

Design and Validation of an Enhanced Cooling System Model for the High-Luminosity LHC Beam Dump at CERN

Giovanni Angelo Roberto Binda Notarianni

Mechanical Engineering Master

Master Thesis FS 2025

Institute of Fluid Dynamics
ETH Zurich

Supervisor: Gabriel Banks (CERN)

Professor: Patrick Jenny

Abstract

This thesis aims to investigate the thermal performance of the High-Luminosity Large Hadron Collider (HL-LHC) external beam dump cooling system, with a focus on identifying upgrade strategies that are both effective and practically feasible under the increased energy deposition expected during HL operation.

A computational framework capable of modeling complex multi-material geometries and incorporating a temperature-dependent thermal contact resistance (TCC) model was developed in *Ansys Fluent*. Multiple cooling layouts were evaluated through mesh sensitivity analyses, variations in material properties and flow distribution studies. Key performance indicators included internal energy evolution, energy extraction efficiency and evaluation of thermal-mechanical coupling via contact pressure mapping.

Among the explored configurations, the option involving an additional set of cooling ducts placed above the beam dump, referred to as the top ducts configuration, was identified as the most balanced solution. While other configurations involving doubling the inlet massflow achieved the highest total cooling performance, they required substantial modifications to the existing ventilation system. In contrast, the top ducts are expected to reduce peak temperatures in critical regions, enhance thermal contact at the interface between dump components and improve the mechanical stability of internal components during beam dump events. Furthermore, this solution also eliminates the need for air handling unit modifications and simplifies installation within radiologically controlled areas.

In conclusion, the top ducts upgrade appears to be the most promising upgrade strategy for HL-LHC beam dump integration, offering a practical compromise between performance and feasibility. This outcome was made possible by the computational framework developed in this work, which can also support dump block monitoring during HL operations and future cooling system benchmarking.

Acknowledgments

I would like to thank my supervisor Gabriel Banks, for the many pieces of advice given and the patience exerted throughout all the duration of my thesis and more.

Furthermore, I am also grateful to Nicola Solieri, Professor Patrick Jenny and Dr. Daniel Werner Meyer-Masseti for providing the opportunity to participate in this master project. Additionally I would like to express my gratitude to Asbjoern Lassen Lund, for his expertise and continuous support throughout the project, to the entire SY-STI-TCD TDE Team that made it possible to learn so many interesting things in such a light and fun environment, and to the whole SY-STI-TCD Section for the warm and much appreciated welcome.

Finally I would also like to thank CERN for allowing me to work with the Organization and providing me with all needed facilities and computational resources.

Contents

List of Figures	vi
List of Tables	xi
1 Introduction	1
2 Literature Review	5
2.1 Beam Intercepting Devices	5
2.1.1 Challenges	6
2.2 Target Dump for Ejected beam	7
2.2.1 LS2 Studies & Upgrades	8
2.2.2 High Luminosity Studies	12
2.2.3 Cooling System	14
3 Methods	17
3.1 Model	17
3.1.1 Computational Domain	17
3.1.2 Geometry	18
3.1.3 Perfect Contact Assumption and Contact Pressure Calculations	24
3.1.4 Meshing	24
3.1.5 Boundary Conditions	27
3.2 Simulation Setup	34
3.3 Model Checks	35
3.3.1 Basic Checks	35
3.3.2 Material Properties Sensitivity Analyses	36
3.3.3 Mesh Sensitivity Analyses	38
4 Results & Discussion	43
4.1 Run 3 Limits	43
4.2 Preliminary HL Investigation on Spare geometry	48
4.2.1 Baseline: 1x Flowrate Shrink Fitted CFC LD Section	49
4.2.2 Baseline: 2x Flowrate Shrink Fitted CFC LD Section	49
4.2.3 2x Flowrate Non-Shrink Fitted CFC LD Section	49
4.2.4 Top Ducts Addition	50
4.2.5 2x Flowrate Central Bottom Cooling	52
4.2.6 Top Ducts and Optimized Central Bottom Cooling	53
4.2.7 Practical Ducts	53
4.3 HL Dump Block Geometry	55
4.3.1 Considered Scenarios	55
4.3.2 Uncertainties	56

4.3.3	Internal Energy and Heat Flux Analysis	57
4.3.4	Contact Pressure Analyses	64
4.3.5	Sensitivity to Massflow Distribution	71
5	Conclusion & Outlook	73
5.1	Conclusion	73
5.2	Outlook	74
A		78
A.1	Simplified Steady State Calculations	78
A.2	Dummy Elements Layers Insertion	81
A.3	Calculations for Dummy Layer Conductivities	81
B		83
B.1	Previous Work	83
B.1.1	Nozzle Radius Study	83
B.1.2	Nozzle Density Study	84
B.1.3	Inclined Nozzles	85
B.1.4	Obstructed Ducts	86
C		95
C.1	AHU Datasheets	96

List of Figures

1.1	Final installation of the Run 3 beam dump assembly, shown with the complete cooling system and instrumentation in place, enclosed within the shielding structure.	1
1.2	Schematic representation of LHC underground point 6 where the two caverns with dead ends highlighted in red are present [5].	2
1.3	Transverse section of the beam dump where the sweep pattern of the beam is observable. Beam energy: 7 TeV [7].	3
2.1	Schematic of the extraction system and detail on the sweep pattern projected on the upstream window from [9]. For a detailed description of the extraction and dilution systems, the reader is referred to the introduction of [9].	7
2.2	Longitudinal cut plane showing the graphite positioning and densities [7].	8
2.3	Simulated energy density profiles along the TDE beam dump length for different LHC runs, assuming a nominal beam sweep pattern [6]. Energy density [J/g] is the amount of energy from the proton beam that is deposited into the dump material at each point and its peak indicates the physical location within the assembly where the most intense heating occurs.	8
2.4	Cross section of the dump installation until the end of Run 2. The blue connecting line can be observed connecting the orange extracting line and the beam dump block. (From [15]).	9
2.5	Picture showing the dump block within its shielding and the connecting line connected to it, from [15].	9
2.6	Left: peak temperature profiles along the length of the dump block in case of MKBH malfunction. Right: transversal cut in the maximum temperature's location showing the modified sweep pattern in case of 2 MKBH malfunction [5]. . .	10
2.7	Run 3 dump block installed in its new support system.	11
2.8	Crack in the extruded graphite disk develops from outgassing groove.	11
2.9	Evolution of the maximum vessel stress in the low density section for both SS 316LN (Red) and Ti Gr5 (Blue) with HL beam conditions and their respective yield strength.	12
2.10	Picture of the nozzles along the longitudinal ducts running below the dump block.	14
2.11	Schematic of impingement jet flow structure [23].	14
2.12	Schematic representation of the AHU and the external duct work conducting the 5000 m ³ /h of air to the bottom ducts.	14
2.13	Forced convection on the bottom side of the dump.	14
2.14	Air velocity distribution in the cooling ducts. The figure shows two cross-sectional velocity fields at different duct locations alongside the full 3D geometry of the duct assembly with indicated slice positions. Higher air velocities are observed at the downstream nozzles, resulting from pressure gradients along the duct. This uneven distribution highlights the influence of pressure drop on airflow delivery [24]. . . .	15
2.15	Initial geometry for the annular cooling investigation [24].	16

3.1	Detailed simulation workflow.	17
3.2	3D model of the dump block assembly, its shielding, the AHU and the cooling ducts.	18
3.3	Dimensions of the computational domain	18
3.4	Fluid domain geometry modeled in <i>Ansys SpaceClaim</i> , with the outlet highlighted in red.	20
3.5	Longitudinal section of the new HL beam dump block highlighting both the HD isostatic graphite (orange), the CFC stacks (blue) and the windows.	21
3.6	3D model of isostatic graphite block geometry.	21
3.7	CAD representation of a stack of CFC plates, showing the arrangement of 12 central plates and 2 end plates. The detail on the right highlights the interface between adjacent plates and the small gaps introduced by the geometry.	21
3.8	Solid vessel model divided into multiple sections to enable evaluation of temperatures and internal energies at different locations within the vessel.	22
3.9	Core model: isostatic graphite blocks (green) and CFC stacks (pink), separated by cylindrical argon gaps (grey).	22
3.10	Section view of submodel representing the CFC stack with argon filled gaps.	22
3.11	Section view of submodel representing the CFC stack as a monolithic block.	22
3.12	Max CFC temperature for monolithic approximation and gap inclusive model.	23
3.13	Max vessel temperature for monolithic approximation and gap inclusive model.	23
3.14	Percentage difference in maximum CFC temperature between the monolithic approximation and the gap inclusive model.	23
3.15	Percentage difference in maximum vessel temperature between the monolithic approximation and the gap inclusive model.	23
3.16	Initial tetrahedral mesh in <i>Ansys Mechanical</i>	25
3.17	Converted polyhedral mesh in <i>Ansys Fluent</i> Solution mode.	25
3.18	BoIs used to refine the mesh around the nozzle area.	26
3.19	BoIs used at the impingement zone between nozzles and dump surface.	26
3.20	BoIs highlighted in orange added around the support structures.	27
3.21	BoI introduced in the main recirculation region after top duct geometry insertion.	27
3.22	Inlets' locations within the computational domain.	28
3.23	Outlet's location within the computational domain.	28
3.24	Enclosure representing the interior of the shielding and enclosing the entirety of the computational domain.	28
3.25	Visualization of the emissivity values assigned to each component in the thermal model.	29
3.26	Naming and segmentation of the solid dump block sections.	30
3.27	Temperature-dependent evolution of selected thermal properties for the materials used in this study.	32
3.28	Time evolution of key diagnostic quantities, including mass flow rate error, temperature difference between outlet and inlet, and total pressure drop across the domain.	35
3.29	Reported values of k_z	36
3.30	Extrapolated values of k_z	36
3.31	Maximum temperatures for core and vessel for the 3 simulated k_z	37
3.32	Residual total internal energy within the dump block for the 3 simulated k_z	37
3.33	Considered derived temperature dependent k_{xy} curves.	37
3.34	Maximum temperatures for core and vessel for different CFC in plane conductivity values.	38
3.35	Residual total internal energy within the dump block.	38
3.36	y^+ values at the wall of the 4 inflation layers fluid mesh.	38

3.37	y+ values at the first cell near the wall of the extremely refined boundary layer. .	39
3.38	y+ values at the first cell near the wall of the coarse boundary layer.	39
3.39	Maximum temperatures for core and vessel for the different boundary layer configurations.	40
3.40	Residual total internal energy within the dump block.	40
3.41	Residual internal energy in the highest energy core segment of the dump block and in the vessel over three consecutive dump events. Solid lines represent results using an extremely refined near-wall fluid mesh while dotted lines show results with only four boundary layers.	40
3.42	Internal energy percentage difference between y+=1 mesh and the 4 inflation layers mesh.	41
4.1	LD top and HD 1 top locations used for thermal analysis, shown at their physical locations on the vessel.	44
4.2	Vessel/core temperature evolution for 2-hour spaced dumps at full bunch intensity (1.8×10^{11} protons).	45
4.3	Vessel/core temperature evolution for 2-hour spaced dumps at reduced beam intensity (1.6×10^{11} protons).	45
4.4	Vessel/core temperature evolution under constant 2-hour repetition rate with 1.2×10^{11} protons bunch intensity.	46
4.5	Vessel/core temperature evolution under fixed bunch intensity (1.6×10^{11} protons) with 3h repetition rate.	46
4.6	Thermal response under partial beam sweep (left-side only).	47
4.7	Hotspot analysis and temperature field within the end of the low density section and the start of the HD 1 block in the partial sweep scenario.	47
4.8	Thermal and energy extraction comparison of nine beam dump cooling configurations/upgrades. Each configuration is evaluated based on the peak and average temperatures in core and vessel, shown as horizontal bars, as well as black diamond markers that indicate the % of beam's stored energy extracted per dump averaged over three beam dump events.	48
4.9	Duct schematic showing the presence of top ducts in the model.	50
4.10	Time evolution of maximum vessel (and core) temperatures across three consecutive dumps for the 1x flowrate configuration, comparing baseline and top duct cases.	51
4.11	Time evolution of maximum vessel (and core) temperatures across three consecutive dumps for the 2x flowrate configuration, comparing baseline and top duct cases.	51
4.12	Dump assembly showing the generated central optimized bottom cooling ducts. .	52
4.13	Dump assembly showing the generated central optimized bottom cooling ducts coupled with the addition of top ducts upgrade.	53
4.14	Duct schematic showing the presence of the developed practical top ducts in the model.	54
4.15	HL dump geometry with current operational duct configuration.	55
4.16	HL dump geometry with top duct cooling configuration.	55
4.17	Longitudinal temperature contour cut through the dump assembly highlighting the artificially elevated temperatures in the region corresponding to the cradle (circled).	57

4.18	Thermal and energy extraction comparison of the 4 beam dump cooling configurations/upgrades. Each configuration is evaluated based on the peak and average temperatures in the core and vessel, shown as horizontal bars, as well as black diamond markers that indicate the % of beam's stored energy extracted per dump averaged over three beam dump events.	58
4.19	Heat flux vector field of a longitudinal cut plane of the LD sector in Top Ducts 1x configuration at peak vessel temperature timestep.	59
4.20	Heat flux vector field of a longitudinal cut plane of the LD sector in Operational 2x configuration at peak vessel temperature timestep.	59
4.21	HTC contour on the bottom surface. Top Ducts 1x configuration.	59
4.22	HTC contour on the bottom surface. Operational 1x configuration.	59
4.23	HTC contour on the top surface. Top Ducts 1x configuration.	60
4.24	HTC contour on the top surface. Operational 2x configuration.	60
4.25	Residual internal energy in HD 1 and vessel. Top Ducts 1x vs. Operational 1x (dashed).	60
4.26	Residual internal energy in HD 1 and vessel. Operational 2x vs. Operational 1x (dashed).	61
4.27	Time evolution of the heat removal rate by the airflow, along with the cumulative energy extracted by air and lost through radiation. Top Ducts 1x vs Operational 1x (dotted).	62
4.28	Time evolution of the heat removal rate by the airflow, along with the cumulative energy extracted by air and lost through radiation. Operational 2x vs. Operational 1x (dotted).	62
4.29	Time evolution of the heat removal rate by the airflow, along with the cumulative energy extracted by air and lost through radiation. Top Ducts 2x vs. Operational 1x (dotted).	62
4.30	Total cumulative extracted energy over time. Top Ducts 1x vs. Operational 1x (dotted).	63
4.31	Total cumulative extracted energy over time. Operational 2x vs. Operational 1x (dotted).	63
4.32	Total cumulative extracted energy over time. Top Ducts 2x vs. Operational 1x (dotted).	63
4.33	Standard 3 dumps cycle highlighting reference points A, B and C.	64
4.34	Detail on evaluation point D just after the energy deposition of the 3rd dump event.	64
4.35	Side view of the contact pressure contour map for the CFC LD sector in the Operational 1x configuration at point C. The surface was reconstructed from point cloud data using a Delaunay 3D algorithm, which introduced artifacts (grey regions and fictitious high values of contact pressure) in the gaps between stacked CFC stacks. Black areas indicate instead regions with contact pressures at or near 0 MPa.	65
4.36	Side view of the contact pressure contour map for the CFC LD sector in the Top Ducts 1x configuration at point C. On top of the grey area artifacts introduced by Delaunay 3D, the black area near CFC 2 corresponds to the modeling artifact discussed in Section 4.3.2, not to an actual loss of contact pressure.	65
4.37	Side view of the contact pressure contour map for the CFC LD sector in the Operational 2x configuration at point C. The contour is visibly improved compared to the 1x case, with no regions at or near 0 MPa (except model artifact).	65
4.38	2D box plots of contact pressure versus temperature for the Operational 1x configuration at point C. The overlaid curve represents the fitted $P(T)$ relationship, with the corresponding equation reported.	66

4.39	2D box plots of contact pressure versus temperature for the Top Duct 1x configuration. Data are matched point wise between the 2 maps and grouped by dump sector.	66
4.40	Maximum vessel and core temperatures over time for the Operational 1x configuration, evaluated under three different thermal contact assumptions: perfect contact, floating core (no contact) and a variable TCC(T) derived from the elastic contact model.	68
4.41	Dump block residual internal energy over time for the Operational 1x configuration, under the same three thermal contact assumptions as in Figure 4.40.	68
4.42	Predicted maximum allowed axial acceleration for xy slice of LD core $\overline{a_{\max}}(z)$ (blue) and weighted polynomial fit (orange) compared to actual peak axial acceleration (green) for the Operational 1x configuration.	70
4.43	Predicted maximum allowed axial acceleration for xy slice of core $\overline{a_{\max}}(z)$ (blue) and weighted polynomial fit (orange) compared to actual peak axial acceleration (green) for the Top Ducts 1x configuration.	70
4.44	Time evolution of maximum core and vessel temperatures for different massflow distribution scenarios. Thermal performance improves as more flow is directed to the top ducts, with both core and vessel temperatures decreasing progressively.	72
4.45	Residual internal energy in the dump block over time for different top/bottom duct flow splits.	72
A.1	Simplified fluid-only duct geometry used to estimate massflow distribution under symmetry assumption. Inlets highlighted in orange.	78
A.2	Updated geometry including preliminary upstream ducting for refined massflow predictions.	79
A.3	Swirling phenomena observable within the ducts.	80
A.4	CAD representation of the CFC core (pink) and the surrounding dummy layers (green) added to implement temperature-dependent thermal contact behavior.	81
A.5	Corresponding solid mesh of one block in the model, showing the interface dummy layers (green) between CFC (red) and the titanium vessel.	81
B.1	Maximum vessel temperature vs. nozzle radii with fitted quadratic regression.	83
B.2	Top view of the duct assembly showing sectioning lines (in red) used to define regions for localized nozzle density augmentation.	84
B.3	Percentage difference in maximum vessel temperature between uniform baseline (T_1) and 50% increase in central nozzle density (T_2).	85
B.4	Progressively inclined nozzle configuration generated using the adapted parametric script.	85
B.5	Velocity contours showing persistent downstream flow concentration despite nozzle inclination.	86
B.6	Geometry of the obstructed duct configuration, showing 50% area reductions at two locations aligned with the cradle positions.	86
B.7	Velocity magnitude contours at the upstream obstruction, illustrating localized pressure build-up and partial acceleration at nozzle exits.	87
B.8	Velocity magnitude contours at the downstream obstruction.	87

List of Tables

1.1	Absorption percentages for peak performance parameters in Run 3 (stored beam Energy: 539 MJ).	2
1.2	Beam parameters over the years [6]. Each LHC beam is made of many tightly packed groups of protons called <i>bunches</i> . The <i>bunch intensity</i> , given in protons (p) or protons per bunch (ppb), refers to the number of protons contained in each of these groups.	3
2.1	Material properties for SS 318LN (AISI 318LN EN 1.4462 / X2CrNiMoN22-5-3) and Titanium Grade 5 (ASTM B381 Grade F-5 / UNS R56400).	13
3.1	Assumptions used in both the elastic and plastic asperity deformation contact model. Contact pressure magnitude assumption was based on interference calculations and material properties.	24
3.2	Comparison of FLUKA energy per pulse and resulting FLUENT internal energy across dump block solid sections.	30
3.3	Assumed constant density of the considered materials.	32
3.4	Parameters of the boundary layer (BL) mesh used in the mesh sensitivity analysis. All simulations include a boundary layer mesh with a fixed growth rate of 1.2. . .	39
3.5	Comparison of FLUKA energy per pulse and resulting FLUENT internal energy across dump block solid sections for different solid domain meshes. For solid section naming: Figure 3.26.	41
4.1	Naming, bunch intensities and frequencies of the different considered scenarios. .	43
4.2	Overview of the tested airflow distribution scenarios for the massflow sensitivity study, showing the assigned massflows to bottom and top ducts and the resulting pressure drops.	71
A.1	Massflow distribution at each duct outlet in the upstream-inclusive simulation. .	79
A.2	Summary of duct surface roughness cases and their impact on total pressure drop.	80
B.1	Summary of thermal performance across nozzle radii and dump events.	88
B.2	Percentage differences (calculated based on temperature values in Kelvin) between baseline and augmented nozzle density configuration for key thermal metrics over three consecutive dump events. Only highest deviations are reported for readability.	89
B.3	Percentage differences (calculated based on temperature values in Kelvin) between baseline and inclined nozzle configuration for key thermal metrics over three consecutive dump events. Only highest deviations are reported for readability. . .	89
B.4	Percentage differences (calculated based on temperature values in Kelvin) between baseline and obstructed ducts configurations for key thermal metrics over three consecutive dump events. Only highest deviations are reported for readability.	89

B.5	Key temperature and energy metrics for the 2x flowrate central bottom cooling configuration , reported as percentage differences with respect to the 2x shrinkfitted baseline. The table highlights how the configuration becomes progressively more effective across consecutive dumps. Only the most significant values are shown for clarity.	90
-----	---	----

Chapter 1

Introduction

539 MJ: the kinetic energy of a Train à Grande Vitesses (TGV) traveling at 192 km/h with a mass of 380 tons. Remarkably, it is also the maximum energy stored in a single beam circulating within the Large Hadron Collider (LHC) at the European Organization for Nuclear Research (CERN).

The LHC, which first began operation in 2008 [1], is the largest particle accelerator in the world. In order to achieve such high levels of stored energy, protons are accelerated to speeds approaching the speed of light within two ultra-high-vacuum tubes, each 27 kilometers in circumference. The beams are guided and kept in circular motion by powerful superconducting magnets and accelerated by radiofrequency cavities [2]. These beams, rotating in clockwise and counterclockwise directions, are finally collided at four distinct locations along the accelerator's circumference. Large detectors are installed in these position to capture the results of high energy particle interactions.




Figure 1.1: Final installation of the Run 3 beam dump assembly, shown with the complete cooling system and instrumentation in place, enclosed within the shielding structure.

Specialized protection devices are essential for the safe operation of the LHC [3], particularly those that ensure the controlled absorption of the beam. To safely dissipate the large amounts of energy carried by the high intensity beams before each new fill [4], dedicated beam dumps are used. For this purpose, two Target Dump for Ejected beam (TDE) assemblies, one per counter-rotating beam, are installed in the LHC (Figure 1.1).




Figure 1.2: Schematic representation of LHC underground point 6 where the two caverns with dead ends highlighted in red are present [5].

During operation, if the beam has to be dumped for any reason, each beam is redirected by kicker magnets into a transfer line that ends in a dead-end equipped with a beam dump device as can be observed in Figure 1.2. The particle beam, after being diluted in a sweep pattern (Figure 1.3) to reduce peak energy density (amount of energy deposited per unit mass), collides with the dump producing a multitude of secondary particles and forming a particle shower. These secondary particles are subsequently absorbed in diverse amounts by the dump and its surrounding environment, which includes concrete shielding, instrumentation, support mechanisms, the cooling system etc. (Table 1.1 from [6]).

This process effectively transforms the kinetic energy of the particles into thermal energy, which is managed by the TDE and dissipated by its cooling system.

	Fraction of Beam Energy	Energy Deposition (MJ)
Dump:		
Graphite	73.6%	397
Shell	4.2%	23
Windows, Flanges, etc.	0.04%	0.2
Total	77.9%	420
Environment:		
Shielding	17.4%	94
Air	0.015%	0.08
Cavern	0.12%	0.6
Molasse, Rock, etc.	0.04%	0.2
Total	17.6%	95

Table 1.1: Absorption percentages for peak performance parameters in Run 3 (stored beam Energy: 539 MJ).




Figure 1.3: Transverse section of the beam dump where the sweep pattern of the beam is observable. Beam energy: 7 TeV [7].

