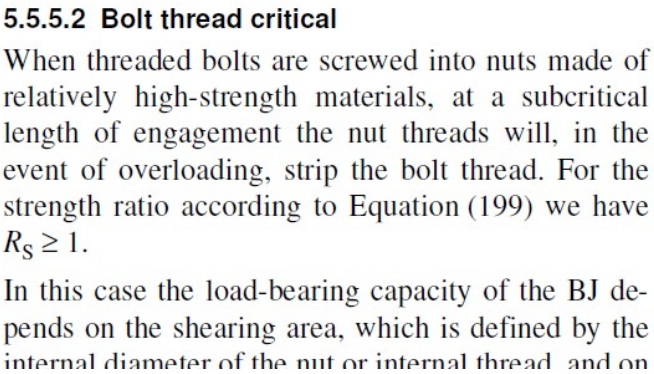

5.5.5.2 Bolt thread critical

When threaded bolts are screwed into nuts made of relatively high-strength materials, at a subcritical length of engagement the nut threads will, in the event of overloading, strip the bolt thread. For the strength ratio according to Equation (199) we have $R_S \geq 1$.

In this case the load-bearing capacity of the BJ depends on the shearing area, which is defined by the internal diameter of the nut or internal thread, and on

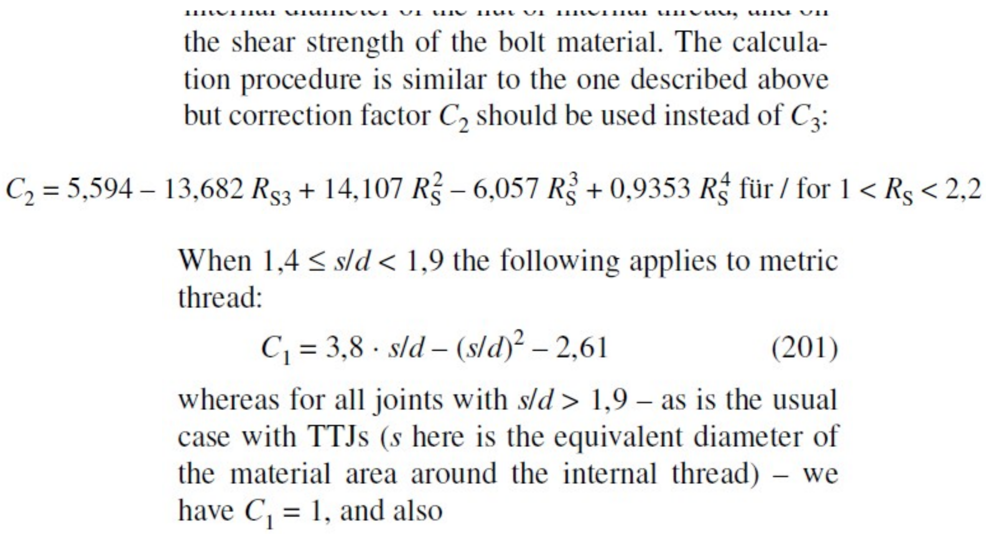

internal diameter of the nut or internal thread, and on the shear strength of the bolt material. The calculation procedure is similar to the one described above but correction factor C_2 should be used instead of C_3:

$$C_2 = 5{,}594 - 13{,}682\ R_{S3} + 14{,}107\ R_S^2 - 6{,}057\ R_S^3 + 0{,}9353\ R_S^4 \text{ für / for } 1 < R_S < 2{,}2$$

When $1{,}4 \leq s/d < 1{,}9$ the following applies to metric thread:

$$C_1 = 3{,}8 \cdot s/d - (s/d)^2 - 2{,}61 \qquad (201)$$

whereas for all joints with $s/d > 1{,}9$ – as is the usual case with TTJs (s here is the equivalent diameter of the material area around the internal thread) – we have $C_1 = 1$, and also

$$C_1 := 3.8 \cdot \frac{s}{d} - \left(\frac{s}{d}\right)^2 - 2.61 = 0.778$$

$$C_2 := 5.594 - 13.682 \cdot R_S + 14.107 \cdot {R_S}^2 - 6.057 \cdot {R_S}^3 + 0.9353 \cdot {R_S}^4 = 0.936$$

Minimum required effective thread length:

$$m_{ges} := \frac{R_{mM} \cdot A_S \cdot P}{C_1 \cdot C_2 \cdot \tau_B \cdot \left(\frac{P}{2} + \left(d_2 - D_1\right) \cdot \tan\left(30\ deg\right)\right) \cdot \pi \cdot d} + 1.2 \cdot P = 9.612\ mm$$

Comparing tensile strength of bolt vs. the threaded hole

Thread length in the design $m_{gesvorh} := 20\ mm$

Non effective engagement length (subtract 2 threads at each end) $m_{effvorh} := m_{gesvorh} - 2 \cdot P = 16\ mm$

The threaded length is longer than the minimum required. $m_{effvorh} = 16\ mm \quad > \quad m_{ges} = 9.612\ mm$

A.13.2 Finite Element Model of Window-Flange Connection

This section describes the window-flange finite element submodel used to validate the joint design using inputs from the analytical calculations presented in the previous section.