With the upcoming High-Luminosity (HL) LHC upgrades, the demands on the LHC's infrastructure will increase significantly. This upgrade aims to deliver up to ten times the integrated luminosity of its predecessor [8]. As part of the upgrade, the stored energy in the beams will increase by more than 30% compared to the previous run, resulting in increased thermal loads on the beam dump system (Table 1.2). Additionally, the functional specification for the HL TDE explicitly states that *"the dumps must withstand these operational conditions even in case of several consecutive high intensity dumps"* [9]. This requirement justifies evaluating the impact of repeated beam dumps at an increased frequency, specifically, one every 2.5 hours. This reflects a realistic operational cycle in which a beam may be re-injected, ramped up, and subsequently dumped again within that timeframe. The *emergency scenario* assumes up to three such full-intensity dumps in succession, representing the most demanding thermal conditions expected during HL-LHC operation. While such a sequence is unlikely to occur often, it defines a conservative upper bound for temperature rise within the dump assembly. Finally, re-injection should only be permitted if the dump is in a safe thermal state and can reliably absorb another beam at any moment. Ensuring that temperatures reduce sufficiently between dump events is therefore critical and the design must withstand this worst-case repeated dumps scenario without risking material degradation or mechanical failure.

	Run 1 (2009-2013)	Run 2 (2015-2018)	Run 3 (2022-)	High Luminosity
Proton Energy	4 TeV	6.5 TeV	6.8 TeV	7 TeV
Number of Bunches	1380	2556	2748	2760
Bunch Intensity	1.7e11 (p)	1.2e11 (p)	1.8e11 (p)	2.3e11 (p)
Beam Energy	150 MJ	320 MJ	539 MJ	712 MJ

Table 1.2: Beam parameters over the years [6]. Each LHC beam is made of many tightly packed groups of protons called *bunches*. The *bunch intensity*, given in protons (*p*) or protons per bunch (*ppb*), refers to the number of protons contained in each of these groups.

Addressing these various challenges has been the focus of several previous studies, which have investigated key aspects of beam dump design and cooling system performance. In particular,

thermal computational fluid dynamic (CFD) modeling approaches have been used to identify factors influencing heat dissipation, assess material behavior under different beam conditions and evaluate potential upgrades to the cooling infrastructure to meet increased operational demands. These contributions offer valuable insights that are discussed in detail in the literature review (Section 2).

Building on this foundation, the present thesis aims to contribute with a practical and enhanced CFD model of the HL-LHC beam dump cooling system, tailored to the latest TDE design and with the ultimate goal to assess the performance of potential upgrades. The model integrates improved meshing strategies, thermal contact resistance representations and a flexible architecture that facilitates future design modifications.

Chapter 2

Literature Review

This literature review serves as an introduction to the topic, providing both a framework for the analysis presented in the following sections and highlighting the need to investigate the TDE cooling system in preparation for HL operations. It includes insights into the research and development of beam intercepting devices, and especially the TDE, examining the materials, configurations and innovations developed over the years. The review also investigates the historical upgrades of the TDE, focusing on the constant upgrading necessary to continuously meet the increasing demands in term of beam intensity and energy deposition. Finally the attention is finally shifted on the multifaceted problem that the TDE represents by highlighting the relation between energy deposition profiles, mechanical behavior of the system and its cooling.

2.1 Beam Intercepting Devices

Beam intercepting devices interact with high energy particle beams in various ways, including energy absorption for beam termination, beam shaping for cleaning and control, and the production of secondary particles. They are necessary for both scientific research purposes and for safe operation of the accelerator from both the human and machine protection standpoints. This macro-category of devices can be further divided into three different subcategories based on their function [10]:

- **Beam dumps and stoppers:**
Designed to safely absorb high energy particle beams in nominal operation and emergency scenarios. They are usually designed with dense materials, complex instrumentation and sophisticated cooling systems.
- **Collimators, strippers and scrapers:**
They serve mainly as beam control devices in order to achieve clean beams, more focused and/or with specific charge state. This allows for modification of the beam characteristic and/or removal of beam halo in order to protect downstream equipment.
- **Targets:**
Focused on the physics' scientific research function, these devices interact with the primary particle beam to produce secondary particles for dedicated experiments. They also require robust designs, materials with high thermal conductivity, mechanical strengths and radiation resistance as they are exposed to high intensity beams with high energy.

2.1.1 Challenges

Designing a beam intercepting device has multiple technical and operational challenges. These devices must withstand both routine and accidental load conditions, enduring extreme thermal mechanical stress levels in high radiation environments. As they are often the most radioactive components in a particle accelerator, they must be designed considering the entirety of their life cycle, from production and installation to dismantling. For example, in some internal beam dump applications, particularly those located in ultra-high vacuum (UHV) sections, additional design constraints arise from stringent vacuum requirements [11]. This significantly restricts material selection and assembly techniques, as even minor outgassing or contamination could influence the accelerator’s vacuum integrity. Other types of beam intercepting devices instead introduce additional layers of complexity. For example, collimators often include movable components that must achieve and maintain extremely precise positioning in order to function correctly. Targets, on the other hand, must meet strict physics-driven requirements that often dictate the use of materials selected for their nuclear interaction properties rather than for their thermal or mechanical performance, further complicating the design process. Balancing these multifaceted requirements is mandatory to ensure both the device’s reliability and its performance throughout its lifetime. This is not always easy since functional requirements for such devices are often incomplete or evolve during the development phase, increasing the complexity of the problem [10], and, as it is often the case, uncertainty demands a flexible and iterative approach where simulations and prototyping are vital for the design phase.

For instance, understanding beam-matter interaction through advanced simulations is crucial for predicting the distribution of energy deposition within a device. This contributes to better understand the effects of the beam on materials, and on the assembly, from both the thermo-mechanical standpoint and from the point of view of residual activation of the materials (section 2.2.1.1). Energy deposition studies are primarily conducted using the Monte Carlo code FLUKA [12]. Already in the early efforts to design a beam dump for the LHC, such as those documented in LHC Project Report 80 [13], the focus was placed on evaluating the thermal behavior of core materials under beam impact. These analyses also identified the need for implementing a beam sweep pattern to distribute energy more evenly and reduce localized heating. Later studies expanded on this foundation by incorporating finite element analysis to model the transient thermal response of the system [14]. These efforts demonstrated the importance of coupling energy deposition simulations with detailed thermo-mechanical and fluid dynamic analyses, a multidisciplinary approach that was already taking shape before the LHC became operational.

To summarize, the beam intercepting device design process is highly iterative: changes in geometry or material properties can significantly modify the energy deposition profiles, which in turn affects mechanical stresses and thermal loads. While cooling configurations have limited influence on energy deposition itself, they play a critical role in determining the system’s ability to extract heat efficiently. Each iteration refines therefore the system by balancing these interconnected factors leading progressively towards an optimized design.

2.2 Target Dump for Ejected beam

The TDE is a specialized beam-intercepting device designed to safely absorb high energy LHC beams during both routine operation and accidental extraction scenarios. Unlike internal dumps or collimators, the TDE operates outside the UHV envelope, allowing for greater design flexibility and fewer constraints.

Each TDE is installed in a dedicated underground cavern located approximately 600 meters downstream from the main ring on either side of Point 6 (Figure 1.2). The beam is first extracted from the LHC ring by the extraction kickers into the tangential extraction line where it is further deflected by the septum field. Finally the dilution kickers sweep the beam on the surface of the upstream window of the TDE [2] (Figure: 2.1).

The dump block itself is a cylindrical assembly weighing 6.2 tonnes and measuring about 8.5 meters in length. It consists of stainless-steel blanks welded together to form a vessel that houses sections of high- and low-density graphite blocks and sheets (Figure 2.2). This internal structure is arranged to reduce peak energy deposition, proportional to material density, and consequently limit peak temperatures within the dump core. High density blocks are shrink-fitted into the vessel, while low-density sheets are loosely installed due to their flexibility and maintained into position by extruded graphite disks.




Figure 2.1: Schematic of the extraction system and detail on the sweep pattern projected on the upstream window from [9]. For a detailed description of the extraction and dilution systems, the reader is referred to the introduction of [9].




Figure 2.2: Longitudinal cut plane showing the graphite positioning and densities [7].

While the core geometry and functional layout of the TDE dump block remained largely unchanged throughout Run 1 and Run 2, the system as a whole underwent several upgrades to meet the increasing demands of beam intensity and energy (Table 1.2, Figure 2.3). The most substantial improvements were implemented during Long Shutdown 2 (LS2), a planned accelerator-wide maintenance and upgrade period, that took place between 2019 and 2021, to address reliability and operational issues that had emerged during earlier operation.




Figure 2.3: Simulated energy density profiles along the TDE beam dump length for different LHC runs, assuming a nominal beam sweep pattern [6]. Energy density [J/g] is the amount of energy from the proton beam that is deposited into the dump material at each point and its peak indicates the physical location within the assembly where the most intense heating occurs.

2.2.1 LS2 Studies & Upgrades

Before LS2, the TDE was connected to the beam extraction line through a connecting line (Figure 2.4 & 2.5). UHV windows acted as divisions between the vacuum of the beam pipe and the nitrogen environment inside the TDE. Furthermore chain clamps and metallic gaskets, to secure a leak-proof connection between the aforementioned components, and bellows, to maintain flexibility, were employed. Finally the dumps were mechanically supported within the shielding by steel frames and held in place solely by frictional forces, without any active restraint system [7].




Figure 2.4: Cross section of the dump installation until the end of Run 2. The blue connecting line can be observed connecting the orange extracting line and the beam dump block. (From [15]).



Figure 2.5: Picture showing the dump block within its shielding and the connecting line connected to it, from [15].

Operational challenges arose during LHC Run 2 as nitrogen leakage was observed at the dump interfaces. This was eventually linked to damage in the sealing system caused by mechanical vibrations and transient displacements of the dump block under beam impact. The original design did not account for this dynamic behavior, which resulted in extreme mechanical loads induced by beam deposition as the beam-induced thermal expansion and rapid pressure wave propagation led to accelerations exceeding 200 g [5]. These issues signaled the need for upgrades, which were addressed in the upgrades carried out during LS2.

2.2.1.1 Energy Deposition and FLUKA Monte Carlo Simulations

Summarized in the Engineering Design Report for the LHC External Beam Dump [5], studies highlighted that approximately 96% of the total beam energy is deposited in the beam dump assembly and surrounding shielding. Furthermore the graphite core absorbs the majority of this energy, with localized peak temperatures reaching up to 1500°C in Run 3 and projected to exceed 1850°C for HL-LHC operation (Chapter 4 in [5]), initiating new research concerning HL upgrades (section 2.2.2). Additionally, in the worst-case scenario of a failure involving two horizontal dilution kicker magnets (MKBH) during HL-LHC operation (Figure 2.6), peak temperatures were predicted to rise to as high as 2660°C, raising concerns about potential material degradation [16]. However, subsequent investigations conducted in the HiRadMat-65 experiment confirmed that such temperatures do not lead to critical damage, mitigating those initial concerns (Section 2.2.2.2).




Figure 2.6: Left: peak temperature profiles along the length of the dump block in case of MKBH malfunction. Right: transversal cut in the maximum temperature's location showing the modified sweep pattern in case of 2 MKBH malfunction [5].

Based on these results critical design upgrades were decided:

- New capacitors to reduce the risk of flashover (unintended electrical discharges within the kicker magnet system) and consequent worst case scenario event.
- New consequent dilution sweep patterns were introduced in Run 3.

In parallel to these energy deposition studies, additional investigations focused on the thermo-mechanical behavior of the dump assembly under the newly calculated physical loads. These assessments, presented in Section 2.2.1.2, played a key role in justifying structural upgrades by ensuring that the dump could safely withstand the thermal stresses and mechanical deformation expected during Run 3 operation.

2.2.1.2 Thermo-Mechanical Studies Driven Upgrades

Using Finite-Element Analysis (FEA), the simulation campaign helped to understand and provided critical insights into the structural response of the dump. This led to specific design modifications to improve the reliability of the device during operation with higher stored energy foreseen for Run 3.

First of all, over-constraining by the existing supports, and excessive stress transmission to the sealing interfaces and connection elements were identified as contributing factors to the nitrogen leaks. An enhanced support system was therefore developed: specifically, a wire-rope suspension system (Figure 2.7) now decouples the dump block's movement from its support frame, eliminating over-constraint. Additionally, the removal of the rigid connection line between the dump and the upstream beam extracting line prevents force transmission to upstream components during beam impact events.



Figure 2.7: Run 3 dump block installed in its new support system.

Removing the connection line between the dump block and the beamline isolates the beamline from any movement of the dump, thereby reducing the risk of leaks at that interface. As part of this redesign, the beam windows also required replacement. The upstream window in particular had previously been shared with the connection line and its integration posed a risk of leakage. Furthermore, thermo-mechanical assessments showed that the existing windows were not expected to withstand the increased thermal and mechanical loads foreseen during Run 3. For these reasons, new windows made of titanium grade 5 (Ti-6Al-4V or Ti Gr5) were implemented, offering significantly improved resistance to high intensity beam impacts. The upgraded dumps also featured monitoring instrumentation. These include: 16 temperature sensors, 13 strain gauges, three Linear Variable Displacement Transducer (LVDT), a Laser Doppler Vibrometers (LDV) and three accelerometers. These are crucial to support operational monitoring, by providing real-time response data, and to benchmark the FE models that are developed to simulate the dump's response. While these upgrades addressed known integration and monitoring challenges, further insights emerged from post-irradiation examinations (PIE) conducted after the decommissioning of the Run 2 dumps [15]. These investigations revealed structural degradation in some components of the dump core, particularly in carbon-based materials. For instance, extruded graphite disks were found to have cracked during operation (Figure 2.8), compromising their role in containing and compressing the low-density sector. In addition, surface alterations were observed in the isostatic graphite discs, attributed to beam impact. These findings highlight the importance of continued material-specific studies and validate the relevance of thermo-mechanical simulation tools in predicting such failure modes. While not all issues revealed by PIE were addressed during LS2, the results underscore the need for further design evolution and material qualification ahead of HL-LHC operation.




Figure 2.8:
Crack in the extruded graphite disk develops from outgassing groove.

2.2.2 High Luminosity Studies

With stored beam energy reaching approximately 700 MJ, thermal loads within the dump core and vessel will exceed those observed in Run 3. FLUKA Monte Carlo simulations with HL beam parameters developed in [6] predict that, under nominal operations, peak energy densities in the graphite core will reach approximately 3400 J/g, corresponding to a temperature rise of around 2000°C. However, in failure scenarios, such as the design case of two horizontal dilution kickers missing, energy densities could rise to 5700 J/g reaching core temperatures above 3200°C. These extreme conditions require careful material evaluation to ensure their compliance with the ten years of operation that are currently foreseen for the HL era [17].

In addition to the challenges within the core, Banks et al. [18] highlight the consequences of the increased loading. The existing stainless-steel (SS 318LN) vessel, used in the operational dumps in Run 3, is expected to experience higher energy deposition levels under HL-LHC conditions, further increasing the mechanical stress and vibrational issues.

2.2.2.1 Structural Integrity and Vessel Material Developments

Thermo-mechanical assessments by Banks et al. [18], based on FLUKA energy deposition maps under HL-LHC beam parameters, revealed increased thermal loads and vessel stresses. Additionally, structural concerns were raised due to a beam-induced longitudinal fast vibrational mode at approximately 200 Hz. The low-density sector, in particular, experiences the highest stress levels (Figure 2.9), as it coincides with both the center of the longitudinal vibrational mode and the region of maximum energy deposition within the vessel. Consequently, the predicted peak stresses during HL-LHC operation approach, or in some scenarios exceed, the yield strength of SS 318LN at 150°C. This led to the conclusion that SS 318LN may not provide sufficient structural margin over the dump’s operational lifetime, prompting further material and design studies to identify more suitable alternatives for HL-LHC conditions.




Figure 2.9: Evolution of the maximum vessel stress in the low density section for both SS 318LN (Red) and Ti Gr5 (Blue) with HL beam conditions and their respective yield strength.

Among the materials evaluated, Ti Gr5 has emerged as the leading candidate [18]. Compared to SS 318LN, Ti Gr5 offers significantly lower thermal expansion, higher yield strength, and reduced interaction with secondary particles, primarily due to its lower density, which effectively lowers energy deposition in the vessel walls (Table 2.1). As a result, replacing SS 318LN with Ti Gr5 would lead to the following improvements:

- a 48% reduction in peak longitudinal extension,

- a more than 56% decrease in peak stresses in critical vessel regions,
- stress levels throughout the vessel remaining well below the estimated titanium fatigue limit of 500 MPa [19].

While the operational conditions are not directly comparable, it is worth noting that Ti Gr5 has already been successfully used for beam dump windows since the LS2 upgrades (Chapter 2.2.1.2). Although the geometry and loading scheme differ significantly from that of the vessel, this prior experience demonstrates the material’s proven reliability under extreme cyclic and thermal loading in high energy physics environments.

	Density [g/cm ³]	Coefficient of Thermal Expansion @ 25°C [10 ⁻⁶ /K]	Young’s Modulus @ 200°C [GPa]	Yield Stress @ 25°C [MPa]
SS 318LN	7.8	13.0	186	450
Ti Gr5	4.43	2.7	113	828

Table 2.1: Material properties for SS 318LN (AISI 318LN EN 1.4462 / X2CrNiMoN22-5-3) and Titanium Grade 5 (ASTM B381 Grade F-5 / UNS R56400).

While the properties of Ti Gr5 seem particularly suitable for this application, the adoption of this material introduces new engineering challenges, particularly in fabrication and assembly. Previous studies [20] have highlighted concerns regarding stress concentrations at weld junctions, particularly in designs employing partial penetration welds. To address this, ongoing research is investigating full penetration welds, with initial tests of electron beam welding (EBW) already conducted on planar plates. At the same time, while studies on welding performance and dismantling procedures are still in the pipeline, confidence has been established regarding the procurement, machining and shrink-fitting of titanium blanks, demonstrating the material’s feasibility for large scale prototyping.

2.2.2.2 Graphitic Material Evolution: from Sigraflex[®] to CFC

With the transition to Ti Gr5 now confirmed, parallel R&D on graphitic materials continued, combining simulations and dedicated experiments to identify materials capable of withstanding HL-LHC energy densities.

Initial studies evaluated carbon based materials under relevant energy deposition conditions. The HiRadMat-56 experiment [21] tested several candidate materials at energy densities up to 3.2 kJ/g, with up to 4 beam impacts per specimen to investigate operational conditions simulating realistic beam impact conditions. While stacked Sigraflex[®] sheets survived, individual sheets exhibited out-of-plane deformations and previous tests (HiRadMat-43 [22]) had already revealed severe delamination issues. In contrast, all tested Carbon Reinforced Carbon (CFC) materials showed promising performance, with no observable damage.

A subsequent HiRadMat-65 experiment was conducted to test CFC, Sigraflex[®] and high density isostatic graphite under different boundary conditions (stacked and free-face). Specimens were exposed to their respective nominal (50 beam impacts) and accidental (single beam impact) energy densities. The results, though unpublished at the time of writing, confirmed that all tested CFC and high density isostatic graphite specimens survived the highest intensities. All tested CFC grades were therefore qualified for possible use in the next HL TDE design except one that did not due to manufacturability limitations. On the other hand, Sigraflex[®] showed extensive damage even at the lowest tested energy density (2.4 kJ/g), ruling it out for HL TDE use and confirming the shift toward CFC and high density graphite as the most promising materials for the next-generation beam dump core.

2.2.3 Cooling System




Figure 2.10: Picture of the nozzles along the longitudinal ducts running below the dump structure block.




Figure 2.11: Schematic of impingement jet flow structure [23].

The TDE cooling system, designed during the early phases of LHC development, has remained largely unchanged throughout its various operational runs. Although the original conceptual design considered a water-based cooling system [13], this approach was eventually discarded primarily due to concerns over potential radioactive water leakage in case of failure. The implemented system relies instead on air impingement cooling (Figure 2.11): nozzles direct high velocity air jets onto the underside of the dump vessel to extract heat by forced convection (Figure 2.10).

The system operates with an air handling unit (AHU) providing $5000 \text{ m}^3/\text{h}$ of airflow, that are forced inside external duct work and distributed through two cooling ducts equipped with nozzles (Figures 2.12 & 2.13).




Figure 2.12: Schematic representation of the AHU and the external duct work conducting the $5000 \text{ m}^3/\text{h}$ of air to the bottom ducts.




Figure 2.13: Forced convection on the bottom side of the dump.

This design was operational during Run 1, Run 2 and Run 3, but as HL conditions introduce 30% higher total deposited energy, it had to be reconsidered. Recent studies [24] developed a CFD model of the Run 3 dump using *Ansys Fluent*, analyzing the thermal response of the system under HL operational conditions and exploring preliminary upgrade concepts. The analysis confirmed that the cooling system effectively extracted heat from the vessel, maintaining safe temperatures after a single dump event. However, it also highlighted significant inhomogeneities

in cooling distribution, particularly between the upper and lower vessel surfaces. A key limitation was identified in the air distribution within the cooling ducts, where pressure distributions along the ducts caused higher airflow rates at downstream nozzles compared to upstream ones (observable in Figure 2.14). This resulted in uneven cooling along the vessel's length, contributing to localized thermal gradients and highlighting an imbalance in the cooling system where some critical areas might not receive sufficient cooling while less critical regions experience excessive cooling flow.




Figure 2.14: Air velocity distribution in the cooling ducts. The figure shows two cross-sectional velocity fields at different duct locations alongside the full 3D geometry of the duct assembly with indicated slice positions. Higher air velocities are observed at the downstream nozzles, resulting from pressure gradients along the duct. This uneven distribution highlights the influence of pressure drop on airflow delivery [24].

Although the existing cooling system was experimentally validated for Run 3 conditions, further analysis revealed operational constraints even under current conditions, hinting at areas where cooling efficiency could be improved (Section 4.1). At the same time, predictions in previous studies for HL operation [24] revealed additional thermal limitations, showing that the cooling in its current configuration would not be sufficient to handle the increased heat loads given by HL beam parameters for consecutive dump events. The asymmetric cooling distribution, caused by impingement jets only targeting the lower vessel surface, leave the upper surface significantly less cooled, leading to temperature differentials that could contribute to mechanical stress and loss of thermal contacts. Additionally, the author concluded that the higher beam energy and more frequent dump cycles foreseen for the HL-LHC would intensify thermal accumulation, further extending cooldown times and raising concerns about heat buildup during consecutive dump events. Since the existing cooling system was not designed to deal with these more demanding conditions, it became clear that simply reusing the Run 3 setup would not be sufficient. While energy deposition and material studies focus on how the dump absorbs and resists high intensity beams, a dedicated CFD evaluation of the cooling system was essential to ensure effective heat removal and reliable operation under HL-LHC conditions.

To address these concerns, the aforementioned studies [24] (still under review for internal approval at the time of writing) also explored potential upgrades to the cooling system. This was not yet centered on a finalized HL beam dump design but rather on identifying general cooling system

optimization strategies that could guide future improvements. One of the primary considerations was increasing the airflow capacity, with simulations exploring the effect of doubling the airflow from 5'000 m³/h to 10'000 m³/h. While this was shown to enhance convective heat transfer and surface heat transfer coefficient, pressure drop limitations in the AHU made this solution impossible without additional ventilation infrastructure modifications. This options was later reassessed for the HL beam dump block design in this current work (Section 4.2.2), where it was found to be acceptable from a cooling efficiency standpoint but ultimately deemed too expensive from a complexity and financial perspective compared to the solution developed here, which will be presented later in the report. Another previously investigated concept was introducing an annular cooling sleeve around the vessel, which showed potential for improving temperature uniformity.




Figure 2.15: Initial geometry for the annular cooling investigation [24].

However, this approach was ultimately discarded due to the significant complexity of integration and the project timeline constraints, which did not allow for its implementation within the required schedule. Finally, alternative low density core materials, such as CFC, were here first considered. Although it enhanced heat transfer efficiency, this approach led to increased vessel temperatures, requiring complementary improvements in external cooling. While the results confirmed that no single solution alone would be sufficient, they identified promising directions for improvement.

Chapter 3

Methods

This section describes the computational framework used in this study, including the computational domain, boundary conditions, geometry, meshing procedure and simulation setup. Figure 3.1 provides a visual overview of the complete workflow, illustrating the key steps performed in this study, from CAD pre-processing through to post-processing and evaluation. Additionally, model checks are presented to ensure accuracy and reliability.

3.1 Model

This subsection details the core aspects of the computational model, including its assumptions, domain definition, boundary conditions, material properties and meshing strategy. The model is designed to be both accurate and repeatable, allowing for efficient updates in response to design changes.

3.1.1 Computational Domain

The computational domain is defined to enclose all relevant components of the system while excluding regions that do not contribute to heat transfer and/or are not relevant to the study. The domain is delimited by a surface representing the interior of the shielding, which physically limits the flow but does not participate in the convective thermal exchange. This surface encloses a modeled representation of the entire assembly, including the beam dump block, the support structures/cables and the cooling system. The model was developed without exploiting any symmetry axis, as none is present due to the sweep pattern observable in Figure 1.3.

The dimensions of the computational domain match the available design space within the shielding (Figure 3.2). The dimensions are as follows and can be observed in Figure 3.3:




Figure 3.1: Detailed simulation workflow.

- 10.22 m length, fully enclosing the entire length of the beam dump,
- 1.14 m height corresponding to the distance between the surface where the longitudinal cooling ducts are installed and the top shielding block,
- 1.04 m width spanning the lateral extent of the system.

This domain ensures that all relevant thermal and flow interactions are accounted for while maintaining as much computational efficiency as possible.




Figure 3.2: 3D model of the dump block assembly, its shielding, the AHU and the cooling ducts.




Figure 3.3: Dimensions of the computational domain

3.1.2 Geometry

At the start of this study, simulations were based on the Run 3 spare dump geometry (Section 4.1 & 4.2), a near-identical variant of the operational Run 3 dumps with only minor differences such as a slightly longer vessel and larger gaps between the core blocks, since the HL-LHC design had not yet been finalized. This choice ensured compatibility with earlier modeling efforts, particularly those reported in [24], where geometric simplifications for operational and spare Run 3 dumps had already been developed and validated.

Once the HL-LHC geometry became available, the study transitioned fully to this updated configuration. All geometry modeling and simplification procedures presented in this section were defined to match the specific layout and requirements of the HL design. The domain itself is therefore divided into two main regions:

- The solid domain that includes only the beam dump block, consisting of the vessel and core. These are the primary components where heat conduction is of interest.
- The fluid domain that includes all regions where fluid can flow, including the interior of cooling ducts and nozzles within the shielding.

Additional components, such as ducts and support structures, are modeled as closed geometric surfaces that define the flow path and act as physical obstructions to the fluid. As their internal thermal behavior is not of interest, they are not represented as solid or fluid regions and are excluded from heat transfer modeling.

3.1.2.1 Fluid Geometry

In this section, the fluid domain of the operational cooling system is described. Other geometric variations explored in this study, such as changes in duct layout or nozzle design, follow the same modeling logic and therefore will not be re-explained in detail.

It is important to note that each geometric variation results in a distinct fluid domain and would ideally require an individual mesh sensitivity analysis. However, due to time constraints, this analysis has been carried out only for the base geometry of the current cooling system. The meshing guidelines derived from this reference case are then applied consistently across all other configurations. While preliminary results were obtained prior to the completion of the mesh sensitivity study, this should not be viewed as a limitation: in each case, mesh parameters were chosen based on validated best practices from previous modeling experience [24], ensuring that early thermal and fluid predictions were sufficiently reliable. Accordingly, the mesh sensitivity analysis presented in Section 3.3.3 is intended to confirm the soundness of the adopted meshing approach, rather than to serve as an exploratory step.

The fluid domain was defined to include all regions within the shielding not occupied by solid components, effectively representing the air-filled space surrounding the TDE and its support structures. A practical modeling strategy is applied throughout this work: objects are subtracted from the fluid domain rather than meshed as solids, with the exception of the dump block itself, which is part of the solid domain. This allows for the representation of impenetrable boundaries without explicitly modeling the solid objects. Specifically simplified representation of the windows are excluded from the fluid domain to act as flow barriers, the support frame is modeled as an obstacle and the duct and nozzles assembly is cut out from the fluid domain to define the inflow-guiding walls that shape the internal airflow. Furthermore different Bodies of Influence (BoIs) have been introduced in order to achieve local mesh refinements (see Section 3.1.4.2).




Figure 3.4: Fluid domain geometry modeled in *Ansys SpaceClaim*, with the outlet highlighted in red.

The fluid domain shown in Figure 3.4 resembles that used in a previous study [24], and the mesh sensitivity analysis performed there remains valid for this setup. Nevertheless, new sensitivity studies, such as those involving boundary layer resolution, will be discussed in Section 3.3.3. Additionally, since the solid domain used in this work is entirely new (see Section 3.1.2.2), a specific mesh sensitivity analysis for the solid domain is also presented in Section 3.3.3.

3.1.2.2 Solid Geometry

The solid geometry of the HL beam dump consists of a titanium vessel housing a core composed of cylindrical isostatic graphite blocks and stacks of round CFC plates. The vessel design features titanium windows, which are bolted directly onto the vessel (observable in Figure 3.5), and is constructed from welded titanium blanks forming the vessel assembly. Inside this vessel, the isostatic graphite blocks (see Figure 3.6) and CFC plates stacks (see Figure 3.7) are shrink-fitted both to prevent relative movement during operation and to maintain good thermal contact with the vessel. Additionally, the vessel is filled with a protective atmosphere of argon at a slight overpressure of 1.15 bar (absolute pressure). This inert environment prevents oxidation of the core materials during high temperature operation.




Figure 3.5: Longitudinal section of the new HL beam dump block highlighting both the HD isostatic graphite (orange), the CFC stacks (blue) and the windows.




Figure 3.6: 3D model of isostatic graphite block geometry.

Figure 3.7: CAD representation of a stack of CFC plates, showing the arrangement of 12 central plates and 2 end plates. The detail on the right highlights the interface between adjacent plates and the small gaps introduced by the geometry.

Given the complexity of the real geometry, both core and vessel were simplified in the model. The detailed vessel was simplified as a hollow cylinder (Figure 3.8), capturing its overall thermal and structural function without unnecessary geometric details. Additionally, since their contribution to heat transfer between the vessel and the cooling flow is negligible in comparison to the vessel, the titanium windows were removed from the solid domain and instead modeled as obstructions in the fluid domain (closed surfaces). This approach maintains accuracy while reducing computational complexity. Regarding core modeling, the high density graphite blocks were modeled as solid cylinders (observable in green in Figure 3.9), preserving the bulk thermal behavior while omitting finer geometrical features such as out-gassing grooves and chamfers, which were previously validated as negligible [24].

Figure 3.8: Solid vessel model divided into multiple sections to enable evaluation of temperatures and internal energies at different locations within the vessel.




Figure 3.9: Core model: isostatic graphite blocks (green) and CFC stacks (pink), separated by cylindrical argon gaps (grey).

Regarding core simplifications, the real design of the LD sector consists of stacked cylindrical CFC plates that are shrink-fitted into the vessel (see Figure 3.7 on the left). These plates include small gaps between them (visible in Figure 3.7, right) to prevent plate migration caused by out-of-plane thermal expansion in the beam-impacted area. Simulations of earlier designs without gaps showed that such expansion could cause the stacks to permanently push themselves apart. This behavior is undesirable due to potential wear at contact surfaces, uncertainty in the core's exact position within the assembly and the resulting unpredictable changes in the energy deposition profile. A sub-modeling approach was therefore used to analyze the thermal impact of these gaps, comparing a monolithic approximation, treating the stack as a continuous solid, and a gap-inclusive model, explicitly accounting for argon-filled spaces between the plates.




Figure 3.10: Section view of submodel representing the CFC stack with argon filled gaps.




Figure 3.11: Section view of submodel representing the CFC stack as a monolithic block.

Both models were computed using a constant convective heat transfer coefficient of $10 \text{ W}/(\text{m}^2 \cdot \text{K})$ applied on the vessel surface. They also shared the same FLUKA energy deposition profile (Section 3.1.5.1), boundary conditions and initial conditions as the only difference between them was

the presence or absence of argon-filled gaps (Figure 3.10 & 3.11).

Figure 3.12: Max CFC temperature for monolithic approximation and gap inclusive model.




Figure 3.13: Max vessel temperature for monolithic approximation and gap inclusive model.





Figure 3.14: Percentage difference in maximum CFC temperature between the monolithic approximation and the gap inclusive model.




Figure 3.15: Percentage difference in maximum vessel temperature between the monolithic approximation and the gap inclusive model.

Figure 3.12 presents the maximum temperature in the CFC core for both models over time. The corresponding percentage difference graph (Figure 3.14) shows a maximum deviation of approximately 5% during the initial energy deposition phase, with the monolithic model exhibiting higher temperatures and indicating that, although the monolithic model slightly overestimates the peak temperatures, the deviation is moderate and progressively diminishes over time.

The second set of graphs shows the maximum temperature in the vessel for both modeling approaches (Figure 3.13), along with the percentage difference (Figure 3.15). Here, the deviation remains within 1.75% and becomes noticeable after the end of energy deposition, particularly during the diffusion-dominated phase from 0.1 h to 0.3 h. Interestingly, the percentage difference then temporarily turns negative, indicating that in the cool down phase the monolithic model is initially more effective. However, this trend reverses after 0.5 h, with both models converging rapidly toward very similar results.

These results show that the monolithic approximation performs similarly to the gap-inclusive model in capturing the thermal behavior of the system. While small deviations exist, they are within an acceptable range. Moreover, the monolithic model tends to slightly overestimate the temperature, which is conservative from a performance evaluation perspective.

Finally, considering the significant simplification it brings to the setup (avoiding the need to model 56 additional interfaces with the vessel and 116 interfaces with the argon gaps), the monolithic approach was selected and the modeled monolithic cylinders can be observed in pink

in Figure 3.9. This allowed for reduced modeling complexity while preserving the accuracy necessary for reliable thermodynamic assessment.

3.1.3 Perfect Contact Assumption and Contact Pressure Calculations

The assumption of perfect thermal contact between components in the core/vessel assembly was evaluated during the model development phase. To support this simplification, both elastic and plastic contact resistance models described in the *Heat Transfer Handbook* (Chapter: *Thermal Spreading and Contact Resistances* [25]) were applied. The input assumptions used in these models are summarized in the following Table 3.1:

Assumption	Value
Titanium Thermal Conductivity (W/m·K) [26]	6.7
Carbon-Carbon Thermal Conductivity (W/m·K) [9]	12
Titanium Hardness (MPa) [27]	3432
Carbon-Carbon Hardness (MPa) (assumed)	490
Titanium Young's Modulus (Pa) [26]	1.2×10^{11}
Carbon-Carbon Young's Modulus (Pa) [28]	1.7×10^{10}
Titanium Poisson's Ratio [26]	0.35
Carbon-Carbon Poisson's Ratio [28]	0.12
Mean Profile Slope ($\tan \theta$) (assumed)	0.2
Std. Dev. of Profile Heights (m) (assumed)	5×10^{-6}
Titanium Thickness (m)	0.012
Contact Pressure (Pa) (assumed)	5×10^6

Table 3.1: Assumptions used in both the elastic and plastic asperity deformation contact model. Contact pressure magnitude assumption was based on interference calculations and material properties.

The resulting TCC was calculated to be 1771 W/m² K for the elastic asperity deformation model and 5221 W/m² K for the plastic deformation model. In comparison, the bulk thermal conductance of the titanium vessel wall in this application, given the 12 mm thickness of the vessel and a thermal conductivity of 6.7 W/m K, amounts to approximately 558 W/m² K. Since both contact models yield TCC values significantly higher than the through-thickness bulk conductance of the titanium, it was concluded that the thermal resistance at the interface is negligible. Although this provides a strong justification for adopting a perfect contact assumption, a more detailed investigation of this perfect contact assumption is presented in the Section 4.3.4.

3.1.4 Meshing

As shown in the overall workflow (Figure 3.1), the meshing strategy adopted in this study was based on generating the solid and fluid domains using two separate workflows. This separation allowed each domain to be meshed independently according to its specific requirements. Once completed, the two meshes were appended to form the final simulation case, which included both solid and fluid regions. This modular approach ensured greater flexibility during the design process, enabling independent refinement and/or easier updates to either domain.

3.1.4.1 Solid Domain Meshing

The solid domain includes the beam dump core and the surrounding titanium vessel. To generate the computational mesh, the solid geometry was first imported into *Ansys Mechanical*, where a tetrahedral mesh was created using the patch-independent meshing method. A maximum element size of 24 mm was applied to the core, while a finer resolution of 12 mm was used for the vessel. The resulting mesh, illustrated in Figure 3.16, consisted of approximately 7.5 million elements.




Figure 3.16: Initial tetrahedral mesh in *Ansys Mechanical*.

Since polyhedral meshes typically result in fewer elements, are more efficient to solve numerically and have also demonstrated improved numerical stability in the context of this application [24], the original tetrahedral mesh was converted into a polyhedral mesh within *Ansys Fluent* in solution mode using the built-in automatic conversion tool. Whilst the automatic conversion process simplifies the workflow, it comes with a limitation: mesh refinement settings are inherited from the original tetrahedral mesh, meaning that polyhedral cell sizing cannot be directly controlled either before or after conversion.

The conversion results in a polyhedral mesh with approximately 1.5 million cells, achieving a minimum orthogonal quality of 0.21 and average of 0.9. The quality and structure of the final mesh are shown in Figure 3.17.




Figure 3.17: Converted polyhedral mesh in *Ansys Fluent* Solution mode.

3.1.4.2 Fluid Domain Meshing

The fluid domain was meshed using the *Ansys Fluent* mesher with the watertight geometry workflow. During the geometry preparation phase, dedicated solid bodies were inserted to define mesh refinement regions using BoIs in areas where strong velocity gradients were expected. BoIs were therefore added around the nozzles (observable in Figure 3.18) with a target mesh size of 3.5 mm, and in the impingement jet zones between the nozzle and the dump surface (Figure 3.19), where a 5 mm mesh size was imposed.




Figure 3.18: BoIs used to refine the mesh around the nozzle area.




Figure 3.19: BoIs used at the impingement zone between nozzles and dump surface.

During the initial solving attempts, the warning "*turbulent viscosity limited to viscosity ratio*" appeared. This message indicates that, in certain regions of the domain, the turbulent viscosity μ_T calculated by the turbulence model became excessively high compared to the molecular viscosity μ , triggering *Fluent's* internal stability limiter.

In this study, the $k-\omega$ *SST* turbulence model was used. According to the *ANSYS Fluent* Theory Guide [29], the turbulent viscosity is computed as

$$\mu_T = \frac{\rho k}{\omega} \cdot \frac{1}{\max\left(\frac{1}{\alpha^*}, \frac{SF_2}{a_1 \omega}\right)} \quad (3.1)$$

where:

- ρ is the local fluid density,
- k the turbulent kinetic energy,
- ω the specific dissipation rate and
- $\frac{1}{\max\left(\frac{1}{\alpha^*}, \frac{SF_2}{a_1 \omega}\right)}$ is a term, based on model specific constants and functions, and the strain rate magnitude, that controls the behavior of the model between near-wall and free-stream flow regions.

Furthermore *Ansys Fluent* limits the *turbulent viscosity ratio*

$$\frac{\mu_T}{\mu} \quad (3.2)$$

to a maximum value of 10'000 to ensure numerical stability, potentially affecting solution accuracy in localized regions where the mesh is too coarse. In order to localize these problematic areas, turbulent viscosity iso-surfaces were visualized and high values were identified near the support structures, where the narrowing of the flow cross-section led to elevated velocities and steep velocity gradients. To improve local resolution and solver stability, two BoIs with a mesh size of 3.5 mm were therefore introduced around the supports (Figure 3.20).

Furthermore given the sensitivity of the solution to fluid domain geometry, this refinement workflow based on iso-surface visualization was repeated for multiple geometries and mesh variants throughout the study. While most cases are not discussed in detail, one particularly relevant example is reported here (its significance is further explained in Section 4): a BoI with 3.5 mm

mesh resolution was added in the region where the insertion of top ducts induced stronger velocity gradients and lead to the formation of a major recirculation zone (Figure 3.21).




Figure 3.20: BoIs highlighted in orange added around the support structures.




Figure 3.21: BoI introduced in the main recirculation region after top duct geometry insertion.

In addition to BoIs, local face sizing refinements were applied to critical regions:

- Dump surface: 20 mm
- Supports (guides): 50 mm
- Enclosure: 60 mm
- Ducts: 30 mm
- Inlets: 20 mm
- Outlet: 30 mm

and curvature-based refinement was also applied to the dump surface using the following parameters:

- Minimum element size: 10 mm
- Maximum element size: 40 mm
- Curvature normal angle: 1.7°
- Scope: edges only

Depending on the specific geometry and scenario simulated, the final polyhedral meshes contained between 9.7 and 12.5 million elements, all with a minimum orthogonal quality greater than 0.21, ensuring high quality fluid mesh across all cases.

3.1.5 Boundary Conditions

The boundary conditions used in this study are defined to represent the physical configuration of the HL beam dump.

Starting with the inlets (Figure 3.22), a mass flow inlet condition is applied for each of them with an inlet temperature of 21.7°C and a mass flow rate varying between 0.07 and 1.7 kg/s depending on the size/number of inlets and the boundary conditions configuration. These values

are scenario-specific and are further detailed in Section 4.2 & 4.3. At the outlet (Figure 3.23), a gauge pressure of 0 Pa is imposed with the backflow temperature set equal to the inlet temperature to maintain thermal consistency in case of minor recirculation.

Turbulence intensity and turbulent viscosity ratio were set to 5% and 10, respectively. The turbulence intensity and turbulent viscosity ratio used in this study were adopted from previous Run 3 models, as the upstream duct geometry and consequent flow history, which strongly influence inlet turbulence, were not represented in the current model and were not fully defined at the time of simulation. These values align with common assumptions for internal duct flows and offer a reasonable approximation in the absence of detailed upstream data. Nonetheless, because turbulence can significantly affect flow development and heat transfer near the walls, it may be useful to assess the sensitivity of key results to these assumptions in future studies.




Figure 3.22: Inlets' locations within the computational domain.




Figure 3.23: Outlet's location within the computational domain.




Figure 3.24: Enclosure representing the interior of the shielding and enclosing the entirety of the computational domain.

Furthermore, all walls within the domain are treated as no-slip boundaries. For the external enclosure (Figure 3.24), a zero conductive and convective heat flux condition is applied. Radiative heat exchange is considered from the vessel to the enclosure walls using a surface-to-surface (S2S) radiation model. This setup effectively treats the enclosure as thermally insulated to conduction and convection, while enabling radiative heat loss from the vessel surface. The shielding outside the enclosure is not included in the model and is assumed to have a very high thermal mass, such that its temperature remains constant. As the flow interacts with the shielding only after

impingement cooling, any convective heat exchange with it is considered negligible and is not included in the analysis. Finally, radiation within the dump block has been neglected due to its geometric complexity and relatively small contribution compared to conduction. In the S2S model, air is treated as a non-participating medium: it does not emit, absorb, or scatter radiation, and its absorptivity is set to zero. As the net radiative heat exchange between two surfaces i and j is described by

$$q_{i \rightarrow j} = \frac{\sigma \cdot (T_i^4 - T_j^4)}{\left(\frac{1-\varepsilon_i}{\varepsilon_i A_i} + \frac{1}{A_i F_{i \rightarrow j}} + \frac{1-\varepsilon_j}{\varepsilon_j A_j} \right)} \quad (3.3)$$

where:

- σ is the Stefan-Boltzmann constant,
- T_i, T_j the surface temperatures,
- A_i, A_j the surface areas,
- $F_{i \rightarrow j}$ the view factor between surfaces i and j and,
- $\varepsilon_i, \varepsilon_j$ the emissivities of the surfaces,

an emissivity of 0.9 is assigned to the shielding, with its rough concrete surface, while the titanium vessel, being a smoother and machined metallic surface, has a lower emissivity of 0.4 (see the Engineering Toolbox [30] and Figure 3.25).




Figure 3.25: Visualization of the emissivity values assigned to each component in the thermal model.

Finally it is worth noting that the cradle and cooling ducts have been excluded from the radiation model. Generating the necessary view factor files for these components using the S2S method is computationally expensive and, since the thermal phenomena occurring in these regions are not of interest for the current analysis, excluding them significantly simplifies the setup and optimizes computational resources.

3.1.5.1 Energy Deposition

To model the thermal load within the HL beam dump, energy deposition maps were generated using a Monte Carlo based particle transport code called FLUKA [12]. This tool simulates how protons interact with the dump geometry and materials accounting for the beam sweep pattern described in Section 2.2.1.1. While previous studies employed FLUKA maps on earlier geometries with reduced values of energy density, this work considers the new HL-LHC dump block and HL beam parameters.

Since FLUKA calculates energy deposition per primary particle, the raw output must be converted and scaled to reflect realistic beam conditions. This is achieved using an adapted Python script that was developed within the SY-STI-TCO section: "FLUKA to Ansys tables" [31]. This script both converts the FLUKA output into an *Ansys Fluent* compatible format, scales it according to the beam intensity and distributes it according to given pulse duration. For the HL scenario, the following values based on nominal beam parameters stated in the TDE functional specifications [9] were used:

- **Beam intensity:** 6.35×10^{14} protons per pulse
- **Pulse duration:** 86 μ s

The resulting heat source was interpolated onto the fluid mesh using cell-center interpolation, in line with the discretization approach described in Section 3.2. This method allowed for consistent integration of the FLUKA energy deposition maps into the CFD model.

Furthermore, a comparison between the total energy deposited in the FLUKA solid domain and the energy effectively transferred to the mesh is shown in Table 3.5, with a cumulative error of approximately 5% observed in the current setup. For reference the solid section segmentation and naming convention used are observable in Figure 3.26.

Solid Section	FLUKA [MJ]	Fluent [MJ]	Error
Vessel	17.7	16.0	9.6%
HD US 1	6.2	4.7	24.2%
HD US 2	22.9	21.0	8.3%
LD 1	45.6	42.9	5.9%
LD 2	66.7	63.2	5.2%
LD 3	77.6	74.2	4.4%
LD 4	76.7	73.5	4.2%
HD 1	97.8	94.7	3.2%
HD 2	68.1	65.8	3.4%
HD 3	40.1	38.9	3.0%
HD 4	22.1	20.6	6.8%
HD 5	11.3	10.0	11.5%
Dump Block	552.8	525.5	4.9%

Table 3.2: Comparison of FLUKA energy per pulse and resulting FLUENT internal energy across dump block solid sections.



Figure 3.26: Naming and segmentation of the solid dump block sections.

The relatively coarse mesh used for the solid domain, combined with the limited number of FLUKA bins (FLUKA equivalent of *Ansys Fluent* mesh cells) and the interpolation method used to map the energy deposition field to Fluent, leads to a small underestimation of energy deposited in the dump block. However, this underprediction remains within acceptable limits and does not compromise the objectives of the present study. Unlike in thermo-mechanical damage assessments, where local temperature spikes could drive material failure, the focus here is on the overall thermal behavior of the system, particularly during the heating and cooldown phases. Moreover, this slight underestimation is implicitly covered by the safety margins built into the design, including conservative thresholds for graphitic materials and allowances on vessel yield strength and fatigue resistance. As a possible mitigation, the energy deposition field obtained from FLUKA could be scaled up by a small factor with "FLUKA to Ansys tables" [31] before mapping it to the Fluent model, to compensate for the underestimation introduced by mesh and binning limitations. However, this correction was not implemented in the current study, as it would not provide a rigorous solution: the overall underestimation does not apply uniformly across all regions, and applying a global scaling factor would neglect local variations. In practice, each segment would require its own tailored correction, making the approach both imprecise and impractically complex for the present scope.

Finally the chosen solid mesh, consisting of approximately 1.5 million elements, was selected as a compromise between accuracy and computational costs (Section 3.1.4). A dedicated mesh sensitivity study, presented in Section 3.3.3, confirms the robustness and the convenience of this choice for the intended scope of analysis.

3.1.5.2 Material Models

This subsection reports the material models used exclusively in the new HL beam dump geometry. For other simulations and corresponding material data, such as those used for Run 3 studies (Section 4.1), the reader is referred to the corresponding report [24].




Figure 3.27: Temperature-dependent evolution of selected thermal properties for the materials used in this study.

Material	Density [kg/m ³]
Air	Ideal Gas
Argon (modeled as Solid)	1.78
CFC	1200
HD Graphite	1730
Titanium	4430

Table 3.3: Assumed constant density of the considered materials.