Figure A.19 shows the geometry of the model. The submodel represents a 5° circumferential segment of the window-flange connection containing a single bolt. Both the window and the flange include knife edges that plastically indent the copper gasket during bolt preload to create a metal-to-metal vacuum seal. The detailed knife geometry is explicitly included in the model.

Partitions are introduced primarily to enable local mesh refinement at the knife-gasket interface, where accurate resolution of the contact pressure is required. The bolt is modelled as a solid body.

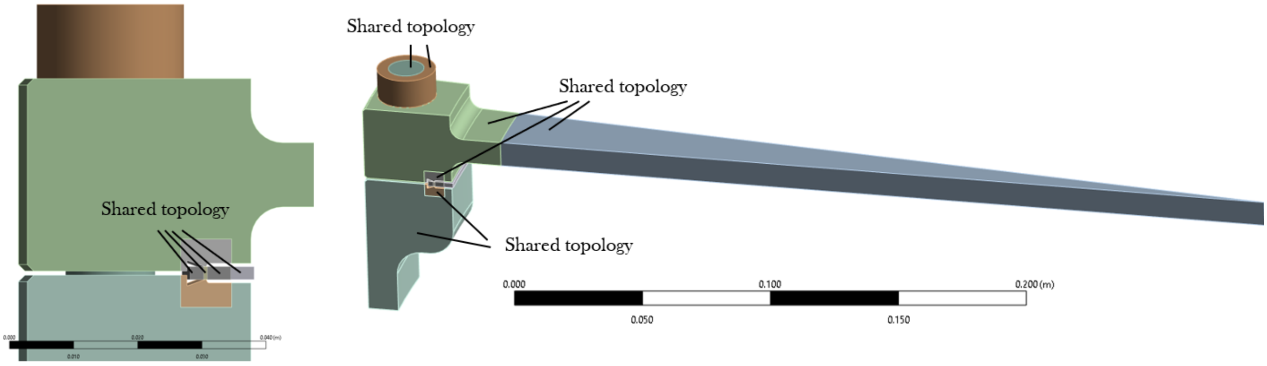

Figure A.19: Finite element model of the window–flange submodel. A 5° circumferential segment containing a single bolt is used to represent the full joint. Both the flange and the window include a knife-edge feature that plastically compresses the copper gasket during bolt preload to form the seal.

Figure A.20 shows the mesh used in the submodel. The model contains approximately 111k nodes and 175k elements. Linear elements are used throughout most of the model, while quadratic elements are applied in the knife-gasket region to improve solution accuracy. The mesh consists of a combination of hexahedral, tetrahedral, pyramidal, and wedge elements. Local mesh refinement is introduced at the knife-gasket interface and at all contact surfaces to ensure an accurate representation of contact forces and the resulting contact pressures.

Table A.16 summarizes the contact definitions used in the model, while Figure A.21 shows their locations.

The screw is connected to the threaded hole in the flange using a bonded contact, representing the engagement of the threads. The contact between the knife edge and the copper gasket is modelled as frictional with a coefficient of friction of 0.1. An augmented Lagrange formulation is used for this interface to better resolve the contact kinematics during the plastic indentation of the gasket.

The contacts between the flange and the window, and between the bolt head and the window, are modelled as frictional using a penalty-based formulation with a coefficient of friction of 0.12. Since the dominant loads act normal to these surfaces, the exact value of the friction coefficient has a limited influence on the results.

Table A.16: Contact definitions used in the window–flange finite element submodel.