Figure 3.27 and Table 3.3 present the thermal properties implemented in the current model. For graphitic materials, the values were taken from [32], which reports distinct property sets for different material grades, each with its own characteristic density. In this work, a fixed density value was selected to represent the graphite used. It should be noted that different material assumptions could influence both FLUKA energy deposition results and the thermal balance

calculations.

The plots in Figure 3.27 clearly highlight the temperature dependence of both specific heat capacity, that was calculated based on measured axial thermal diffusivity and assumed POCO graphite properties, and thermal conductivity for the considered materials. Notably, CFC is modeled as anisotropic, with distinct behaviors in conductivity in the xy and z directions. The out-of-plane (z) conductivity was fitted from experimental measurements reported in an internal document [32]. However, the in-plane (xy-direction) conductivity was not available in the referenced documentation and was instead estimated based on engineering experience and historical data provided by both expert team members and external suppliers [33]. A dedicated sensitivity study on the impact of CFC material properties is nonetheless presented in Section 3.3.2 and will be further explored in that context.

As it can be observed in the aforementioned graphs, isostatic graphite and titanium also exhibit strong temperature dependence. Unlike the CFC data, the thermal properties of isostatic graphite have already been benchmarked in prior simulation studies [24] and are considered validated for current use. Furthermore Ti-6Al-4V properties are extracted from a supplier datasheet and air is treated as an ideal gas with density that is only dependent on the operating pressure of 101'325 Pa.

Finally, it is important to note that the argon gaps are modeled as solids with constant properties of argon in standard conditions. Although this represents a considerable simplification, especially in the context of the thermal contact resistance analysis presented in Chapter 4, it is consistent with the approach adopted in previous studies where gaseous nitrogen gaps were also modeled as solids and the results were subsequently validated experimentally. This assumption is maintained here to simplify both the setup and solving of the model, as well as the meshing of the argon-filled gaps between graphite components. A dedicated sub-model could be developed to include convection in the argon gaps. While this would likely improve heat transfer between graphite components and slightly reduce local temperature gradients and peak temperatures, the impact is expected to be limited as axial heat transfer has been deemed less significant. However, accurately modeling this effect would require a detailed fluid simulation, adding substantial complexity and computational cost. Due to time constraints, this is not pursued in the present work.

3.2 Simulation Setup

As part of the simulation setup, key quantities are defined using report definitions to monitor thermal and fluid behavior over time throughout the computational domain. As discussed in Section 3.1.2.2, the solid geometry, particularly the vessel, is divided into different sections as this allows for efficient evaluation of internal energies, maximum and average temperatures in both the core and vessel, as well as the heat transfer coefficient on the vessel surface. These quantities are recorded at each timestep and written to a text file, allowing for targeted post processing. In parallel, case and data files are periodically written to allow the creation of checkpoints and to support additional visualization aids, such as generating different contour plots.

The simulation workflow itself, shown in Figure 3.1, begins with a steady-state flow solution, which provides a fully developed velocity field to initialize the transient simulation. Once this initial condition is established, the energy deposition maps are applied in the transient simulation over a $86\ \mu\text{s}$ interval discretized into 20 time steps of $4.3\ \mu\text{s}$ each. Following this intense energy deposition phase (the beam dump event), the timestep size is progressively increased up to 60 seconds, allowing the solver to remain stable while capturing both the thermal gradients generated immediately after the beam impact and the long-term thermal diffusion during the cooldown phase.

This level of temporal control and numerical precision is achieved thanks to the computational resources available at CERN as all simulations are executed on the CERN HPC facility, using SLURM and journal files to allocate resources. A transient run spanning 2.5 hours of flow time typically requires, depending on the mesh, between 6 and 9 hours of computation on 8 nodes, each equipped with 20 CPU cores.

Finally the simulations use the $k\text{-}\omega$ SST model which effectively captures both near-wall effects and free-stream behavior (Section 3.1.4.2). A pressure-based solver is employed with a coupled pressure-velocity scheme, all transport equations are discretized using the second-order upwind scheme and gradient calculation is carried out using the least squares cell-based method.

3.3 Model Checks

In this section, a series of checks, sensitivity analyses on selected material properties and mesh refinement studies for both the fluid and solid domains are presented. These were conducted specifically for the current HL beam dump geometry (Figure 3.5) with the currently installed cooling system (Figure 2.10).

3.3.1 Basic Checks

Figure 3.28 includes a set of diagnostic quantities that give useful insight into whether the simulated behavior remains consistent with expected physics across all the simulated timeframe.




Figure 3.28: Time evolution of key diagnostic quantities, including mass flow rate error, temperature difference between outlet and inlet, and total pressure drop across the domain.

The top plot shows the percentage difference in mass flow rate between the inlet and outlet throughout the full simulation time. Apart from a few sharp spikes (later further explained), the flow imbalance remains consistently around 0%, confirming effective mass conservation.

The middle plot instead illustrates the temperature difference between the outlet and inlet. The curve reflects the presence of three distinct beam dump events as expected, with thermal transient behaviors consistent with the thermal trends observed in the solid domain (see Section 4).

Finally, the last plot displays the total pressure at the inlet minus that at the outlet. Overall, the values remain stable with minor oscillations, except for the already mentioned distinct spikes that occur at the same time as the beam dump events. These spikes are attributed to the sudden, localized heat release within a very short timescale, which can induce brief pressure transients. However, this does not affect the credibility of the simulation, as the model is primarily intended to capture the system's behavior over longer cooling timescales. Given that the beam energy deposition occurs over just 86 μs , whereas the thermal diffusion and convective cooling evolve

over several hours, the flow accuracy during such a short pulse is of marginal importance to the objectives of this work.

3.3.2 Material Properties Sensitivity Analyses

Given the uncertainties in the CFC material properties discussed in Section 3.1.5.2, two sensitivity tests were performed: one on the in-plane conductivity and one on the out-of-plane conductivity.

$$k_z = \rho \cdot c_p \cdot \alpha_z \quad \left[\frac{\text{kg}}{\text{m}^3} \cdot \frac{\text{J}}{\text{kg} \cdot \text{K}} \cdot \frac{\text{m}^2}{\text{s}} = \frac{\text{W}}{\text{m} \cdot \text{K}} \right] \quad (3.4)$$

For the sensitivity analysis on the axial thermal conductivity of CFC, thermal diffusivity α_z and POCO graphite assumed c_p were taken from reference [32]. Based on the available data, multiple conductivity curves were plotted using Equation 3.4, extrapolated based on engineering assumptions and are observable in the following figures.




Figure 3.29: Reported values of k_z .




Figure 3.30: Extrapolated values of k_z .

As can be observed in Figure 3.31 and 3.32, the differences between the tested cases appears to have little to no influence on the global thermal response under the considered conditions.




Figure 3.31: Maximum temperatures for core and vessel for the 3 simulated k_z .




Figure 3.32: Residual total internal energy within the dump block for the 3 simulated k_z .

Regarding the sensitivity analysis for the in-plane thermal conductivity of CFC, the tested curves were generated by applying the constant ratio between k_{xy} and k_z of 4.5 provided by a material supplier, and multiplying it with the previously extrapolated k_z curves shown in Figure 3.30 resulting in the curves observable in Figure 3.33.




Figure 3.33: Considered derived temperature dependent k_{xy} curves.

The effects of varying k_{xy} are clearly observable in the thermal response of both the core and the vessel, with increased k_{xy} leading to lower maximum temperatures in the core as the heat is more efficiently transferred to the vessel (Figure 3.34). As a result, the vessel reached higher peak temperatures especially during the diffusion-dominated phase following the beam dump and, as the thermal gradient with the surrounding cooling air augments, faster heat extraction was promoted.

This improved thermal coupling is also observable in both the internal energy balance (Figure 3.35) and the maximum vessel temperature curves: the cases with higher k_{xy} showed a lower remaining internal energy (or a greater amount of energy extracted from the dump block over

time) and a faster decline of vessel temperatures during the cooldown phase in the high conductivity cases.




Figure 3.34: Maximum temperatures for core and vessel for different CFC in plane conductivity values.




Figure 3.35: Residual total internal energy within the dump block.

Despite these trends, the differences remained relatively limited in the context of the overall modeling uncertainties. However, to ensure a conservative approach, the lowest in-plane conductivity curve was selected for the main simulations. This conservative choice results in slightly reduced radial heat extraction, leading to higher core temperatures, greater residual internal energy in the system and slightly higher vessel temperatures toward the end of the dump. Although the peak vessel temperature is marginally lower, the overall heat removal is less efficient, providing a lower-bound scenario for thermal performance evaluation.

3.3.3 Mesh Sensitivity Analyses

Given the turbulence model used (Section 3.2) and its sensitivity to near-wall resolution, a mesh sensitivity analysis was carried out for the fluid domain. The baseline configuration, introduced in Section 3.1.4.2, uses a four-layer inflation layer mesh grown at the dump surface and results, for the biggest part of the interest part of the dump, in y^+ values between 5 and 30, as shown in Figure 3.36. To better understand the influence of wall resolution on the results, two additional meshes were generated.




Figure 3.36: y^+ values at the wall of the 4 inflation layers fluid mesh.

The first mesh was refined to bring y^+ as close as possible to 1 (Figure 3.37) and to get 15 layers in the viscosity-affected region, aiming to fully resolve the boundary layer. The second mesh was instead coarsened to reach y^+ values greater than 30 in most of the region of interest (Figure 3.38) to utilize a wall-function for the boundary layer solution. The parameters and mesh characteristics of all three configurations are summarized in Table 3.4.




Figure 3.37: y^+ values at the first cell near the wall of the extremely refined boundary layer.




Figure 3.38: y^+ values at the first cell near the wall of the coarse boundary layer.

Fluid domain	# Cells	# Inflation layers	Thickness first BL [mm]	y^+ values
Reference Mesh	9.8	4	0.7-1.1	5 -30
Low y^+	12.1	15	0.06 - 0.1	0 - 4
High y^+	9.4	2	4.2 - 9.3	30+

Table 3.4: Parameters of the boundary layer (BL) mesh used in the mesh sensitivity analysis. All simulations include a boundary layer mesh with a fixed growth rate of 1.2.

As shown in Figures 3.39, the evaluation was conducted across multiple consecutive dump events and the maximum temperature, as well as remaining internal energy in different components, were the main evaluation metric (particularly considering that these value will, in future work, be cross-validated against experimental data from installed PT100 sensors). As no unexpected behavior were observed, the boundary layer calculations only slightly affected the predicted maximum vessel temperatures. Additionally, the internal energy evolution within different sections of the dump block shows minimal deviation across the three meshes. Figures 3.40 and 3.41 confirm that the difference in remaining internal energy in each case remains within a 5% range (Figure 3.42), which is considered acceptable given the broader modeling uncertainties discussed in Section 4.3.2.




Figure 3.39: Maximum temperatures for core and vessel for the different boundary layer configurations.




Figure 3.40: Residual total internal energy within the dump block.




Figure 3.41: Residual internal energy in the highest energy core segment of the dump block and in the vessel over three consecutive dump events. Solid lines represent results using an extremely refined near-wall fluid mesh while dotted lines show results with only four boundary layers.




Figure 3.42: Internal energy percentage difference between $y+=1$ mesh and the 4 inflation layers mesh.

To conclude, while thermal behavior and local gradients are important drivers, the primary motivation for the subsequent solid mesh sensitivity analysis is the observed discrepancy between the calculated total energy deposited by FLUKA and the internal energy interpolated by Fluent after a dump event. Across simulations, this deviation consistently falls around 5% (Section 3.1.5.1).

To rule out misalignment, the FLUKA energy deposition maps were carefully reviewed and verified against the modeled geometry and no inconsistencies were identified. As prior studies showed that interpolation accuracy improves significantly with solid mesh resolution, two refined solid meshes, one with approximately 3 million elements and one with 5 million, were tested, each coupled to the same baseline fluid mesh. The results are summarized in Table 3.5.

Solid Section	FLUKA [MJ]	Number of Cells in the Solid Domain					
		1.5 Million		3 Million		5 Million	
		Fluent [MJ]	Error	Fluent [MJ]	Error	Fluent [MJ]	Error
Vessel	17.7	16.0	9.6%	16.1	9%	16.1	9%
HD US 1	6.2	4.7	24.2%	4.9	21%	3.8	38.7%
HD US 2	22.9	21.0	8.3%	21.4	6.6%	20.6	10%
LD 1	45.6	42.9	5.9%	43.4	4.8%	42.2	7.5%
LD 2	66.7	63.2	5.2%	63.9	4.2%	62.5	6.3%
LD 3	77.6	74.2	4.4%	74.8	3.6%	73.4	5.4%
LD 4	76.7	73.5	4.2%	73.9	3.7%	72.6	5.3%
HD 1	97.8	94.7	3.2%	95.1	2.8%	94.1	3.8%
HD 2	68.1	65.8	3.4%	66	3.1%	64.9	4.7%
HD 3	40.1	38.9	3.0%	39	2.7%	37.9	5.5%
HD 4	22.1	20.6	6.8%	20.6	6.8%	19.6	11.3%
HD 5	11.3	10.0	11.5%	10	11.5%	8.9	21.2%
Dump Block	552.8	525.5	4.9%	529.1	4.3%	516.6	6.5%

Table 3.5: Comparison of FLUKA energy per pulse and resulting FLUENT internal energy across dump block solid sections for different solid domain meshes. For solid section naming: Figure 3.26.

The findings confirm that sections with low energy deposited show proportionally the highest underestimation. However, the hotspot region is already well represented as the total difference between the energy deposited by FLUKA and the one interpolated in Fluent remains within the 3 to 4% range, even for the 1.5 million element mesh and does not significantly improve with further refinement. Finally considering the increased computational cost, the 1.5 million cell solid mesh was maintained across all main simulations as a reasonable compromise between accuracy and efficiency.

Chapter 4

Results & Discussion

This chapter presents the results obtained throughout the different phases of the study. It identifies the key parameters influencing thermal behavior and assess the robustness of the different potential HL cooling system upgrades under worst-case operating conditions.

First the analysis of the limitations of the existing Run 3 cooling configuration with Run 3 beam parameters is presented. This is followed by a preliminary investigation using the geometry of the Run 3 spare dump, adapted into multiple configurations. These configurations and their evolution reflect the iterative design process used to evaluate potential cooling system upgrades while the final HL dump geometry was still under development. Finally the core of the chapter focuses on the final HL dump block possible cooling configuration, where a range of simulation scenarios is explored in detail, including thermal performance assessments, residual internal energy trends, sensitivity analyses and contact pressure evaluations.

4.1 Run 3 Limits

The current Run 3 operational beam dumps consist of a SS 318LN vessel, with a core made from loose fitted Sigraflex[®] sheets retained by extruded graphite plates (see Section 2.2). To preserve structural integrity, a conservative temperature limit of 250°C was imposed on the vessel. This threshold was established by comparing vessel integrity simulations from Banks et al. [18] with the temperature-dependent yield strength curve of SS 318LN as, at this temperature, the mechanical stress caused by beam-induced vibrations dangerously approaches the material's yield strength, raising concerns over the reliability of the beam dump vessel's structure.

The main objective of this study was to understand the performance limits of the Run 3 beam dump system. Specifically, to identify the maximum number of acceptable dump events under different operating conditions, determine the intensity threshold at fixed frequency and establish the frequency threshold at fixed intensity. The key considered scenarios are summarized in Table 4.1.

Scenario	Bunch Intensity (#p)	Frequency (h)
# Dumps at max intensity	1.8×10^{11}	2
# Dumps at reduced intensity	1.6×10^{11}	2
Intensity threshold at constant frequency	1.2×10^{11}	2
Frequency threshold at constant intensity	1.6×10^{11}	3
Partial sweep failure case	1.6×10^{11}	2

Table 4.1: Naming, bunch intensities and frequencies of the different considered scenarios.

All thermal assessments presented in this section were based on the monitored regions LD top and HD 1 top (refer to Figure 4.1). These two locations were selected because previous energy deposition studies have consistently identified them as the most critical for hotspots under Run 3 operating conditions.




Figure 4.1: LD top and HD 1 top locations used for thermal analysis, shown at their physical locations on the vessel.

Simulation results confirm the limited thermal capacity of the current beam dump. At the maximum intensity of 1.8×10^{11} protons, the vessel reaches 250 °C after just two dumps (Figure 4.2), indicating that short intervals between high intensity dumps are not feasible. Reducing the intensity to 1.6×10^{11} protons slightly improves the situation, with the critical temperature likely reached after three dumps (Figure 4.3).




Figure 4.2: Vessel/core temperature evolution for 2-hour spaced dumps at full bunch intensity (1.8×10^{11} protons).




Figure 4.3: Vessel/core temperature evolution for 2-hour spaced dumps at reduced beam intensity (1.6×10^{11} protons).

For both the intensity threshold (Figure 4.4) and the frequency threshold (Figure 4.5) cases, the vessel temperature stabilizes after several dumps, but the temperature in the Sigraflex[®] core continues to rise gradually, suggesting progressive accumulation of heat in the low-density region even at moderate operational loads.




Figure 4.4: Vessel/core temperature evolution under constant 2-hour repetition rate with 1.2×10^{11} protons bunch intensity.




Figure 4.5: Vessel/core temperature evolution under fixed bunch intensity (1.6×10^{11} protons) with 3h repetition rate.

Finally, a scenario was considered, in which only the left half of the beam sweep is applied (Figure 4.6). This corresponds to only the first half of the beam being filled with protons and the second half left empty. In this case, the energy deposition is asymmetrically localized in the left part of the dump, but the overall vessel temperature remains lower than in the full sweep scenario. Visualization of the hotspot region (Figure 4.7) shows that the loosely fitted Sigraflex[®] in the core plays an important role in limiting peak vessel temperature formation under these asymmetric load conditions.




Figure 4.6: Thermal response under partial beam sweep (left-side only).




Figure 4.7: Hotspot analysis and temperature field within the end of the low density section and the start of the HD 1 block in the partial sweep scenario.

These results clearly show that the current Run 3 beam dump system operates near its thermal limits even under nominal conditions. As such, both a redesigned dump block and an upgraded cooling system are essential to manage the increased beam stored energy of HL-LHC.

4.2 Preliminary HL Investigation on Spare geometry

At the early stages of the HL dump development, the exact geometry was not yet defined. To support preliminary investigations, the Run 3 spare model geometry was therefore adopted. This setup enabled early assessments of potential cooling system improvements, focusing on airflow distribution and the system's thermal response under the HL "*emergency scenario*" [9]: three consecutive 711 MJ dump events separated by 2.5 hour cooling intervals. These investigations built directly on insights from a prior work (see Appendix B.1), which had shown that optimizing only the bottom ducts would likely lead to limited thermal improvements. At the same time, radiological safety discouraged major interventions in activated areas, such as replacing the bottom ducts entirely.

In parallel, the HL dump design evolved, with updated material choices: a low-density CFC core and a Ti Gr5 vessel. Accordingly, all simulations in this phase were carried out using these updated materials. FLUKA energy deposition maps were retained for the core, given the similar density between CFC and Sigraflex[®], while new maps were generated for the titanium vessel to reflect its lower energy absorption. Two primary airflow categories were considered: standard ($1x = 5'000 \text{ m}^3/\text{h}$) and increased ($2x = 10'000 \text{ m}^3/\text{h}$), each with multiple variants involving top and bottom duct modifications or their combinations. As the CFC blocks were shrink-fitted within the vessel, a perfect thermal contact assumption was applied in all simulations, replacing previous interface dummy elements with ideal contact conditions, except for the case discussed in Section 4.2.3.

Finally, a summary of all configurations and corresponding results is provided in Figure 4.8, with each scenario discussed in detail in the following subsections.




Figure 4.8: Thermal and energy extraction comparison of nine beam dump cooling configurations/upgrades. Each configuration is evaluated based on the peak and average temperatures in core and vessel, shown as horizontal bars, as well as black diamond markers that indicate the % of beam's stored energy extracted per dump averaged over three beam dump events.

4.2.1 Baseline: 1x Flowrate Shrink Fitted CFC LD Section

The baseline configuration consists of a fully shrink fitted CFC monolithically modeled LD core, a Ti Gr5 vessel and the standard nominal airflow rate of 5000 m³/h. This setup reflects the first model combining HL-relevant materials and outdated Run 3 spare geometry and is used as a reference for evaluating all subsequent upgrade scenarios with 1x airflow configuration.

Previous investigations [24] have already shown that switching the vessel material from SS 318LN to Ti Gr5 has only a negligible effect on maximum vessel temperatures. However, a more significant change arises from the introduction of perfect thermal contact in the LD section. Compared to the simulations with Sigraflex[®] (for example the one discussed in Section 4.1 or B.1), where thermal contact resistances were present, the perfect contact condition in this case leads to a marked increase in heat transfer from the core to the vessel. As a result, vessel peak temperatures rise significantly, reaching up to 345°C. The location of the thermal hotspot remains consistent with earlier findings, appearing at the transition between the LD block and HD 1 (See Figure 4.1 for reference on the section naming).

The total energy extracted per dump averaged over 3 dump events in this case is approximately 53.6%, which serves as a thermal efficiency benchmark for all subsequent configurations. However, this scenario also highlights the limitations of the current nominal airflow rate, as both peak and average temperatures remain elevated throughout the dump sequence.

4.2.2 Baseline: 2x Flowrate Shrink Fitted CFC LD Section

To directly evaluate the impact of increasing the airflow, the same configuration as the baseline case was simulated with a doubled flowrate of 10'000 m³/h. This case serves as baseline for all 2x flowrate simulations. Achieving this flowrate requires a new, dedicated AHU and, since the system is yet to be defined, it can be tailored to the target flow and pressure drop, offering wider design margins compared to scenarios constrained by existing equipment.

The increased convective heat transfer on the bottom of the dump leads to a clear improvement in thermal performance compared to the 1x flowrate case. The average vessel temperature is reduced by more than 35°C, while the average core temperature also decreases noticeably. The maximum vessel temperature drops by nearly 44°C. Finally the total energy extracted per dump rises to approximately 68%, confirming the benefits of enhanced airflow even without additional changes to the cooling geometry.

This configuration demonstrates that doubling the flowrate alone provides a substantial improvement in both average temperatures and energy extraction under the scenario of three consecutive high intensity HL dumps.

4.2.3 2x Flowrate Non-Shrink Fitted CFC LD Section

To evaluate the impact of thermal contact resistance in the LD region of the dump, a worst case scenario was simulated using the same 2x airflow configuration but without shrink fitting the CFC monolithic block. Dummy interface elements were retained between the CFC core and the titanium vessel with a TCC value of 8 W/m²·K. This corresponds to a loosely fitted core, “floating” inside the vessel, with an effective thermal contact equivalent to having an argon gap of approximately 2 mm to the vessel, assuming the thermal conductivity of argon at room temperature. As a result, the average core temperature rises by almost 50% and similarly, the total energy extracted per dump drops from approximately 68% to 40%. The peak vessel temperature decreases from 301°C to 282°C, which can be attributed to the reduced heat flux from the core to the vessel due to the presence of thermal resistance. Finally, the average vessel temperature remains nearly unchanged indicating that, while peak heating at localized regions is reduced, the system as a whole retains more thermal energy.

This case clearly demonstrates the effect of the strong thermal contact and efficient heat extrac-

tion given by shrink fitting. Even with doubled airflow, the thermal contact resistance in this case strongly limits heat extraction. A more detailed investigation of thermal contact behavior and its modeling in the simulations is presented in Section 4.3.4.

4.2.4 Top Ducts Addition

Until this point, all investigations focused on boundary condition variations, specifically inlet mass flow, and assumptions on thermal contact resistance in the low-density section. The next step was to evaluate geometric upgrades, starting with the addition of top cooling.

The new configuration includes a set of ducts placed above the dump block (Figure 4.9), with the same cross-sectional area as the bottom duct. The bottom duct geometry was not altered. This upgrade was motivated by two observations:

- the hotspots consistently formed on the upper surface of the dump and
- bottom-side-only impingement cooling showed limited effectiveness (Appendix B.1).

While this is an idealized layout, since the available space within the shielding is limited, it allows for testing the full thermal potential of direct top-side impingement. Finally, based on nozzle outlet area, 27.5% of the total inlet mass flow was redirected to the top ducts, with the remaining 72.5% kept through the bottom ducts.




Figure 4.9: Duct schematic showing the presence of top ducts in the model.

Compared to the 1x flowrate baseline, the addition of top ducts shows clear benefits. The maximum vessel temperature drops from 345°C to 279°C, and the energy extracted per dump increases by nearly 4%. The average vessel temperature decreases from 281°C to 244°C, a significant reduction that can be attributed to the direct targeting of the hotspot area. The upgrade increases the surface area directly cooled by the nozzles, enhancing local heat removal at hotspot regions and redistributing the temperatures more evenly across the dump surface. This effect is clearly visible in Figure 4.10, where both core and vessel temperature are consistently reduced across all three dump events. Additionally, the cooling rate between dumps steepens, even at the nominal 5'000 m³/h airflow condition.




Figure 4.10: Time evolution of maximum vessel (and core) temperatures across three consecutive dumps for the 1x flowrate configuration, comparing baseline and top duct cases.

The second configuration combines the addition of top ducts with a 2x flowrate, coupling the advantages of increased convective capacity with improved airflow placement. The flow split is again 27.5% to the top ducts and 72.5% to the bottom.

Compared to the 2x flowrate baseline, the maximum vessel temperature drops from 301°C to 243°C, while the average vessel temperature falls from 243°C to 207°C. These results highlight the importance of placing cooling where it is most effective to redistribute heat more uniformly across the vessel. Figure 4.11 confirms this, showing consistently lower peak temperatures and faster cooldown throughout the three dump cycles in both the vessel and the core.




Figure 4.11: Time evolution of maximum vessel (and core) temperatures across three consecutive dumps for the 2x flowrate configuration, comparing baseline and top duct cases.

Finally, the addition of top ducts increases the total inlet area, reducing the overall pressure drop by nearly 40%. This helps maintain system performance and allows more flexibility for combining with other upgrades (see Subsection 4.2.6).

4.2.5 2x Flowrate Central Bottom Cooling

The goal of this configuration was to explore whether focusing cooling efforts locally, while closing the majority of upstream and downstream nozzles, could offer a simplified yet effective upgrade. The idea originated once again from the observation that the hotspots region consistently appear in the longitudinally central part of the dump between the support structures. This study was conducted under a 2x airflow scenario, consistent with the enhanced baseline configuration. The 1x flowrate case was not pursued, as the reduced number of open nozzles would likely cause a high pressure drop, exceeding the limits of the existing AHU. In contrast, since a new AHU would be required to support 2x flowrate operation, it could be sized to accommodate the pressure demands of the central cooling layout as part of the overall system design. The parametric nozzle density script (Appendix B, Listing B.1) was therefore reused to generate this geometry, where 100 nozzles are concentrated in the central section of the duct, with only a few sparse nozzles retained upstream and downstream. The final layout is shown in Figure 4.12.




Figure 4.12: Dump assembly showing the generated central optimized bottom cooling ducts.

While the total energy extracted per dump is lower than in the 2x flowrate baseline, the maximum vessel temperature still decreases and this benefit becomes more pronounced across the sequence of three dumps as reported in Appendix B, Table B.5. This suggests a shift in the internal heat transfer path, with cooling better targeted at high energy regions and less flow reaching areas with lower energy deposition. Despite extracting less total energy, this rebalancing keeps peak vessel temperatures lower, indicating improved local thermal managing rather than increased overall cooling. On the downside, concentrating the cooling flow through a reduced number of nozzles results in a significant increase in system pressure drop: nearly 25% compared to the 2x flowrate baseline.

4.2.6 Top Ducts and Optimized Central Bottom Cooling

For this combined upgrade, the idea was to test whether the combination of localized bottom-side cooling and targeted top-side impingement could improve thermal performance while mitigating the pressure rise associated with central bottom cooling alone. Two simulations were run: 1x and 2x flowrate. The geometry used for this configuration is shown in Figure 4.13 and, as with the earlier central cooling case, approximately 100 bottom nozzles were concentrated in the central region between the support structures, while 2 top ducts delivered 27.5% of the total airflow on the assembly's theoretical hotspot.




Figure 4.13: Dump assembly showing the generated central optimized bottom cooling ducts coupled with the addition of top ducts upgrade.

The thermal metrics are comparable to those achieved with top ducts alone, both in the 1x and 2x flow scenarios. In fact, vessel and core temperatures are slightly lower, while the total energy extracted is marginally reduced. The main advantage of this combined configuration lies in its pressure behavior: by reducing the number of active bottom nozzles and introducing a second inlet path via the top ducts, the total pressure drop increased by only 1% compared to the baseline, significantly lower than the +25% observed in the central bottom-only case. This favorable pressure profile could enable reuse of the existing AHU at 1x flowrate and simplify AHU dimensioning in the 2x flowrate scenario, without compromising thermal performance. Despite these positive results, the configuration was ultimately excluded from the final upgrade strategy due to implementation complexity. Specifically, any solution requiring modification of the bottom duct layout was deemed impractical, as it would involve extended work in radiologically activated areas.

4.2.7 Practical Ducts

This configuration is the modified version of the top duct concept, adapted to fit within the available space inside the dump's shielding. This upgrade was considered given:

- the proven thermal effectiveness in targeting hotspot regions of the top ducts (see Section 4.2.4),
- the need to respect radiological constraints and
- the potential to reduce overall pressure drop, making it possible to redistribute flow through the top ducts without requiring changes to the existing AHU and optimizing both pressure

losses and thermal performances (see Section 4.3.5).

The following constraints were therefore defined in collaboration with the design team.

- Given the uncertainties regarding the shielding positioning, the nozzles' extremities had to remain at least 4 cm away from the dump surface, to avoid mechanical interference.
- The ducts had to extend upstream up to the shielding's opening to allow direct connection to external ducting and the AHU.
- Standard 104 mm diameter tubes were to be used to simplify manufacturing.
- The duct layout had to avoid interference with lifting points and support structures, particularly near the top of the cradle assembly.

A nozzle layout was iteratively developed to balance manufacturing feasibility, cooling performance requirements and the spatial distribution needed to extend the high HTC zones across the dump top surface in critical locations. The final design, shown in Figure 4.14, features a top-to-bottom inlet area ratio of approximately 1:6.25, with top ducts measuring 104 mm in diameter and equipped with 50 pairs of nozzles per duct, each 14 mm in diameter.




Figure 4.14: Duct schematic showing the presence of the developed practical top ducts in the model.

To better understand the assigned flow distribution, simplified steady-state simulations (Appendix A.1) were carried out. These simulations revealed a massflow distribution ratio of approximately 1:7.5 between the top and bottom inlets, corresponding to 5'000 m³/h total flowrate. When comparing this practical design to the earlier 1x flowrate idealized top duct configuration, results were surprisingly similar. Despite the different geometry and the multiple spatial constraints imposed by both the shielding and the support structures, the thermal performance of the practical implementation remained very close to that of the ideal case. Maximum and average vessel temperatures were slightly higher, and the energy extracted was marginally lower, but the overall differences remained within an acceptable margin (Figure 4.8). These results confirm that a realistic, shielding-conforming top duct layout achieves nearly the same thermal performance as the ideal case, with only minor losses that are offset by its greater feasibility under the radiological and mechanical constraints of the system. Furthermore, the solution offers advantages in terms of system pressure behavior, flow control flexibility (Section 4.3.5) and potential integration with other upgrades (for example the 2x flowrate upgrade, see Section 4.3.1).

These results marked the conclusion of the preliminary upgrade exploration phase, which established the thermal performance potential and practical constraints of various configurations using the Run 3-based placeholder geometry.

4.3 HL Dump Block Geometry

This section presents the results obtained using the finalized HL dump block geometry described in Section 3.1.2. The model includes the complete integration of all updated design features, including the Ti Gr5 vessel, shrink-fitted CFC modeled as monolithic blocks and isostatic graphite blocks. Regarding energy deposition the simulations reflect realistic HL-LHC full intensity beam conditions and build upon the findings of the earlier upgrade studies, now applied to the real HL dump design.

The analysis is structured around key performance metrics and sensitivities. It begins with a review of the considered scenarios and the main sources of uncertainty in both the thermal modeling and boundary conditions. Subsequent subsections present results on internal energy and heat flux distribution (Section 4.3.3), 3D contact pressure analysis (Section 4.3.4) and targeted sensitivity studies to massflow distribution between top and bottom ducts (Section 4.3.5).

Finally, a set of simulations is dedicated to exploring worst-case scenarios for TCC. First, the perfect contact assumption is removed and replaced with physically motivated TCC estimates. Then, by combining contact pressure results from mechanical simulations with temperature outputs from CFD, a temperature-dependent contact resistance model is introduced, providing a first step toward coupling the two domains in a more realistic framework.

4.3.1 Considered Scenarios

Following the conclusions of the preliminary investigations (Section 4.2), two upgrade strategies were identified as the most promising: doubling the airflow and implementing the practical top duct layout. This subsection assesses their performance within the finalized HL dump geometry, using updated FLUKA maps and a fully reworked model.

Although the core upgrade concerned the solid domain, the transition to the finalized geometry required significant changes to the fluid domain due to updated dimensions, material layout, and interface boundaries (see Section 3.1.4).

Two geometric configurations were modeled, each with a distinct external cooling setup:

- *Operational*: current duct layout (Figure 4.15)
- *Top Ducts*: newly developed top-side cooling (Figure 4.16)




Figure 4.15: HL dump geometry with current operational duct configuration.




Figure 4.16: HL dump geometry with top duct cooling configuration.

Each geometry was simulated under two flowrate conditions, resulting in four primary scenarios:

- Operational 1x Flowrate; current duct layout with nominal Run 3 flowrate (5'000 m³/h); serves as baseline.
- Operational 2x Flowrate; same duct layout with upgraded AHU delivering 10'000 m³/h.
- Top Ducts 1x Flowrate; split nominal flowrate between top and bottom ducts (see Section A.1).
- Top Ducts 2x Flowrate; combined top duct cooling and doubled total flowrate; simulates a best-case configuration.

These scenarios allow direct comparison between the two upgrade paths and their combination. This flexible framework also supports future evaluations: if HL nominal beam energy is further increased during Run 4, the model can quickly determine how further cooling via combined upgrades would perform.

4.3.2 Uncertainties

This subsection acknowledges and discusses the key sources of uncertainty in the developed model, both in terms of physical assumptions and numerical limitations.

A first source of uncertainty lies in the material properties, in particular those of the CFC used in the low-density section of the core. While efforts were made to quantify these effects and evaluate their influence through sensitivity analyses (see Section 3.3.2), certain values remain based on extrapolated or manufacturer supplied data.

Another important approximation relates to the modeling of argon, which is treated as a solid in the current simulations. This actually excludes convective and buoyancy driven effects as well as neglecting gravity. A separate sub-model or experimental validation would be required but due to time constraints no dedicated study was performed. Nonetheless it is important to state that previous simulations involving nitrogen in the Run 3 dump geometry, where similar assumptions were applied, showed good agreement with experimental benchmarking, providing some precedent for the current approach.

Further uncertainties also arise from the energy deposition maps, which are imported from FLUKA Monte Carlo simulations. As with all statistical models, spatial uncertainty increases along the beam path. While the total integrated energy deposition remains more reliable, it introduces significant variability at the peak level: in some regions towards the downstream beam dump window, local uncertainties on peak energy density can reach up to 50%. This uncertainty is compounded by a systematic 5% energy undershoot introduced during the interpolation of the FLUKA results onto the CFD mesh, as described in Section 3.1.5.1.

Considering now the model related uncertainties, additional uncertainty is introduced by the resolution of the boundary layer, which depends on the structure and quality of the fluid mesh. This has been addressed and quantified through a dedicated mesh sensitivity study (see Section 3.3.3).

Furthermore, the response of the current AHU introduces uncertainty in the massflow boundary condition. Since the AHU operates at constant rotations per minutes, the actual delivered air-flow varies with the pressure drop induced by the cooling duct network. This effect should be further investigated, once the upgrade is installed and the model benchmarked, by extending the sensitivity analysis described in Section 4.3.5, using both reference data provided in Appendix C and measured velocities/pressure drops.

Another noteworthy limitation is the absence of the cradle structure in the thermal model. As shown in the longitudinal section in Figure 4.17, the region corresponding to the cradle,

highlighted by the notch within the black ellipse, contains no solid elements, meaning that no conduction or convection occurs in that area.




Figure 4.17: Longitudinal temperature contour cut through the dump assembly highlighting the artificially elevated temperatures in the region corresponding to the cradle (circled).

This leads to a localized artificial hotspot, to an artificial reduction in contact pressure at the vessel/core interface, given their different CTE, and is also observable in the attempts to correlate pressure to temperature in Section 4.3.4. While this limitation is conservative, since it reduces heat extraction only locally and does not affect a critical region of the dump, it should still ideally be addressed to improve realism. A practical solution could be to either introduce a simplified solid volume representing the cradle or impose an artificial local heat flux sink to allow basic heat transfer.

Finally, and most critically, uncertainty arises from the assumption of perfect thermal contact at the CFC core/vessel interface. While this assumption has been justified analytically (see Section 3.1) and used consistently throughout the study, its physical validity must ultimately be confirmed through experimental validation. To better understand its implications, additional simulations were performed with reduced TCC values based on contact pressure feedback from mechanical models (Section 4.3.4) but nonetheless, this remains one of the most significant areas of uncertainty and will require future experimental verification for full confidence in the model's predictive capability.

4.3.3 Internal Energy and Heat Flux Analysis

This subsection presents an assessment of the thermal performance of the four HL dump configurations studied: Operational 1x Flowrate, Top Ducts 1x Flowrate, Operational 2x Flowrate and Top Ducts 2x Flowrate. All analyses are conducted relative to the Operational 1x baseline case, enabling consistent evaluation of improvements in energy extraction, temperature monitoring and system behavior during a three full intensity HL dumps scenario.

The comparative performances of the configurations is first summarized in Figure 4.18, where both temperature based indicators (maximum and average temperatures for vessel and core) and energy extraction percentage are illustrated.




Figure 4.18: Thermal and energy extraction comparison of the 4 beam dump cooling configurations/upgrades. Each configuration is evaluated based on the peak and average temperatures in the core and vessel, shown as horizontal bars, as well as black diamond markers that indicate the % of beam’s stored energy extracted per dump averaged over three beam dump events.

Among the upgrade scenarios, Top Ducts 1x and Operational 2x emerge as particularly interesting as, although they rely on different strategies, their thermal performance is comparable in several aspects. The Top Ducts 1x configuration achieves a slightly lower maximum vessel temperature (approximately 276°C vs. 298°C) and reduced average vessel temperature (238°C vs. 241°C) compared to Operational 2x. However, these gains are offset by higher core temperatures, with both average and peak values exceeding those in the Operational 2x case. This suggests a slightly different thermal behavior: the top ducts act as a “cold patch” directly targeting the hotspot region, while the Operational 2x case promotes increased convective transfer across the entire bottom surface. This distinction is well captured in the longitudinal heat flux vector fields recorded at the timestep corresponding to the peak vessel temperature for each case. The Top Ducts 1x configuration (Figure 4.19) shows intensified vertical heat flux toward the top surface, whereas the Operational 2x configuration (Figure 4.20) reveals a more distributed and higher-magnitude flux concentrated across the bottom surface.




Figure 4.19: Heat flux vector field of a longitudinal cut plane of the LD sector in Top Ducts 1x configuration at peak vessel temperature timestep.




Figure 4.20: Heat flux vector field of a longitudinal cut plane of the LD sector in Operational 2x configuration at peak vessel temperature timestep.

This difference in cooling strategy is further supported by the surface heat transfer coefficient maps. As visible in Figures 4.21 and 4.22, the Operational 2x configuration achieves significantly higher surface HTC across the entire bottom surface of the dump, augmenting therefore overall convective heat transfer. On the other hand, Figures 4.23 and 4.24 confirm that the Top Ducts 1x configuration increases cooling effectiveness at the top of the dump, but offers limited improvement the bottom surface.




Figure 4.21: HTC contour on the bottom surface. Top Ducts 1x configuration.




Figure 4.22: HTC contour on the bottom surface. Operational 1x configuration.




Figure 4.23: HTC contour on the top surface. Top Ducts 1x configuration.




Figure 4.24: HTC contour on the top surface. Operational 2x configuration.

Internal energy trends in key regions also align with this interpretation. As shown in Figures 4.25 and 4.26, the HD 1 core region, selected as it consistently accumulates the most energy, shows higher internal energy retention in the Top Ducts 1x case compared to Operational 2x. This is also true for the vessel, which retains more energy under the Top Ducts 1x case, confirming that the Operational 2x case is more effective at extracting thermal energy across both the core and vessel.




Figure 4.25: Residual internal energy in HD 1 and vessel. Top Ducts 1x vs. Operational 1x (dashed).




Figure 4.26: Residual internal energy in HD 1 and vessel. Operational 2x vs. Operational 1x (dashed).

In addition valuable insight is also provided by examining the outlet power balance and cumulative energy extraction through the air and radiation channels. To compute energy flows and track the system thermal evolution, the average air temperature at the outlet was first used to calculate the air density at each timestep based on the ideal gas law

$$\rho_{\text{out}} = \frac{p_{\text{abs}}}{R \cdot T_{\text{out}}}$$

where $p_{\text{abs}} = 101'325$ Pa (Section 3.1.5.2), $R = 287.058$ J/kg·K and T_{out} is the average temperature of the outlet itself. This allows for the estimation of the instantaneous mass flow rate $\dot{m}$ from the volumetric flow rate and, given a fixed inlet temperature and the computed $\Delta T = T_{\text{outlet}} - T_{\text{inlet}}$, the instantaneous power extracted by the air:

$$\dot{Q}_{\text{air}} = \dot{m} \cdot c_p \cdot \Delta T$$

with $c_p = 1006.43$ J/kg·K (Section 3.1.5.2). Radiation power $\dot{Q}_{\text{rad}}$ through the enclosure is then added to obtain the total extracted power

$$\dot{Q}_{\text{total}} = \dot{Q}_{\text{air}} + \dot{Q}_{\text{rad}}$$

and energy increments are then obtained at each timestep multiplying by Δt , where Δt corresponds to the time between simulation outputs. Finally the cumulative energy values are derived, by summing these increments over time, and are plotted together with the outlet power balance in the comparative figures for each configuration (Figure 4.27-4.29) offering a direct comparison between the baseline (Operational 1x Flowrate) and each upgraded case.

In all upgraded configurations, the steeper power curves following a dump event during the cooldown phase, especially in the Top Ducts 2x and Operational 2x scenarios, reflect a stronger cooling efficiency and higher HTC enabling faster energy extraction.




Figure 4.27: Time evolution of the heat removal rate by the airflow, along with the cumulative energy extracted by air and lost through radiation. Top Ducts 1x vs Operational 1x (dotted).




Figure 4.28: Time evolution of the heat removal rate by the airflow, along with the cumulative energy extracted by air and lost through radiation. Operational 2x vs. Operational 1x (dotted).




Figure 4.29: Time evolution of the heat removal rate by the airflow, along with the cumulative energy extracted by air and lost through radiation. Top Ducts 2x vs. Operational 1x (dotted).

A similar trend can be seen in the cumulative air energy (red lines), where more efficient configurations extract more heat (for example Figure 4.29). Notably, the radiative energy (magenta line) extracted from the system, decreases as we move from less to more effective cooling setups. This is explained by the fact that lower overall temperatures result in diminished thermal radiation from the dump's outer surfaces. The improved convective cooling reduces radiative output due to the T^4 dependence of radiative flux. However, as clearly demonstrated in the cumulative total energy plots (Figures 4.30-4.32), this reduction in radiative output does not outweigh the gain in convective extraction: total cumulative energy at the outlet still increases with more efficient configurations.




Figure 4.30: Total cumulative extracted energy over time. Top Ducts 1x vs. Operational 1x (dotted).




Figure 4.31: Total cumulative extracted energy over time. Operational 2x vs. Operational 1x (dotted).




Figure 4.32: Total cumulative extracted energy over time. Top Ducts 2x vs. Operational 1x (dotted).

In conclusion, all performance metrics consistently indicate that the Top Ducts 2x configuration offers the best overall cooling efficiency, extracting approximately 18.2 % more energy per dump than the baseline Operational 1x configuration. The Operational 2x and Top Ducts 1x cases also show notable improvements, with increases of 14.43% and 3.82%, respectively. While the overall energy extracted in the Top Ducts 1x case is smaller, this configuration delivers a distinct advantage in terms of maximum vessel temperature reduction. By directly targeting the hotspot region on the upper surface, the added top ducts serve as a local heat sink, lowering critical peak temperatures without requiring any modifications to the existing AHU infrastructure. In contrast, the Operational 2x configuration promotes more cooling across the dump, reducing internal energy accumulation more broadly through enhanced convective heat transfer. The Top Ducts 2x configuration, combining both strategies, remains the most effective overall. These findings will be further discussed in Section 5, where final recommendations are made based on both thermal performance and implementation feasibility.

4.3.4 Contact Pressure Analyses

This subsection explores the thermal and mechanical behavior at the vessel/CFC core interface, focusing on how variations in contact pressure influence both heat transfer efficiency and assembly integrity. On the thermal side, contact pressure directly affects the TCC, which in turn governs the efficiency of heat extraction from the carbon core to the titanium vessel. Mechanically, the focus shifts to the much faster timescale of a beam dump event, where the structural response to the beam impact of the core becomes critical. Simulations conducted by the SY-STI TCD TDE team indicate that beam impacts can generate accelerations in the core up to 100 g , potentially leading to plates displacement inside the dump.

To investigate the thermo-mechanical coupling, temperature fields were extracted from the CFD simulations at key points during the transient simulations. These include point A (maximum vessel temperature of the 2nd dump), B (just before the 3rd dump), C (maximum vessel temperature over the entire simulation) and D (exactly at the end of the energy deposition of the 3rd dump event). Figure 4.33 & 4.34 illustrates the timeline/placement of these points for reference. The corresponding thermal fields were used as input for mechanical simulations performed by an expert within the SY-STI-TCD TDE team. This process, outlined schematically in Figure 3.1, allowed feedback from the mechanical model in the form of contact pressure maps. These maps were then used both to determine the mechanical integrity of the assembly and, in combination with temperature maps, to explore a pressure-temperature correlation for estimating temperature-dependent TCC and assess the influence of the perfect contact assumption outlined in Section 3.1.3.




Figure 4.33: Standard 3 dumps cycle highlighting reference points A, B and C.




Figure 4.34: Detail on evaluation point D just after the energy deposition of the 3rd dump event.

Given that the thermal CFD simulations assumed perfect contact and that the structural mechanical model considered the loosest possible fit between vessel and core considering the manufacturing tolerances, the resulting contact pressure maps represent a conservative, worst-case scenario. For instance, the pressure map shown in Figure 4.35 (Operational 1x, point C) reveals many regions at or near 0 MPa, clearly highlighting that perfect contact cannot be guaranteed. In contrast at the same time point for the Top Ducts 1x case (Figure 4.36), the number of low contact pressure regions is visibly reduced (aside from the artifact present in the section near CFC 2, see discussion in Section 4.3.2). Finally the pressure map for the Operational 2x case (Figure 4.37) shows even better contact retention.

Both upgraded configurations outperform the baseline Operational 1x case when evaluated at point C. More importantly, further analysis of the contact pressure maps reveals that both upgraded setups maintain better contact retention after three dumps (point C) than the baseline configuration does already after two dumps (point A), highlighting their significantly improved mechanical behavior throughout the cycle.




Figure 4.35: Side view of the contact pressure contour map for the CFC LD sector in the Operational 1x configuration at point C. The surface was reconstructed from point cloud data using a Delaunay 3D algorithm, which introduced artifacts (grey regions and fictitious high values of contact pressure) in the gaps between stacked CFC stacks. Black areas indicate instead regions with contact pressures at or near 0 MPa.




Figure 4.36: Side view of the contact pressure contour map for the CFC LD sector in the Top Ducts 1x configuration at point C. On top of the grey area artifacts introduced by Delaunay 3D, the black area near CFC 2 corresponds to the modeling artifact discussed in Section 4.3.2, not to an actual loss of contact pressure.