Contact Interface	Contact Type	Formulation	Friction Coefficient
Screw threads – hole	Bonded	—	—
Knife edge – copper gasket	Frictional	Augmented Lagrange	0.10
Flange – window	Frictional	Penalty method	0.12
Bolt head – window	Frictional	Penalty method	0.12

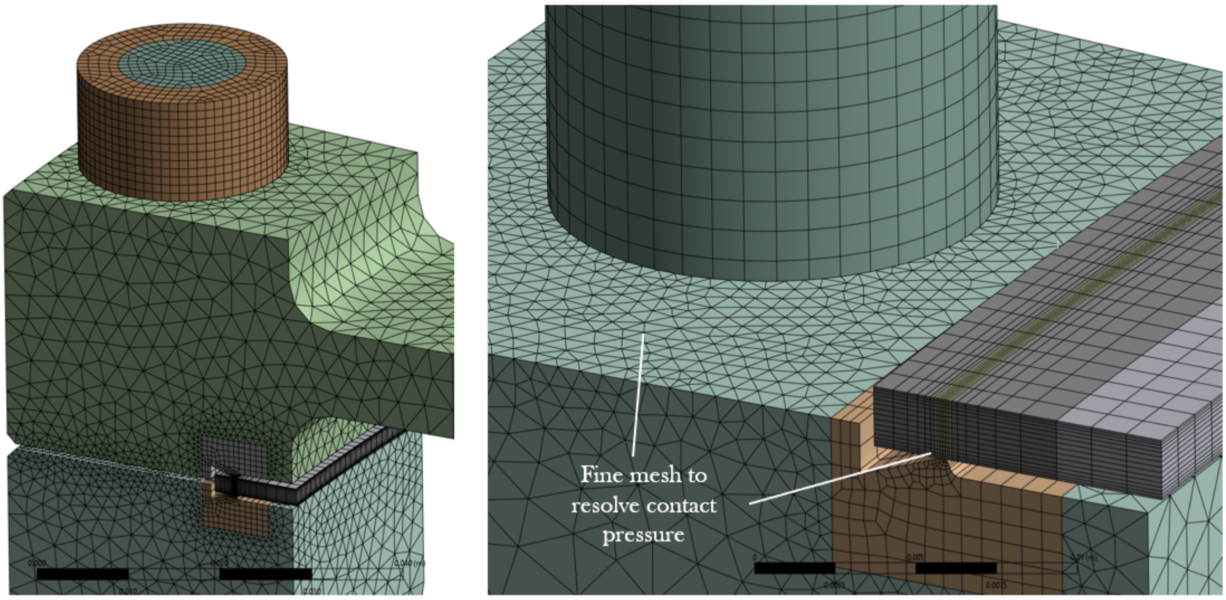

Figure A.20: Overview of the mesh used in the submodel. A refined mesh is applied at the knife–gasket interface to accurately resolve the contact pressure, and along the contact surface between the flange and the window.

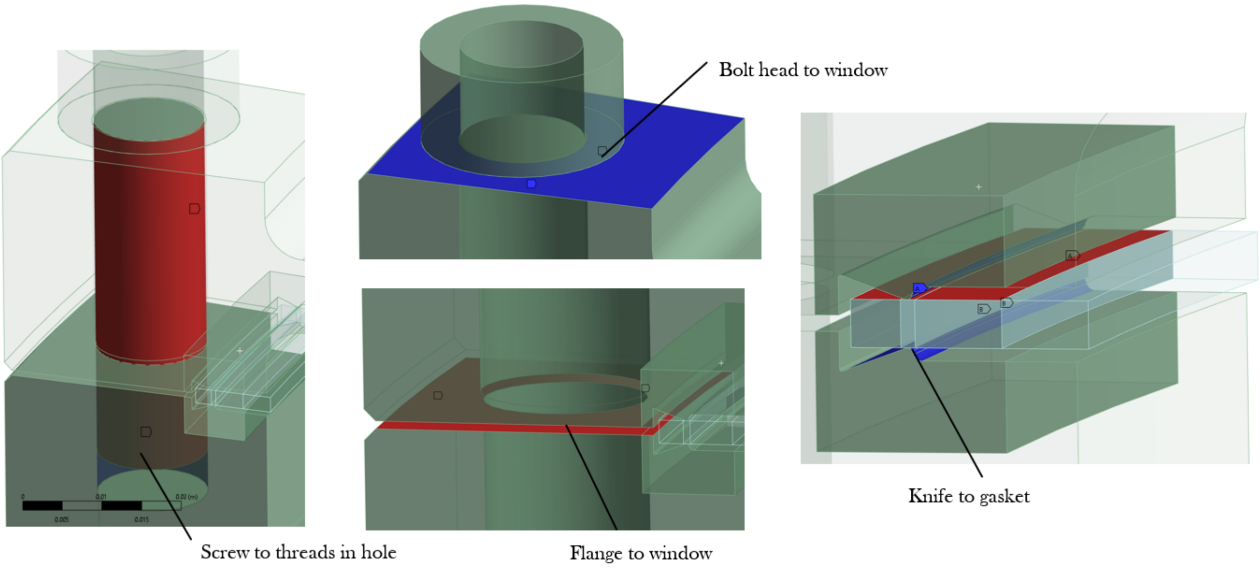

Figure A.21: Locations of the contact interfaces defined in the window–flange submodel.

Figure A.22 shows the boundary conditions applied in the model. These are summarized below:

- Symmetry boundary conditions are applied on the two radial faces of the 5° model segment to represent the periodic continuation of the full circular flange.
- A frictionless support is applied on the bottom surface of the flange to prevent displacement normal to the flange plane while allowing in-plane deformation.
- A cylindrical support is applied on the inner diameter of the gasket to improve convergence and stabilize the deformation during compression. Additionally, remote displacement constraints restricting rotations

are applied on the outer top and bottom surfaces of the gasket. These constraints are necessary to stabilize the otherwise weakly constrained gasket while it is plastically compressed between the two knife edges.

- An internal pressure of 0.2 bar is applied to all internal surfaces to represent the overpressure of the argon atmosphere inside the vessel.
- A bolt preload of 76.5 kN is applied to the screw.
- A force of 0.4 kN is applied to the window surface that contacts the flange, acting away from the flange. This represents the inertial force of the window when the vessel contracts. This loading condition is conservative for the bolt since it increases the tensile load in the bolt.

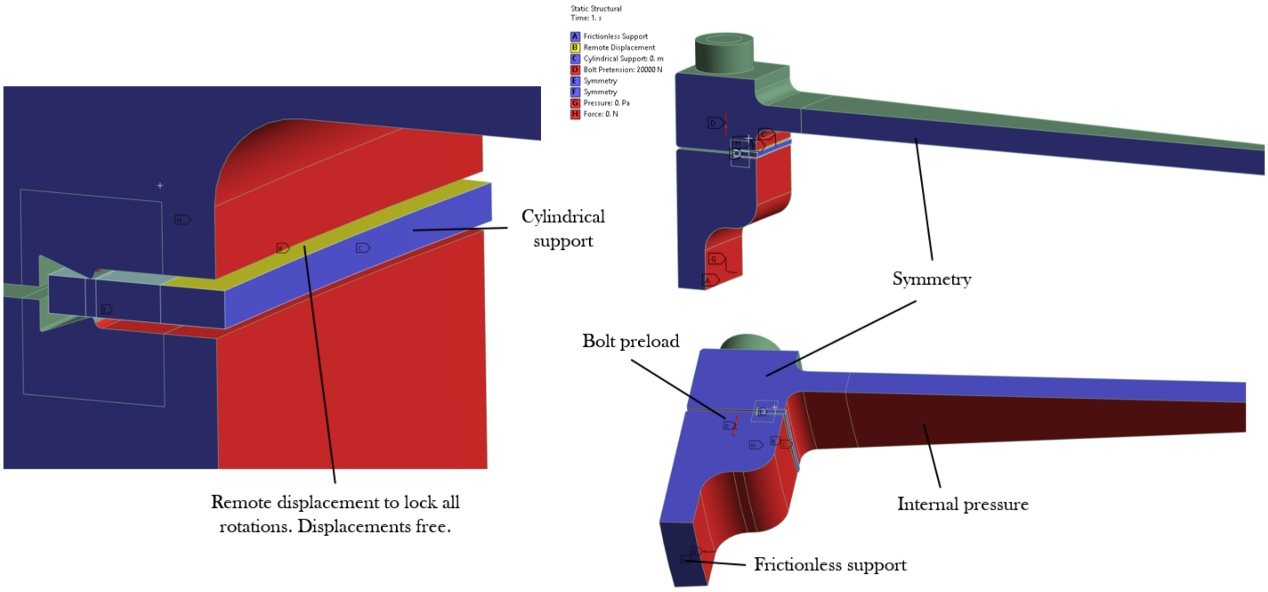

Figure A.22: Overview of the boundary conditions applied in the window–flange finite element submodel.

The material model for titanium grade 5 for the flange, window and screw is consistent with the information provided in Appendix A.4.1. For the copper gasket, the material properties of the elastic-plastic model is provided in Table A.17. Plastic deformation of the copper gasket is modelled using a bilinear kinematic hardening law. The material response is represented by a linear elastic region defined by Young's modulus, followed by a linear plastic region defined by the yield strength and a constant tangent modulus. In the kinematic hardening formulation, the yield surface translates in stress space during plastic deformation. This simplified plasticity model is sufficient to represent the plastic indentation of the copper gasket produced by the knife edges during bolt preload.

Table A.17: Material properties used for the copper gasket in the finite element model.

Property	Value	Unit
Density	3230	$\mathrm{kg\,m^{-3}}$
Young's Modulus	115	GPa
Poisson's Ratio	0.31	–
Bulk Modulus	100.9	GPa
Shear Modulus	43.9	GPa
Yield Strength	35	MPa
Tangent Modulus	500	MPa