Figure 4.37: Side view of the contact pressure contour map for the CFC LD sector in the Operational 2x configuration at point C. The contour is visibly improved compared to the 1x case, with no regions at or near 0 MPa (except model artifact).

While mechanical simulations confirm better contact in upgraded configurations, their thermal implication and influence on the developed CFD model had to be tested. To this end, two constant TCC values were assigned at the vessel/core interface in the entirety of the CFC LD section to simulate contact degradation: $8 \text{ W/m}^2\text{K}$, representing a worst case scenario of a "floating" core with a 2 mm argon gap to the vessel, and $20 \text{ W/m}^2\text{K}$, a more realistic estimate based on past studies [24]. These values were applied in separate simulations for both Operational 1x and Top Ducts 1x, with results confirming a clear trend: lower TCC leads to lower vessel temperatures and higher core temperatures due to reduced heat extraction and to thermal accumulation in the core.

Furthermore, to develop a more dynamic and localized model of thermal contact resistance, a correlation between temperature and contact pressure was attempted. This analysis was limited to two configurations: Operational 1x, being the baseline, and Top Ducts 1x, which showed the least favorable contact pressure distribution among the upgrade cases. The remaining configurations were not considered due to the high computational cost and the assumption that their improved pressure distributions would yield more favorable results. For both selected cases, contact pressure and temperature values were extracted and coupled pointwise, then segmented by dump sector. Each segment was analyzed using double box plots, enabling the derivation of approximate $P(T)$ relationships. The resulting 2D boxplots, along with the fitted functions for Operational 1x and Top Ducts 1x, are shown in Figures 4.38 and 4.39, respectively.




Figure 4.38: 2D box plots of contact pressure versus temperature for the Operational 1x configuration at point C. The overlaid curve represents the fitted $P(T)$ relationship, with the corresponding equation reported.




Figure 4.39: 2D box plots of contact pressure versus temperature for the Top Duct 1x configuration. Data are matched point wise between the 2 maps and grouped by dump sector.

These $P(T)$ functions were then used to compute a temperature-dependent TCC using the Mikic model for elastic asperity deformation [34]:

$$TCC(P) = 1.55 \left(\frac{k \tan \theta}{\sigma} \right) \left(\frac{P\sqrt{2}}{E' \tan \theta} \right)^{0.94}$$

where:

- k is the harmonic mean of the thermal conductivities of titanium and carbon-carbon,
- σ is the standard deviation of the surface profile height,
- θ is the mean profile slope,
- E' is the equivalent elastic modulus of the two materials in contact and
- P is the local contact pressure.

All assumed parameters are listed in Table 3.1. Once computed, the resulting $TCC(T)$ values were converted to effective thermal temperature dependent conductivity to assign to a 4 mm thick dummy layer in the CFD model between the CFC core and titanium vessel (Appendix A.2):

$$k_{\text{dummy}}(T) = TCC(T) \cdot 0.004$$

The extended numerical calculations are summarized in Appendix A.3.

Corresponding results, shown in Figures 4.40 and 4.41, display the evolution of maximum vessel and core temperatures as well as remaining internal energy for three thermal contact conductance scenarios applied to a single configuration: perfect contact (infinite TCC), worst-case floating core contact (8 W/m²K), and a variable TCC as a function of temperature derived from pressure-temperature coupling.

The perfect contact case results in a significant heat transfer from the core to the vessel, as evidenced by a second noticeable rise in maximum vessel temperature after the first dump not present in the other 2 scenarios. This indicates enhanced thermal conduction across the interface, particularly in the carbon-carbon region, where heat from the core is transferred efficiently to the vessel and heats up the corresponding vessel section.

In contrast, the worst-case scenario with a constant TCC of 8 W/m²K shows a clear thermal decoupling between the core and the vessel in the LD section. The core remains highly heated as most of the energy remains trapped within it, while the vessel temperature rise is limited. The maximum vessel temperature coincides with the initial energy deposition in the high density graphite (HD 1 section), and no secondary peak is observed.

The variable TCC case lies between these two extremes. Already towards the end of the first dump, a distinct kink appears in the maximum vessel temperature curve, indicating a shift in hotspot location along the vessel. Specifically, the hottest regions migrate toward the CFC 3 and CFC 4 sections, consistent with the spatial variation in temperature and corresponding contact pressure. Meanwhile, the core temperature for this case remains between the perfect and worst-case scenarios.

This behavior highlights a self-mitigating mechanism: as the vessel heats up, contact pressure decreases, lowering TCC and reducing heat transfer to the vessel. This, in turn, limits further vessel heating. However, as the vessel cools and contracts, contact pressure and TCC recover, enhancing heat extraction once again. The absence of runaway effects or positive feedback loops demonstrates the robustness of the thermal interface under transient conditions and, while a prolonged loss of contact could delay core cooldown after the third dump, the dynamic and intermittent nature of the thermal behavior ensures that contact is eventually reestablished. This interpretation is further confirmed by the internal energy evolution shown in Figure 4.41, which demonstrates that higher TCC yields faster heat removal and lower residual energy over time.




Figure 4.40: Maximum vessel and core temperatures over time for the Operational 1x configuration, evaluated under three different thermal contact assumptions: perfect contact, floating core (no contact) and a variable TCC(T) derived from the elastic contact model.




Figure 4.41: Dump block residual internal energy over time for the Operational 1x configuration, under the same three thermal contact assumptions as in Figure 4.40.

To complement the thermal analysis, a mechanical evaluation was conducted to assess the retention and stability of the CFC plates within the titanium vessel during beam impact events. The primary concern is whether the CFC plates, shrink-fitted into the vessel, remain securely in place under axial accelerations that can reach up to 100 g during a beam impact, as predicted by *Ansys LS-DYNA* simulations. The analysis therefore focuses on a specific point D in the simulation timeline, immediately after the third beam dump (Figure 4.34). This represents for the CFC plates the moment of weakest contact pressure when being subjected to the next beam impact and the consequent axial accelerations.

The mechanical analysis evaluates therefore the maximum axial acceleration sustainable by each contact point before slippage occurs. This is based on a local force balance between the inertial force acting on the plate mass and the frictional resistance, which is estimated from the contact

pressure and an assumed coefficient of static friction. The local contact pressure at each point, $P_{\text{contact}}(x, y, z)$, is used to compute the maximum axial acceleration according to

$$a_{\text{max}}(x, y, z) = \mu \cdot P_{\text{contact}}(x, y, z) \cdot \frac{A_{\text{contact}}}{m_{\text{plate}}}$$

where:

- $\mu = 0.1$ is the static friction coefficient assumed between CFC and titanium,
- $A_{\text{contact}} = 0.1032 \text{ m}^2$ is the contact area per plate,
- $m_{\text{plate}} = 21.1 \text{ kg}$ is the mass of each CFC plate.

This formulation assumes a uniform distribution of mass across the contact interface of each plate and the calculation is applied to the entire contact pressure map. However, this includes regions outside actual contact zones, such as argon gaps or plate edges, which can produce unrealistic values. These outliers were identified and filtered out to avoid affecting the interpretation of the results.

To identify areas most susceptible to slippage, the pointwise values of $a_{\text{max}}(x, y, z)$ were axially averaged over the xy planes to produce a longitudinal z-coordinate dependent profile:

$$\overline{a_{\text{max}}}(z) = \text{mean} (a_{\text{max}}(x, y, z) \mid z = \text{const})$$

This profile provides a spatial map of holding capacity of xy slices of core throughout the dump length, allowing for the localization of sections most vulnerable to slippage during beam impact events. Crucially, by overlaying these predicted maximum allowed axial acceleration profiles $\overline{a_{\text{max}}}(z)$ with the corresponding simulated profile of maximum axial accelerations experienced during beam impact, it is possible to assess whether all axial positions remain within safe margins (Figure 4.42 & 4.43).




Figure 4.42: Predicted maximum allowed axial acceleration for xy slice of LD core $\overline{a_{\max}}(z)$ (blue) and weighted polynomial fit (orange) compared to actual peak axial acceleration (green) for the Operational 1x configuration.




Figure 4.43: Predicted maximum allowed axial acceleration for xy slice of core $\overline{a_{\max}}(z)$ (blue) and weighted polynomial fit (orange) compared to actual peak axial acceleration (green) for the Top Ducts 1x configuration.

These calculations were derived using a static friction coefficient of $\mu = 0.1$, which, at the time of writing, was still undergoing experimental characterization. It is important to note that the estimated holding capacity (orange curve) is directly proportional to this coefficient. Therefore, any refinement in its measured value would significantly shift the orange trend-line up or down, influencing the predicted safety margin.

In both configurations, the artificial drop near $z = 2.4$ m corresponds to the contact modeling artifact previously discussed in Section 4.3.2, highlighting a known limitation in the interface between the support system and the CFC 2 section of the vessel. Additionally, isolated outliers are visible near the end plates of the stacks and within the vessel regions between adjacent CFC stacks, where no physical contact exists. These points were excluded from the weighted polynomial fit and subsequent interpretation.

Finally, despite being generated under the most conservative scenario, using thermal fields from simulations with perfect contact assumptions and mechanical feedback from the loosest plausible shrink-fit, the Operational 1x configuration still maintains a small safety margin with average sustainable acceleration staying above the actual beam-induced load across most of the dump. The Top Ducts 1x configuration shows a substantial improvement in mechanical retention, with the polynomial fit rising above the critical acceleration curve by a wider margin across the full

length of the dump. This confirms that enhanced thermal contact in the upgraded configuration not only improves cooling efficiency but also contributes to greater mechanical stability, reducing the risk of plate displacement under beam impact.

4.3.5 Sensitivity to Massflow Distribution

The addition of top ducts offers significant advantages not only in thermal performance but also in pressure drop reduction. With a larger total outlet area, the same massflow rate can be delivered with less resistance, reducing ΔP across the system. This improved margin opens the door for a potential optimization of the massflow distribution between the top and bottom ducts by inserting tuning valves. To investigate this, four additional configurations (Table 4.2) were simulated alongside the baseline Top Ducts 1x case, varying the proportion of flow sent to the top ducts while maintaining a constant total massflow of 1.7 kg/s.

Name	Bottom duct massflow [kg/s]	Top duct massflow [kg/s]	Total massflow [kg/s]	Top/Total massflow [%]	Total ΔP [Pa]	Comments
Case 1	1.370	0.330	1.700	19.4	167.51	+7.6% to top
Case 2	1.435	0.265	1.700	15.6	173.15	+3.8% to top
Baseline	1.500	0.200	1.700	11.8	165.72	Baseline
Case 3	1.565	0.135	1.700	7.9	167.58	-4.0% to top
Case 4	1.630	0.070	1.700	4.1	174.20	-7.7% to top

Table 4.2: Overview of the tested airflow distribution scenarios for the massflow sensitivity study, showing the assigned massflows to bottom and top ducts and the resulting pressure drops.

The results are presented in Figure 4.44, which shows the evolution of maximum vessel and core temperatures across the three dump cycles for various top/bottom duct splits.

As expected, increased top duct flow consistently leads to improved thermal performance: both maximum core and vessel temperatures decrease as more air is redirected toward the top, confirming that shifting the cooling effort upward enhances overall heat removal. This trend is further supported by Figure 4.45, which displays the residual internal energy in the dump block: as more flow is allocated to the top ducts, less energy remains in the system after each dump, clearly indicating more efficient thermal extraction. However, this trend must be interpreted within practical design constraints. From a fluid dynamics perspective, the more flow is sent to the top ducts, the higher the pressure drop. While beneficial from a cooling standpoint, this pressure drop must remain within the specifications of the current AHU, as detailed in Appendix C. Simultaneously, excessive redirection of flow to the top could eventually under-serve the bottom cooling region, leading to localized overheating in sections that are no longer adequately cooled.

Finally from the plotted results, the spacing between the performance curves becomes progressively narrower as more flow is shifted towards the top ducts, suggesting that the system is approaching an optimum. Once the upgraded top duct system is eventually installed and the real pressure/flow characteristics are benchmarked, the massflow distribution can be fine-tuned using the actual AHU response curve (Appendix C) as a boundary condition input for the model. At that stage, a dedicated optimization study should be carried out to set the correct valve positions. The present study therefore not only demonstrates that this type of fine-tuning is possible, but also that it is meaningful and potentially beneficial. The implementation of fine-tuning should be seriously considered in the final installation phase.




Figure 4.44: Time evolution of maximum core and vessel temperatures for different massflow distribution scenarios. Thermal performance improves as more flow is directed to the top ducts, with both core and vessel temperatures decreasing progressively.




Figure 4.45: Residual internal energy in the dump block over time for different top/bottom duct flow splits.

Chapter 5

Conclusion & Outlook

5.1 Conclusion

This thesis presented a comprehensive CFD modeling framework for the HL-LHC beam dump cooling system. This new computational pipeline is capable of simulating complex multi-material assemblies with varying thermal contact resistance, diverse duct configurations and a flexible structure designed to support upgrades. Building upon a validated Run 3 model and the developed HL configuration, this work enabled a detailed evaluation of multiple cooling layouts, taking into account thermal, mechanical, and operational constraints.

Combining thermal performance gains with operational feasibility, the top ducts upgrade is identified as the recommended solution for HL-LHC integration. It offers improvement including:

- Peak and average vessel temperatures were considerably reduced in key locations, improving safety margins under consecutive dump scenarios;
- TCC between LD core and vessel was improved due to the resulting increase in contact pressures, which is critical in reducing thermal bottlenecks and restoring energy extraction pathways;
- Mechanical retention of the CFC stacks at the core/vessel interface was improved, as thermal contraction of the upper vessel increased contact pressure, mitigating the contact loss observed in the non-upgraded configuration.

Moreover from an engineering integration perspective, the top ducts stand out as the most practical and feasible solution.

- Minimal intervention in the existing AHU infrastructure, ducting or ventilation system is required as there is no increase in total flowrate.
- The top ducts only require integration of additional components within the existing infrastructure, enabling rapid installation without extensive mechanical modifications, minimizing work required inside the cavern and reducing radiation exposure to personnel.
- Compared to doubling the flowrate, the top duct solution offers a low-disruption, high impact alternative that is better aligned with HL-LHC installation time schedule and significantly more cost-effective

The variable TCC model, defined as a function of local contact pressure, represents a major shift from the previous assumption of perfect thermal contact. By analyzing the evolution of pressure

maps and their correlation with temperature fields, the model revealed zones of intermittent contact and enabled a more detailed assessment of effective heat transfer through the dump's layered structure.

Lastly, the framework developed here also enabled critical assessment of various flow-splitting strategies, vessel/core thermal coupling mechanisms, and energy extraction metrics across realistic 3-dump operational cycles, and is now ready to support HL-LHC design monitoring as well as future benchmarking of cooling system configurations.

5.2 Outlook

The simulation framework developed in this thesis provides a solid basis for continued refinement of the HL beam dump model. Based on previously discussed limitations of the model, the following directions are recommended:

- Tune and validate thermal contact models using vessel-mounted PT100 data. The current temperature-TCC correlation may be adjusted or replaced with a plastic contact formulation, depending on post-installation measurements.
- Measure the coefficient of friction between the vessel and CFC plates to improve accuracy of contact pressure predictions and mechanical retention analysis.
- Characterize flow conditions in the ducts through direct measurement of pressure and velocity fields after installation. This data will enable realistic inlet boundary conditions based on the AHU curve and guide top/bottom duct flow balancing and optimization.
- Implement coupled thermal-mechanical simulations to refine contact pressure and temperature maps over time beyond worst-case scenarios.
- Model the thermal behavior of the cradle region either by explicitly introducing it as a solid domain or by applying a simplified approximation, in order to eliminate known artifacts in the CFC 2 and HD 3 sectors.

These developments will further improve the predictive capability and reliability of the model, supporting operational monitoring and possible future iterations of the HL-LHC dump system.

Bibliography

- [1] CERN. "*The Large Hadron Collider*". <https://home.cern/science/accelerators/large-hadron-collider>. Last Accessed: 2025-05-19.
- [2] Oliver S. Bruning et al. "The LHC Main Ring". In: *LHC Design Report Vol.1*. 2004. URL: <https://api.semanticscholar.org/CorpusID:221242712>.
- [3] SYMMETRY dimensions of particle physics. "*Protecting the LHC from itself*". https://www.symmetrymagazine.org/article/december-2007/protecting-lhc-itself?language_content_entity=und. Last Accessed: 2024-12-26.
- [4] Lyndon Evans and Philip Bryant. "LHC Machine". In: *Journal of Instrumentation* 3.08 (2008), S08001. DOI: 10.1088/1748-0221/3/08/S08001. URL: <https://dx.doi.org/10.1088/1748-0221/3/08/S08001>.
- [5] Marco Calviani et al. "*LHC External Beam Dumps (TDE) Modifications and upgrades during LS2 for Run3*". Sept. 2023. DOI: 10.5281/zenodo.8365957. URL: <https://doi.org/10.5281/zenodo.8365957>.
- [6] E Farina et al. "Energy deposition challenges for the HL-LHC beam dump". In: *JACoW IPAC 2023* (2023), MOPA101. DOI: 10.18429/JACoW-IPAC2023-MOPA101. URL: <https://cds.cern.ch/record/2886005>.
- [7] J. Maestre et al. "Design and behaviour of the Large Hadron Collider external beam dumps capable of receiving 539 MJ/dump". In: *Journal of Instrumentation* (2021). DOI: 10.1088/1748-0221/16/11/P11019.
- [8] O. Aberle et al. "*High-Luminosity Large Hadron Collider (HL-LHC): Technical design report*". CERN Yellow Reports: Monographs. Geneva: CERN, 2020. DOI: 10.23731/CYRM-2020-0010. URL: <https://cds.cern.ch/record/2749422>.
- [9] C. Bracco et al. "*Functional Specification: TDE requirements for operation with HL-LHC beams*". <https://edms.cern.ch/document/2694263/1.0>. Last Accessed: 2025-1-16.
- [10] Marco Calviani. "*Beam Intercepting Devices at CERN – Types, Challenges, Design, RD and Operation, Academic Training Lecture Regular Programme ; 2021-2022*". <https://cds.cern.ch/record/27910>. Last Accessed: 2025-03-05.
- [11] Alvaro Romero Francia et al. "Design and early operation of a new-generation internal beam dump for CERN's Super Proton Synchrotron". In: *Physical Review Accelerators and Beams* 27 (Apr. 2024). DOI: 10.1103/PhysRevAccelBeams.27.043001.
- [12] Giuseppe Battistoni et al. "Overview of the FLUKA code". In: *Annals of Nuclear Energy* 82 (2015). Joint International Conference on Supercomputing in Nuclear Applications and Monte Carlo 2013, SNA + MC 2013. Pluri- and Trans-disciplinarity, Towards New Modeling and Numerical Simulation Paradigms, pp. 10–18. ISSN: 0306-4549. DOI: <https://doi.org/10.1016/j.anucene.2014.11.007>. URL: <https://www.sciencedirect.com/science/article/pii/S0306454914005878>.
- [13] J.M. Zazula and S. Peraire. "*LHC BEAM DUMP DESIGN STUDY Part I : Simulation of energy deposition by particle cascades; implications for the dump core and beam sweeping system*". <https://cds.cern.ch/record/315077/files/lhc-project-report-80.pdf?version=1>. Last Accessed: 2024-12-26.

- [14] J.M. Zazula and S. Peraire. "LHC Beam Dump Design Study, Part II: Thermal analysis; implication for abort repetition and cooling system." <https://lib-extopc.kek.jp/preprints/PDF/1997/9702/9702111.pdf>. Last Accessed: 2024-12-26.
- [15] N. Solieri et al. "Decommissioning and Post-Irradiation Examination of the LHC Beam Dumps". 2025. arXiv: 2505.04235 [physics.acc-ph]. URL: <https://arxiv.org/abs/2505.04235>.
- [16] Christoph Wiesner et al. "Machine protection aspects of high-voltage flashovers of the LHC beam dump dilution kickers". In: *10th International Particle Accelerator Conference*. 2019, WEPMP040. DOI: 10.18429/JACoW-IPAC2019-WEPMP040.
- [17] CERN. "Longer term LHC schedule". <https://lhc-commissioning.web.cern.ch/schedule/LHC-long-term.htm>. Last Accessed: 2025-05-20.
- [18] G Banks et al. "Design development and R&D for CERN's HL-LHC external beam dump". In: *JACoW IPAC 2023* (2023), TUPM122.
- [19] K. Wang et al. "Prediction of short fatigue crack growth of Ti-6Al-4V". In: *Fatigue & Fracture of Engineering Materials & Structures* 37.10 (2014), pp. 1075–1086. DOI: <https://doi.org/10.1111/ffe.12177>. eprint: <https://onlinelibrary.wiley.com/doi/pdf/10.1111/ffe.12177>. URL: <https://onlinelibrary.wiley.com/doi/abs/10.1111/ffe.12177>.
- [20] Asbjørn Lassen Lund. "Upgrades of the Target Dump Externals for the Large Hadron Collider". Presented 27 Jan 2023. Aarhus U., 2022. URL: <https://cds.cern.ch/record/2851843>.
- [21] Pablo Andreu Muñoz et al. "Irradiation of Low-Z Carbon-Based Materials with 440 GeV/c Proton Beam for High Energy amp; Intensity Beam Absorbers: The CERN HiRadMat-56-HED Experiment". In: *JACoW IPAC 2022* (2022), pp. 2883–2886. DOI: 10.18429/JACoW-IPAC2022-THP0TK049. URL: <https://cds.cern.ch/record/2845794>.
- [22] Jorge Heredia et al. "Sigraflex® Studies for LHC CERN Beam Dump: Summary and Perspective. Sigraflex® STUDIES FOR LHC CERN BEAM DUMP: SUMMARY AND PERSPECTIVE". In: *JACoW IPAC 2021* (2021), pp. 3571–3574. DOI: 10.18429/JACoW-IPAC2021-WEPAB368. URL: <https://cds.cern.ch/record/2810841>.
- [23] Gréard J Poitras et al. "Aerodynamic and heat transfer analysis of a impinging jet on a concave surface". In: *International Journal of Thermal Sciences* 114 (2017), pp. 184–195.
- [24] Leonhard Gross. "CFD Cooling Simulations for the LHC Beam Dumps under High Luminosity Conditions". Last accessed: 2025-04-18.
- [25] M. M. Yovanovich and E. E. Marotta. "Thermal Spreading and Contact Resistances". In: *Heat Transfer Handbook*. Ed. by Adrian Bejan and Allan D. Kraus. Hoboken, NJ: John Wiley & Sons, 2003, pp. 261–385. ISBN: 0-471-39015-1.
- [26] ASM Aerospace Specification Metals Inc. "Titanium Ti-6Al-4V (Grade 5), Annealed". <https://asm.matweb.com/search/specificmaterial.asp?bassnum=mtp641>. Last Accessed: 2025-05-19.
- [27] S. S. d. Rocha et al. "Vickers hardness of cast commercially pure titanium and ti-6al-4v alloy submitted to heat treatments". In: *Brazilian Dental Journal* 17 (2 2006), pp. 126–129. DOI: 10.1590/s0103-64402006000200008.
- [28] Myungjun Kim and Yongha Kim. "A Thermo-Mechanical Properties Evaluation of Multi-Directional Carbon/Carbon Composite Materials in Aerospace Applications". In: *Aerospace* 9.8 (2022). ISSN: 2226-4310. DOI: 10.3390/aerospace9080461. URL: <https://www.mdpi.com/2226-4310/9/8/461>.
- [29] Ansys FLuent. "Ansys Fluent 12.0 Theory Guide". <https://www.afs.enea.it/project/neptunius/docs/fluent/html/th/node67.htm>. Last Accessed: 2025-03-12.

- [30] Engineering Toolbox. *"Emissivity Coefficients of Common Materials: Data Reference Guide"*. https://www.engineeringtoolbox.com/emissivity-coefficients-d_447.html. Last Accessed: 2025-04-20.
- [31] Franqueira Ximenes. *"Fluka To Ansys Tables"*. https://gitlab.cern.ch/SY_STI_TCD/Fluka-to-ansys-tables. Last Accessed: 2025-03-04.
- [32] Oscar Sacrist an de Frutos Stefan Hoell. *"Characterization of Graphitic materials for the HiLumi Beam Dump Market Survey"*. https://edms.cern.ch/ui/file/3064824/1/EDMS_3064824_Market_survey_on_isostatic_graphite_and_CFC_FINAL-2024-11-18.pdf. Last Accessed: 2025-04-03.
- [33] SGL Carbon. *"SGL Carbon Website"*. <https://www.sglcarbon.com/en/>. Last Accessed: 2025-04-15.
- [34] COMSOL. *"Theory for Thermal Contact"*. https://doc.comsol.com/5.5/doc/com.comsol.help.heat/heat_ug_theory.07.66.html. Last Accessed: 2025-05-19.

Appendix A

A.1 Simplified Steady State Calculations

To better estimate the massflow distribution between top and bottom ducts in the Top Ducts 1x configuration, a set of simplified steady-state CFD simulations was conducted. As the aim was to isolate and simulate only the fluid domain inside the ducting network, reducing computational cost while capturing the essential behavior of the system, the thermal solid domain and transient effects were neglected.

The initial geometry, shown in Figure A.1, includes only a top and a bottom duct omitting the upstream ducting. Inlet boundary conditions were tuned to reflect the inlet surface area differences between top and bottom ducts, aiming to balance the pressure drop between the two networks and reach therefore the equilibrium point of the system. This tuning led to a fixed inlet massflow of 0.1 kg/s at the top duct and 0.75 kg/s at the bottom duct. Doubling these values to account for symmetry results in a total of 1.7 kg/s of air which, assuming standard conditions air density, corresponds to approximately 4'096 m³/h, consistent with the design target of 5'000 m³/h.




Figure A.1: Simplified fluid-only duct geometry used to estimate massflow distribution under symmetry assumption. Inlets highlighted in orange.

To assess the influence of the upstream ducting on the flow and pressure drop behavior, a more realistic geometry that included a preliminary version of upstream ducting was created (Figure A.2). A new steady-state simulation was therefore performed using the same total massflow rate

of 1.7 kg/s. The resulting outlet massflow rates for the four branches (top left, top right, bottom left, bottom right) are summarized in Table A.1. These results show that ducts positioned further away from the main inlet exhibit slightly higher massflow, likely due to reduced flow resistance after the first branching. The deviation from the assumed symmetric configuration remains within 13% for top ducts and 4.5% for bottom ducts, confirming that the original assumption provides a reasonably accurate baseline.

Outlet Name	Massflow rate [kg/s]	Deviation from assumption [%]
Top left	0.1139	13.8
Top right	0.1051	5.0
Bottom left	0.7165	-4.5
Bottom right	0.7656	2

Table A.1: Massflow distribution at each duct outlet in the upstream-inclusive simulation.




Figure A.2: Updated geometry including preliminary upstream ducting for refined massflow predictions.

Additionally, a swirling motion was observed leading to rotational velocity components in the duct (Figure A.3). However, this swirl is not expected to significantly affect the overall cooling performance. The reason lies in the downstream design: as air exits through the nozzles distributed along the dump length, the dominant velocity direction becomes axial through the nozzle exits, effectively dampening the swirl-induced asymmetry before it can influence the impingement pattern. Although the upstream ducting geometry is still under refinement, this analysis highlights its potential impact on inlet boundary conditions and shows that it should be revisited for the final design.




Figure A.3: Swirling phenomena observable within the ducts.

To complete, an additional set of simulations was performed to investigate the effect of internal surface roughness on the overall pressure drop across the duct assembly. The tested cases and corresponding results are summarized in Table A.2. From this data, we observe that adding the upstream ducting geometry increases the total pressure drop significantly, from 165 Pa in the original simplified case (Table 4.2) to 213 Pa with the upstream ducting included. Increasing the roughness height further leads to additional pressure drop, though all cases remain within acceptable limits. This confirms that internal roughness contributes measurably to the total resistance and should be considered in the final design.

Name	Assigned roughness height [m]	Total pressure drop [Pa]
Baseline	0	213.32
Medium	$1.5 \cdot 10^{-4}$	222.82
High	$3.0 \cdot 10^{-4}$	229.93

Table A.2: Summary of duct surface roughness cases and their impact on total pressure drop.

A.2 Dummy Elements Layers Insertion

To enable the implementation of temperature-dependent TCC models between the CFC core and the surrounding titanium vessel, dedicated dummy elements were inserted between each CFC monolithic block and the vessel enclosure. These interface regions, shown in green in Figure A.4, were created by cutting out a 4 mm thick layer from the outer edge of each CFC block (pink), preserving the total volume of the dump geometry. The original CFC block zones were then replaced in the original meshes by two distinct mesh regions: one for the inner CFC core and one for the surrounding dummy interface layer.

The corresponding solid mesh, illustrated in Figure A.5, applies the same mesh-independent control parameters used throughout the model (see Section 3.1.4) to ensure consistency across simulations.

Finally, to maintain energy conservation, the energy deposition profiles originally applied to the full CFC blocks were adapted such that these dummy regions continue to receive the appropriate fraction of deposited energy. In this way, the integrated energy deposition across the block remains invariant. The dummy material inherits all thermal properties from the original CFC except for the thermal conductivity, which was adjusted to replicate the desired temperature-dependent TCC behavior. The corresponding conductivity values used for the dummy material are listed in Appendix A.3.




Figure A.4: CAD representation of the CFC core (pink) and the surrounding dummy layers (green) added to implement temperature-dependent thermal contact behavior.




Figure A.5: Corresponding solid mesh of one block in the model, showing the interface dummy layers (green) between CFC (red) and the titanium vessel.

A.3 Calculations for Dummy Layer Conductivities

This derivation provides an expression for the thermal conductivity $k_{\text{dummy}}(T)$ of a dummy layer of elements, based on a contact model and temperature-contact pressure curve fits from two configurations: Operational 1x and Top Ducts 1x (Figure 4.38 & 4.39). The approach unfolds through the following steps:

1. Fit a curve for pressure as a function of temperature: $P(T)$
2. Use the Mikic model for thermal contact conductance as a function of pressure: $\text{TCC}(P)$

3. Compose to obtain $TCC(T)$
4. Compute the dummy material conductivity as $k_{\text{dummy}}(T) = TCC(T) \cdot \delta$

The contact conductance model, assuming parameters in Table 3.1, is given by:

$$TCC(P) = \frac{1.55 \cdot 8.59893048 \cdot 0.2}{5 \cdot 10^{-6}} \cdot \left(\frac{P \cdot \sqrt{2}}{1.5317 \cdot 10^{10}} \cdot 0.2 \right)^{0.94}$$

simplifying the expression becomes

$$TCC(P) = 532733.25 \cdot (P \cdot 1.846 \cdot 10^{-11})^{0.94}$$

For the **Operational 1x** configuration, the fitted pressure-temperature relation is (Figure 4.38):

$$P(T) = 2.40 \cdot 10^1 \cdot T^2 - 5.08 \cdot 10^4 \cdot T + 2.00 \cdot 10^7$$

so substituting into the TCC model it becomes

$$TCC(T) = 532733.25 \cdot ((2.40 \cdot 10^1 \cdot T^2 - 5.08 \cdot 10^4 \cdot T + 2.00 \cdot 10^7) \cdot 1.846 \cdot 10^{-11})^{0.94}$$

Assuming a dummy element thickness $\delta = 0.004$ m, the dummy material conductivity becomes:

$$k_{\text{dummy}}(T) = 0.004 \cdot TCC(T)$$

therefore

$$k_{\text{dummy}}(T) = 0.004 \cdot 532733.25 \cdot ((2.40 \cdot 10^1 \cdot T^2 - 5.08 \cdot 10^4 \cdot T + 2.00 \cdot 10^7) \cdot 1.846 \cdot 10^{-11})^{0.94}$$

For the **Top ducts 1x** configuration instead, the pressure-temperature fit is (Figure 4.39):

$$P(T) = -2.34 \cdot 10^1 \cdot T^2 - 5.51 \cdot 10^3 \cdot T + 9.28 \cdot 10^6$$

leading to

$$TCC(T) = 532733.25 \cdot ((-2.34 \cdot 10^1 \cdot T^2 - 5.51 \cdot 10^3 \cdot T + 9.28 \cdot 10^6) \cdot 1.846 \cdot 10^{-11})^{0.94}$$

and the corresponding dummy material conductivity is therefore

$$k_{\text{dummy}}(T) = 0.004 \cdot 532733.25 \cdot ((-2.34 \cdot 10^1 \cdot T^2 - 5.51 \cdot 10^3 \cdot T + 9.28 \cdot 10^6) \cdot 1.846 \cdot 10^{-11})^{0.94}$$

Finally to ensure numerical stability in simulation environments, the use of a capped expression is therefore recommended:

$$k_{\text{dummy}}(T) = 0.004 \cdot 532733.25 \cdot (\max(P(T) \cdot 1.846 \cdot 10^{-11}, 0))^{0.94}$$

this ensures that conductivity is never evaluated with a negative pressure value, which would result in nonphysical behavior or divergence in *Ansys Fluent*.

Appendix B

B.1 Previous Work

This set of simulations were carried out using the Run 3 spare beam dump geometry and material configuration. The vessel was modeled in SS 318LN and the low-density region of the core was represented using Sigraflex[®]. TCC values at the interface between the Sigraflex[®] and the vessel were assigned based on previously experimentally validated data [24].

The air flow rate through the cooling ducts was fixed at 5000 m³/h, matching the nominal operational capacity of the current AHU. All scenarios in this set of simulations share the same boundary conditions and material assumptions; the only parameter varied from case to case is the geometry of the cooling ducts and each variation is discussed in detail in the following subsections.

B.1.1 Nozzle Radius Study

Initially, a nozzle size sensitivity study was conducted to understand how the radius of the cooling nozzles influenced the maximum vessel temperature.

A parametric geometry generation script was developed in SpaceClaim (see Listing B.1) to automatically create five geometries with nozzle radii ranging from 10 mm to 18 mm in 2 mm increments, while keeping the total number of nozzles fixed at 138 per duct. A summary of all metrics and simulation results for this study is reported in Table B.1.




Figure B.1: Maximum vessel temperature vs. nozzle radii with fitted quadratic regression.

To summarize, the resulting maximum vessel temperatures for each configuration were fitted with a quadratic regression, yielding a coefficient of determination $R^2 = 0.961$ (Figure B.1). This indicated a strong correlation across the sampled range. However, the total variation between the best and worst-performing cases remained low, at approximately 0.8% in absolute temperature (K). While the data suggested a weak sensitivity with quite a good quadratic fit within the tested radius interval, this trend is not expected to continue indefinitely. As nozzle radii increase, the exit velocity of the cooling air drops, eventually leading to reduced convective cooling effectiveness on the dump surface.

B.1.2 Nozzle Density Study

As the previous nozzle radius study showed only limited influence of nozzle size on thermal performance, the next objective was to understand whether modifying nozzle density and distribution could alter the axial flow distribution along the ducts. The aim was to redirect airflow toward the central portion of the dump, which had experienced suboptimal cooling in previous simulations (Figure 2.14).

The bottom ducts were first conceptually divided into three distinct sections, upstream, middle and downstream, corresponding respectively to the regions before the first support crane, between the two support cranes and after the second crane (Figure B.2).




Figure B.2: Top view of the duct assembly showing sectioning lines (in red) used to define regions for localized nozzle density augmentation.

This segmentation allowed targeted control of nozzle placement within each zone with the script in Listing B.1. To isolate the effect of nozzle distribution, the total number of nozzles was kept constant across all configurations. The nozzle density was then increased in the central section by 25% and 50%, with a proportional reduction in the upstream and downstream sections.

However, the adjustment failed to alter the flow enough to influence the behavior of the maximum observable vessel temperature. The trend observed in previous studies (Figure 2.14) persisted, with higher air velocities still emerging from the downstream section. This confirmed that the flow imbalance was primarily governed by pressure and duct resistance effects rather than nozzle distribution alone.

As a result, the impact on maximum vessel temperatures remained minimal as the observed percentage differences across configurations were within 1%, as summarized in Figure B.3 (comparison metrics observable in Table B.2).




Figure B.3: Percentage difference in maximum vessel temperature between uniform baseline (T_1) and 50% increase in central nozzle density (T_2).

B.1.3 Inclined Nozzles

Following the ineffectiveness of the nozzle redistribution strategy in enhancing central cooling, another approach was explored by introducing inclined nozzles. The objective was to evaluate whether angling the nozzles could reduce the flow resistance caused by the sharp 90° turn from the duct, thereby improving air distribution and ultimately enhancing cooling.

To this end, the existing parametric script was adapted to produce geometries with progressively inclined nozzles, shown in Figure B.4, while keeping the remaining setup unchanged. Nonetheless, the resulting temperature fields revealed no improvement in thermal performance compared to the baseline with the operating cooling duct configuration. The maximum vessel temperatures remained virtually unchanged compared to the baseline case with straight nozzles, indicating that the inclination had no practical effect on the dump's cooling efficiency (results observable in Table B.3). One positive outcome, however, was that the inclined configuration resulted in a 13.8% reduction in the total pressure drop across the system. Despite this, the inclination still failed to keep the flow attached inside the nozzle, and thus did not translate into any significant thermal performance increase.




Figure B.4: Progressively inclined nozzle configuration generated using the adapted parametric script.

This conclusion was further supported by the velocity magnitude contours shown in Figure B.5. As highlighted in the zoomed details, the same axial flow imbalance observed in previous tests is still clearly present. Airflow remains predominantly concentrated toward the downstream nozzles, while the upstream and central regions receive significantly less flow. This persistent distribution pattern demonstrates that the flow dynamics within the duct are largely unaffected by nozzle inclination alone.




Figure B.5: Velocity contours showing persistent downstream flow concentration despite nozzle inclination.

B.1.4 Obstructed Ducts

Given the limited thermal improvement observed with inclined nozzles, a different concept was explored. The goal was to determine whether a local reduction in duct cross-sectional area could generate a favorable local pressure gradient and enhance air ejection at selected positions. To this end, a new geometry was created where the duct cross-section was constricted by 50% at two separate locations, as illustrated in Figure B.6.




Figure B.6: Geometry of the obstructed duct configuration, showing 50% area reductions at two locations aligned with the cradle positions.

The upstream and downstream obstructions were designed to coincide with the dump block supports locations, in order to evaluate how the flow was redirected and how much air was ejected through the corresponding nozzle sections (observable in Figure B.2). The impact on flow distribution is shown in Figure B.7, for the upstream obstruction, and in Figure B.8, for the downstream location. The reduction in duct diameter had a moderate general and highly localized effect. The upstream obstruction successfully generated a pressure build-up that increases local velocity magnitude at the nozzle exit. However, the benefit was partially offset by the sharp 90° turn the air still had to perform from the duct into the nozzle. This abrupt redirection remained a limiting factor, dissipating part of the pressure gain and reducing the overall efficiency of flow redistribution.




Figure B.7: Velocity magnitude contours at the upstream obstruction, illustrating localized pressure build-up and partial acceleration at nozzle exits.




Figure B.8: Velocity magnitude contours at the downstream obstruction.

Moreover, the introduction of these geometric restrictions resulted in a 10.2% increase in total pressure drop compared to the baseline configuration. Such a rise could exceed the operating specifications of the current AHU. While the concept demonstrated that local modifications can influence flow behavior, the corresponding gains in cooling effectiveness were not substantial enough to justify the added pressure losses and increased geometric complexity. A summary of the main results for this configuration is reported in Table B.4.

	Nozzle Radii [mm]														
	10			12			14			16			18		
	Dump event														
	1st	2nd	3rd	1st	2nd	3rd	1st	2nd	3rd	1st	2nd	3rd	1st	2nd	3rd
Max vessel T after 2.5h [°C]	393.4	433.9	458.7	396.1	437.1	462.6	395.5	435.0	464.7	394.4	437.4	467.8	392.7	438.5	469.3
Max core T after 2.5h [°C]	460.4	546.7	600.9	461.9	550.2	606.0	461.6	549.4	605.0	461.7	549.7	605.6	461.8	549.8	605.6
Max vessel T [°C]	505.8	565.8	592.3	506.0	567.0	594.9	505.4	566.9	594.3	505.1	565.9	592.7	504.4	563.8	590.1
Max core T [°C]	2024.4	2092.4	2136.8	2024.4	2092.4	2138.1	2024.4	2093.7	2139.8	2024.4	2094.5	2141.5	2024.4	2094.7	2141.9
Max average vessel T [°C]	472.0	528.3	552.5	472.0	528.9	554.7	471.4	529.4	554.8	471.0	528.2	553.4	470.3	526.4	550.7
Max average Core T [°C]	523.2	633.3	703.1	523.2	633.1	704.1	523.3	634.4	705.5	523.3	634.5	705.9	523.3	634.6	706.0
Max Vessel Total Internal Energy [MJ]	57.9	86.1	100.3	58.4	87.1	102.2	58.6	87.9	103.0	58.7	88.1	103.4	58.7	87.8	103.0
Max Core Total Internal Energy [MJ]	507.0	751.0	893.9	507.0	752.6	901.6	507.0	757.7	907.4	507.0	759.0	910.4	507.0	758.3	909.1
Max Total Internal Energy [MJ]	547.8	840.3	1007.4	547.8	842.4	1016.8	547.8	848.3	1023.6	547.8	849.8	1027.0	547.8	849.0	1025.3
Total Internal Energy dissipated in 2.5h [MJ]	267.2	392.2	458.5	259.7	377.1	443.7	259.2	384.1	451.7	257.7	382.2	450.0	258.5	383.1	450.6
Pressure Drop [Pa]	447.4			284.1			218.6			185.5			173.2		
% raise in max vessel T btw each dump	11.9	4.7		12.1	4.9		12.2	4.8		12.0	4.7		11.8	4.7	
Pearson Corr. btw Max Vessel T and Radius	-0.55														
Spearman Corr. btw Max Vessel T and Radius	-0.4														
% improvement wrt to worst performing geometry evaluated over max vessel T	0.44%				-		0.093%		0.37%		0.8%				

Table B.1: Summary of thermal performance across nozzle radii and dump events.

	1st Dump	2nd Dump	3rd Dump
Max Vessel T			
Max Core T			
Max Vessel T after 2.5h	-0.53%	-0.61%	-0.73%
Max Core T after 2.5h			
Max Average Vessel T			
Max Average Core T			
Max vessel IE			
Max Core IE			
Max Tot IE			
Tot E dissipated in 2.5h			

Table B.2: Percentage differences (calculated based on temperature values in Kelvin) between baseline and **augmented nozzle density configuration** for key thermal metrics over three consecutive dump events. Only highest deviations are reported for readability.

	1st Dump	2nd Dump	3rd Dump
Max Vessel T			
Max Core T			
Max Vessel T after 2.5h			
Max Core T after 2.5h			
Max Average Vessel T			
Max Average Core T			
Max vessel IE		1.51%	1.85%
Max Core IE			
Max Tot IE			
Tot E dissipated in 2.5h	-1.55%		

Table B.3: Percentage differences (calculated based on temperature values in Kelvin) between baseline and **inclined nozzle configuration** for key thermal metrics over three consecutive dump events. Only highest deviations are reported for readability.

	1st Dump	2nd Dump	3rd Dump
Max Vessel T			
Max Core T			
Max Vessel T after 2.5h			
Max Core T after 2.5h			
Max Average Vessel T			
Max Average Core T			
Max vessel IE		1.45%	1.67%
Max Core IE			
Max Tot IE			
Tot E dissipated in 2.5h	-1.39%		

Table B.4: Percentage differences (calculated based on temperature values in Kelvin) between baseline and **obstructed ducts configurations** for key thermal metrics over three consecutive dump events. Only highest deviations are reported for readability.

	1st Dump	2nd Dump	3rd Dump
Max Vessel T	-0.13%	-2.02%	-3.03%
Max Core T			
Max Vessel T after 2.5h			
Max Core T after 2.5h			
Max Average Vessel T			
Max Average Core T			
Max vessel IE			
Max Core IE			
Max Tot IE	0.04%	0.66%	1.76%
Tot E dissipated in 2.5h			

Table B.5: Key temperature and energy metrics for the **2x flowrate central bottom cooling configuration**, reported as percentage differences with respect to the 2x shrinkfitted baseline. The table highlights how the configuration becomes progressively more effective across consecutive dumps. Only the most significant values are shown for clarity.

Listing B.1: Automation script for parametric nozzle generation.

```

#Python Script, API Version = V22

#
#####
LIMITATIONS
#####

# outer_rad_nozzle<= 18.5

#Starting points for sketching
nozzles
#####

tot_nr_nozzles = 69

# number of nozzles created on
single plane (effectively total
number of nozzles = 2*
tot_nr_nozzles)
outer_rad_nozzle = 18

#outer radius nozzles
length_nozzle = 50

# extrusion lenght nozzles

inner_rad_nozzle = outer_rad_nozzle
- 1

inner radius nozzles
rad_boi_impingement=
outer_rad_nozzle*2.73
l_boi_impingement=226
first_nozzle_dist_to_inlet =
outer_rad_nozzle + 150

#
first nozzle distance to the inlet
(change onl figure)
tot_length = 8616

# total lenght where the
nozzles can be create (do not
change)
starting_z = -4257

# z of inlet_r
in global CS (do not change)
l_max = 170+tot_length + starting_z
- outer_rad_nozzle -
first_nozzle_dist_to_inlet #
length over with it will perform
the creation of nozzles (not to
change unless need unconstant

distribution)

# New parameters for guide start
distance and end distance from
inlet (do not change)
start_guide_us = 2359
end_guide_us = 2500
start_guide_ds = 6332
end_guide_ds = 6473

# New parameters for relative
density in sections (sections
defined circle_impingement(plane)
function)
density_section_1 = 1.0
density_section_2 = 1.5
density_section_3 = 1.0

#####
# Plane 1 settings (do not change)
plane1_x = -251.204075541207
plane1_y = -456.787849877095
plane1_z = 4184.79400333488
plane1_dir_x = 0.4999999999999922
plane1_dir_y = -0.866025403784484
plane1_dir_z = -2.5826494325374E-13
# Plane 2 settings (do not change)
plane2_x = -298.053885913417
plane2_y = -410.296292794593
plane2_z = 4184.79400332895
plane2_dir_x = 0.869989468365886
plane2_dir_y = -0.493070304249245
plane2_dir_z = -8.76678549632992E
-14

number_of_nozzles_plane1=0 #
counter of holes done on plane #1
to calculate index of faces/edges
on plane #2

# Function to set the sketch plane
def set_sketch_plane(x, y, z, dir_x
, dir_y, dir_z):
    sectionPlane = Plane.Create(
Frame.Create(Point.Create(MM(x), MM
(y), MM(z)),
Direction.Create(dir_x,
dir_y, dir_z),
Direction.DirZ))
    result = ViewHelper.
SetSketchPlane(sectionPlane, None)

# Function to create concentric
circles at every pitch on a given
plane
def create_concentric_circles(
start_z, l_max, outer_rad,
inner_rad, plane_x, plane_y,
plane_z, plane_dir_x, plane_dir_y,
plane_dir_z):
    global number_of_nozzles_plane1
    global number_of_nozzles_plane2

    #section definition for
relative density parameter

    end_section_1 = start_z +
start_guide_us
    start_section_2 = start_z +
end_guide_us
    end_section_2 = start_z +
start_guide_ds
    start_section_3 = start_z +
end_guide_ds

    total_length_1 = end_section_1
- start_z
    total_length_2 = end_section_2
- start_section_2
    total_length_3 = l_max -
start_section_3

    effective_density_total = (
density_section_1 * total_length_1)
+ (density_section_2 *
total_length_2) + (
density_section_3 * total_length_3)

    # Calcolare il numero di
nozzles per sezione based on the
relative density
    nozzles_1_float = (
density_section_1 * total_length_1
/ effective_density_total) *
tot_nr_nozzles
    nozzles_2_float = (
density_section_2 * total_length_2
/ effective_density_total) *
tot_nr_nozzles
    nozzles_3_float = (
density_section_3 * total_length_3
/ effective_density_total) *
tot_nr_nozzles

    # Arrotondare al numero intero
pi vicino
    nozzles_1 = round(
nozzles_1_float)
    nozzles_2 = round(
nozzles_2_float)
    nozzles_3 = round(
nozzles_3_float)

    # Calcolare la differenza e
aggiustare l'ultima sezione
    total_nozzles = nozzles_1 +
nozzles_2 + nozzles_3
    if total_nozzles !=
tot_nr_nozzles:
        difference = tot_nr_nozzles
- total_nozzles
        nozzles_3 += difference

    # Calcolare i pitch in base al
numero di nozzles arrotondati
    pitch_1 = total_length_1 /
nozzles_1
    pitch_2 = total_length_2 /
nozzles_2
    pitch_3 = total_length_3 /

```

```

nozzles_3

    current_z = start_z +
first_nozzle_dist_to_inlet

    while current_z <= l_max:
        if current_z <=
end_section_1:
            pitch = pitch_1
        elif current_z <=
end_section_2:
            pitch = pitch_2
        else:
            pitch = pitch_3

    # Set Sketch Plane
    set_sketch_plane(plane_x,
plane_y, plane_z, plane_dir_x,
plane_dir_y, plane_dir_z)

    #dont sketch in
correspondence of the guides
    if (start_z +
start_guide_us - outer_rad_nozzle
<= current_z <= start_z +
end_guide_us + outer_rad_nozzle):
        current_z = start_z +
end_guide_us + outer_rad_nozzle + 4
        continue
    if (start_z +
start_guide_ds - outer_rad_nozzle
<= current_z <= start_z +
end_guide_ds + outer_rad_nozzle):
        current_z = start_z +
end_guide_ds + outer_rad_nozzle + 4
        continue

    # Sketch Outer Circle
    origin = Point2D.Create(MM
(0), MM(current_z))
    result = SketchCircle.
Create(origin, MM(outer_rad))

    # Sketch Inner Circle
    origin = Point2D.Create(MM
(0), MM(current_z))
    result = SketchCircle.
Create(origin, MM(inner_rad))

    # Add number of nozzles to
global variable
    if plane_x == plane1_x:

number_of_nozzles_plane1 += 1
    current_z += pitch

# Function to create the sketch/
extrusion of BOI impingement (
called in operate_on_sketch)
def circles_impingement(plane):
    # Change Object Visibility
    selection = BodySelection.
Create(GetRootPart().Bodies[4])
    visibility = VisibilityType.
Hide
    inSelectedView = False
    faceLevel = False
    ViewHelper.SetObjectVisibility(
selection, visibility,
inSelectedView, faceLevel)
    # EndBlock

    # Change Object Visibility
    selection = BodySelection.
Create(GetRootPart().Bodies[21])
    visibility = VisibilityType.
Hide
    inSelectedView = False
    faceLevel = False
    ViewHelper.SetObjectVisibility(
selection, visibility,
inSelectedView, faceLevel)
    # EndBlock

    end_section_1 = starting_z +
start_guide_us
    start_section_2 = starting_z +
end_guide_us
    end_section_2 = starting_z +
start_guide_ds
    start_section_3 = starting_z +
end_guide_ds

    total_length_1 = end_section_1
- starting_z
    total_length_2 = end_section_2
- start_section_2
    total_length_3 = l_max -

```

```

start_section_3

    effective_density_total = (
density_section_1 * total_length_1)
+ (density_section_2 *
total_length_2) + (
density_section_3 * total_length_3)

    # Calcolare il numero di
nozzles per sezione
    nozzles_1_float = (
density_section_1 * total_length_1
/ effective_density_total) *
tot_nr_nozzles
    nozzles_2_float = (
density_section_2 * total_length_2
/ effective_density_total) *
tot_nr_nozzles
    nozzles_3_float = (
density_section_3 * total_length_3
/ effective_density_total) *
tot_nr_nozzles

    # Arrotondare al numero intero
pi vicino
    nozzles_1 = round(
nozzles_1_float)
    nozzles_2 = round(
nozzles_2_float)
    nozzles_3 = round(
nozzles_3_float)

    # Calcolare la differenza e
aggiustare l'ultima sezione
    total_nozzles = nozzles_1 +
nozzles_2 + nozzles_3
    if total_nozzles !=
tot_nr_nozzles:
        difference = tot_nr_nozzles
- total_nozzles
        nozzles_3 += difference

    # Calcolare i pitch in base al
numero di nozzles arrotondati
    pitch_1 = total_length_1 /
nozzles_1
    pitch_2 = total_length_2 /
nozzles_2
    pitch_3 = total_length_3 /
nozzles_3

    current_z = starting_z +
first_nozzle_dist_to_inlet

    while current_z <= l_max:
        if current_z <=
end_section_1:
            pitch = pitch_1
        elif current_z <=
end_section_2:
            pitch = pitch_2
        else:
            pitch = pitch_3

    # Set Sketch Plane
    sectionPlane = Selection.
Create(GetRootPart().DatumPlanes
[6])
    result = ViewHelper.
SetSketchPlane(sectionPlane, None)
    # EndBlock

    if (starting_z +
start_guide_us - outer_rad_nozzle
<= current_z <= starting_z +
end_guide_us + outer_rad_nozzle):
        current_z = starting_z
+ end_guide_us + outer_rad_nozzle +
4
        continue
    if (starting_z +
start_guide_ds - outer_rad_nozzle
<= current_z <= starting_z +
end_guide_ds + outer_rad_nozzle):
        current_z = starting_z
+ end_guide_ds + outer_rad_nozzle +
4
        continue

    # Sketch Outer Circle
    origin = Point2D.Create(MM
(0), MM(current_z))
    result = SketchCircle.
Create(origin, MM(
rad_boi_impingement))
    current_z += pitch

    # Solidify Sketch

```

```

    mode = InteractionMode.Solid
    result = ViewHelper.SetViewMode
(mode, None)
    # EndBlock

    if plane == 1:
        for i in range(0,
number_of_nozzles_plane1, 1):
            # Extrude 1 Face
            selection =
FaceSelection.Create(GetRootPart().
Bodies[22].Faces[0])
            options =
ExtrudeFaceOptions()
            options.ExtrudeType =
ExtrudeType.Cut
            result = ExtrudeFaces.
Execute(selection, MM(-
l_boi_impingement), options)
            # EndBlock

            # Move Solid to BOI
component
            selections =
BodySelection.Create(GetRootPart().
Bodies[21 +
number_of_nozzles_plane1 - i])
            component =
PartSelection.Create(GetRootPart().
Components[3].Content)
            result =
ComponentHelper.
MoveBodiesToComponent(selections,
component, False, None)
            # EndBlock

            # Hide Boi
            selection =
BodySelection.Create([GetRootPart().
.Components[3].Content.Bodies[i]])
            visibility =
VisibilityType.Hide
            inSelectedView = False
            faceLevel = False
            ViewHelper.
SetObjectVisibility(selection,
visibility, inSelectedView,
faceLevel)

        else:
            for i in range(0,
number_of_nozzles_plane1, 1):
                # Extrude 1 Face
                selection =
FaceSelection.Create(GetRootPart().
Bodies[22].Faces[0])
                options =
ExtrudeFaceOptions()
                options.ExtrudeType =
ExtrudeType.Cut
                result = ExtrudeFaces.
Execute(selection, MM(-
l_boi_impingement), options)
                # EndBlock

                # Move Solid to BOI
component
                selections =
BodySelection.Create(GetRootPart().
Bodies[21 +
number_of_nozzles_plane1 - i])
                component =
PartSelection.Create(GetRootPart().
Components[3].Content)
                result =
ComponentHelper.
MoveBodiesToComponent(selections,
component, False, None)
                # EndBlock

                # Hide Boi
                selection =
BodySelection.Create([GetRootPart().
.Components[3].Content.Bodies[i +
number_of_nozzles_plane1]])
                visibility =
VisibilityType.Hide
                inSelectedView = False
                faceLevel = False
                ViewHelper.
SetObjectVisibility(selection,
visibility, inSelectedView,
faceLevel)

            # Change Object Visibility
            selection = BodySelection.
Create(GetRootPart().Bodies[4])
            visibility = VisibilityType.

```

```

Show
    inSelectedView = False
    faceLevel = False
    ViewHelper.SetObjectVisibility(
        selection, visibility,
        inSelectedView, faceLevel)
    # EndBlock

    # Change Object Visibility
    selection = BodySelection.
Create(GetRootPart().Bodies[21])
    visibility = VisibilityType.
Show
    inSelectedView = False
    faceLevel = False
    ViewHelper.SetObjectVisibility(
        selection, visibility,
        inSelectedView, faceLevel)

# Function to operate on sketch and
create nozzles
def operate_on_sketch(plane):
    global length_nozzle
    global outer_rad_nozzle
    global inner_rad_nozzle

    if plane == 1:
        datumPlanePoint = Point.
Create(MM(plane1_x), MM(plane1_y),
MM(plane1_z)) # punto sullo
sketchplane
        datumPlaneDirection =
Direction.Create(plane1_dir_x,
plane1_dir_y, plane1_dir_z) # Usa
la stessa direzione dello sketch
plane per poi ruotare in seguito
    else:
        datumPlanePoint = Point.
Create(MM(plane2_x), MM(plane2_y),
MM(plane2_z)) # punto sullo
sketchplane
        datumPlaneDirection =
Direction.Create(plane2_dir_x,
plane2_dir_y, plane2_dir_z) # Usa
la stessa direzione dello sketch
plane per poi ruotarlo in seguito

        datumPlane = DatumPlaneCreator.
Create(datumPlanePoint,
datumPlaneDirection)

    # Rotate About Z axis
    selection = Selection.Create(
GetRootPart().DatumPlanes[6])
    if plane == 1:
        anchor = Selection.Create(
GetRootPart().Bodies[4].Edges[13].
GetChildren[ICurvePoint]() [1])
    else:
        anchor = Selection.Create(
GetRootPart().Bodies[4].Edges[13].
GetChildren[ICurvePoint]() [0])
        axis = Move.GetAxis(anchor)
        options = MoveOptions()
        result = Move.Rotate(selection,
axis, DEG(90.0000000000002),
options)
    # EndBlock

    # Translate Along X and Z
Handle (X movement is for
creation of impingements' boi)
    if plane == 1:
        selection = Selection.
Create(GetRootPart().DatumPlanes
[6])
        direction = Move.
GetDirection(selection)
        options = MoveOptions()
        result = Move.Translate(
selection, direction, MM(-18.5922-
length_nozzle), options)
        # Translate Along X Handle
direction = Move.
GetDirection(selection, HandleAxis.
X)
        result = Move.Translate(
selection, direction, MM(18.5922),
options)
    else:
        selection = Selection.
Create(GetRootPart().DatumPlanes
[6])

```

```

        direction = Move.
GetDirection(selection)
        options = MoveOptions()
        result = Move.Translate(
selection, direction, MM(18.5211-
length_nozzle), options)
        # Translate Along X Handle
direction = Move.
GetDirection(selection, HandleAxis.
X)
        result = Move.Translate(
selection, direction, MM(-18.5211),
options)
    # EndBlock

    #delete center surface
    if plane == 1:
        start_index = 10
        step = 1
        for i in range(start_index,
start_index +
number_of_nozzles_plane1, step):
            selection =
FaceSelection.Create(GetRootPart().
Bodies[4].Faces[i])
            result = Delete.Execute
(selection)
        else:
            base_index = 10 + int(2 *
number_of_nozzles_plane1)
            step = 1
            for i in range(base_index,
base_index +
number_of_nozzles_plane1, step):
                selection =
FaceSelection.Create(GetRootPart().
Bodies[4].Faces[i])
                result = Delete.Execute
(selection)

    #project to solid
    if plane == 1:
        start_index = 25
        step = 2
        for i in range(start_index,
start_index + 2 *
number_of_nozzles_plane1, step):
            selection =
EdgeSelection.Create(GetRootPart().
Bodies[4].Edges[i])
            target = Selection.
Empty()
            direction = Selection.
Create(GetRootPart().DatumPlanes
[6])
            result = ProjectToSolid
.Execute(selection, target,
direction)
        else:
            base_index = 25 + int(4 *
number_of_nozzles_plane1)
            step = 2
            end_index = int(base_index
+ 2 * number_of_nozzles_plane1)
            for i in range(base_index,
end_index, step):
                selection =
EdgeSelection.Create(GetRootPart().
Bodies[4].Edges[i])
                target = Selection.
Empty()
                direction = Selection.
Create(GetRootPart().DatumPlanes
[6])
                result = ProjectToSolid
.Execute(selection, target,
direction)

    # extrude ring surface to
create boi nozzles
    if plane == 1:
        for i in range(0,
number_of_nozzles_plane1, 1):
            selection =
FaceSelection.Create(GetRootPart().
Bodies[4].Faces[9])
            options =
ExtrudeFaceOptions()
            options.ExtrudeType =
ExtrudeType.Cut
            result = ExtrudeFaces.
Execute(selection, MM(-

```

```

length_nozzle), options)

    # Move Solid to BOI
component
    selections =
BodySelection.Create(GetRootPart().
Bodies[22])
    component =
PartSelection.Create(GetRootPart().
Components[1].Content)
    result =
ComponentHelper.
MoveBodiesToComponent(selections,
component, False, None)
    # EndBlock

    #Hide Boi
    selection =
BodySelection.Create([GetRootPart().
Components[1].Content.Bodies[i]])
    visibility =
VisibilityType.Hide
    inSelectedView = False
    faceLevel = False
    ViewHelper.
SetObjectVisibility(selection,
visibility, inSelectedView,
faceLevel)

    else:
        for i in range(0,
number_of_nozzles_plane1, 1):
            selection =
FaceSelection.Create(GetRootPart().
Bodies[4].Faces[9+2*
number_of_nozzles_plane1])
            options =
ExtrudeFaceOptions()
            options.ExtrudeType =
ExtrudeType.Cut
            result = ExtrudeFaces.
Execute(selection, MM(-
length_nozzle), options)

    # Move Solid to BOI
component
    selections =
BodySelection.Create(GetRootPart().
Bodies[22])
    component =
PartSelection.Create(GetRootPart().
Components[1].Content)
    result =
ComponentHelper.
MoveBodiesToComponent(selections,
component, False, None)
    # EndBlock

    #Hide Boi
    selection =
BodySelection.Create([GetRootPart().
Components[1].Content.Bodies[i+
number_of_nozzles_plane1]])
    visibility =
VisibilityType.Hide
    inSelectedView = False
    faceLevel = False
    ViewHelper.
SetObjectVisibility(selection,
visibility, inSelectedView,
faceLevel)

    # Delete bottom center surface
    if plane == 1:
        for i in range(1, 1 +
number_of_nozzles_plane1, 1):
            selection =
FaceSelection.Create(GetRootPart().
Bodies[21].Faces[1])
            result = Delete.Execute
(selection)
        else:
            for i in range(1, 1 +
number_of_nozzles_plane1, 1):
                selection =
FaceSelection.Create(GetRootPart().
Bodies[21].Faces[1+
number_of_nozzles_plane1])
                result = Delete.Execute
(selection)

    # Extrude interior Circles
Edges
    if plane == 1:
        edge_start_index = 2

```

```

        edge_step = 1
        options =
    ExtrudeEdgeOptions()
        upToSelection = Selection.
    Create(GetRootPart().DatumPlanes
    [6])
        for i in range(
            edge_start_index, edge_start_index
            +number_of_nozzles_plane1,
            edge_step):
            selection =
    EdgeSelection.Create(GetRootPart().
    Bodies[21].Edges[i])
            #secondarySelection =
    FaceSelection.Create(GetRootPart().
    Bodies[4].Faces[8]) ##### da
    modificare in funz al numer di
    buchi
            result = ExtrudeEdges.
    Execute(selection, Direction.Create
    (-0.866025403784454,
    -0.4999999999999974,
    -1.35704729573522E-13),
    upToSelection, Point.Create(MM
    (-249.056671958879), MM
    (-34.8472619845258), MM
    (-1316.40960431812)), ExtrudeType.
    None)
        else:
            edge_base_index = 2 + 2 *
            number_of_nozzles_plane1
            edge_step = 1
            options =
    ExtrudeEdgeOptions()
            upToSelection = Selection.
    Create(GetRootPart().DatumPlanes
    [6])
            for i in range(
                edge_base_index, edge_base_index +
                number_of_nozzles_plane1, edge_step
            ):
                selection =
    EdgeSelection.Create(GetRootPart().
    Bodies[21].Edges[i])
                #secondarySelection =
    FaceSelection.Create(GetRootPart().
    Bodies[4].Faces[8 + number_of
    nozzles_plane1])
                result = ExtrudeEdges.
    Execute(selection, Direction.Create
    (-0.493070304249179,
    -0.869989468365923,
    -2.29890382193861E-13),
    upToSelection, Point.Create(MM
    (-779.083821556032), MM
    (236.127137799035), MM
    (-850.981600189991)), ExtrudeType.
    None)

# Extrude external Circles Edges
    if plane == 1:
        edge_start_index = 24
        edge_step = 1
        options =
    ExtrudeEdgeOptions()
        upToSelection = Selection.
    Create(GetRootPart().DatumPlanes
    [6])
        for i in range(
            edge_start_index, edge_start_index
            +number_of_nozzles_plane1,
            edge_step):
            selection =
    EdgeSelection.Create(GetRootPart().
    Bodies[4].Edges[24])
            #secondarySelection =
    FaceSelection.Create(GetRootPart().
    Bodies[4].Faces[8]) ##### da
    modificare in funz al numer di
    buchi
            result = ExtrudeEdges.
    Execute(selection, Direction.Create
    (-0.866025403784454,
    -0.4999999999999974,
    -1.35704729573522E-13),
    upToSelection, Point.Create(MM
    (135.723145622414), MM
    (-1026.96546430816), MM
    (102.086540670581)), ExtrudeType.
    None)
        else:
            edge_base_index = 24 + 4 *
            number_of_nozzles_plane1
            edge_step = 1
            options =

```

```

    ExtrudeEdgeOptions()
        upToSelection = Selection.
    Create(GetRootPart().DatumPlanes
    [6])
        for i in range(
            edge_base_index, edge_base_index +
            number_of_nozzles_plane1, edge_step
        ):
            selection =
    EdgeSelection.Create(GetRootPart().
    Bodies[4].Edges[24+4*
    number_of_nozzles_plane1])
            #secondarySelection =
    FaceSelection.Create(GetRootPart().
    Bodies[4].Faces[8 + number_of
    nozzles_plane1])
            result = ExtrudeEdges.
    Execute(selection, Direction.Create
    (-0.493070304249179,
    -0.869989468365923,
    -2.29890382193861E-13),
    upToSelection, Point.Create(MM
    (-779.083821556032), MM
    (236.127137799035), MM
    (-850.981600189991)), ExtrudeType.
    None)

        # Extrude Edges to close top of
        nozzles and creates boi
        impingements
        if plane ==1:
            options =
    ExtrudeEdgeOptions()
            edge_start_index = 26
            face_start_index = 9
            step = 3
            for i in range(
                number_of_nozzles_plane1):
                selection =
    EdgeSelection.Create(GetRootPart().
    Bodies[4].Edges[edge_start_index +
    i * step])
                secondarySelection =
    FaceSelection.Create(GetRootPart().
    Bodies[4].Faces[face_start_index +
    i])
                results = ExtrudeEdges.
    Execute(selection,
    secondarySelection, MM(
    outer_rad_nozzle-inner_rad_nozzle),
    options)
                circles_impingement(1)
        else:
            options =
    ExtrudeEdgeOptions()
            edge_start_index = 26+4*
            number_of_nozzles_plane1
            face_start_index = 9+2*
            number_of_nozzles_plane1
            step = 3
            for i in range(
                number_of_nozzles_plane1):
                selection =
    EdgeSelection.Create(GetRootPart().
    Bodies[4].Edges[edge_start_index +
    i * step])
                secondarySelection =
    FaceSelection.Create(GetRootPart().
    Bodies[4].Faces[face_start_index +
    i])
                result = ExtrudeEdges.
    Execute(selection,
    secondarySelection, MM(
    outer_rad_nozzle-inner_rad_nozzle),
    options)
                circles_impingement(2)
        # EndBlock

        # Elimina il datum plane, il
        punto e la direzione
        selection = Selection.Create(
    GetRootPart().DatumPlanes[6])
        result = Delete.Execute(
    selection)
        #EndBlock

    def show_hide_bodies_in_component(
    component_number,
    number_of_bodies_in_component,
    action):
        # Determina la visibilit in
        base all'azione
        if action.lower() == 'show':
            visibility = VisibilityType
            .Show

```

```

        else:
            visibility = VisibilityType
            .Hide

            inSelectedView = False
            faceLevel = False

            # Loop per iterare "
            number_of_bodies_in_component"
            volte
            for i in range(
                number_of_bodies_in_component):
                # Seleziona il corpo
                corrente
                selection = BodySelection.
    Create(GetRootPart().Components[
    component_number].Content.Bodies[i
    ])

            # Imposta la visibilit
            del corpo corrente
            ViewHelper.
    SetObjectVisibility(selection,
            visibility, inSelectedView,
            faceLevel)

    def mirror_bodies_in_component(
    component_number,
    number_of_bodies_in_component,
    mirror_plane_number):
        # Opzioni di mirror comuni
        options = MirrorOptions()
        options.MergeObjects = False
        options.CreateRelationships =
        False

        # Selezione del piano di mirror
        mirrorPlane = Selection.Create(
    GetRootPart().DatumPlanes[
    mirror_plane_number])

        # Loop per iterare "2 *
        number_of_nozzles_plane1" volte
        for i in range(
            number_of_bodies_in_component):
            # Seleziona il corpo
            corrente
            selection = BodySelection.
    Create(GetRootPart().Components[
    component_number].Content.Bodies[i
    ])

            # Esegui il mirror del
            corpo corrente
            result = Mirror.Execute(
            selection, mirrorPlane, options,
            None)

    def move_bodies_between_components(
    source_component_index,
    number_of_bodies_to_move,
    destination_component_index):
        # Opzioni di spostamento comuni
        options = MoveOptions()

        # Seleziona il componente di
        destinazione
        component = PartSelection.
    Create(GetRootPart().Components[
    destination_component_index].
    Content)

        # Loop per iterare sul numero
        di corpi da spostare
        for i in range(
            number_of_bodies_to_move):
            # Seleziona il corpo
            corrente nel componente di origine
            selection = BodySelection.
    Create(GetRootPart().Components[
    source_component_index].Content.
    Bodies[2*number_of_nozzles_plane1])

            # Esegui lo spostamento del
            corpo corrente al componente di
            destinazione
            result = ComponentHelper.
    MoveBodiesToComponent(selection,
            component, False, None)

        # Call the function to create the
        sketch on Plane 1
        create_concentric_circles(
        starting_z, l_max, outer_rad_nozzle
        , inner_rad_nozzle, plane1_x,

```



```

plane1_y, plane1_z, plane1_dir_x,
plane1_dir_y, plane1_dir_z)

# Solidify Sketch
mode = InteractionMode.Solid
result = ViewHelper.SetViewMode(
mode, None)
# EndBlock

#call the function to manipulate
the geometry and create the nozzle
geomtry on plane 1
operate_on_sketch(1)

# Call the function to create the
sketch on Plane 2
create_concentric_circles(
starting_z, l_max, outer_rad_nozzle
, inner_rad_nozzle, plane2_x,
plane2_y, plane2_z, plane2_dir_x,
plane2_dir_y, plane2_dir_z)

# Solidify Sketch
mode = InteractionMode.Solid
result = ViewHelper.SetViewMode(
mode, None)
# EndBlock

#call the function to manipulate
the geometry and create the nozzle
geomtry on plane 2
operate_on_sketch(2)

#fix the geom. unify into single
surface and repair it
# Stitch Surfaces
result = StitchFaces.FindAndFix()
# EndBlock

# Fix Split Edges
result = FixSplitEdges.FindAndFix()
# EndBlock

# Mirroring section to create left
duct+bois
# Create Datum Plane for Miirroring
purposes
selection = Selection.Create(
GetRootPart().CoordinateSystems[0].
Axes[0])

```

```

result = DatumPlaneCreator.Create(
selection, False, None)

# Show bois that were hidden
#show_hide_bodies_in_component(1,
2*number_of_nozzles_plane1,'show')
#show_hide_bodies_in_component(3,
2*number_of_nozzles_plane1,'show')

# Mirror duct
selection = BodySelection.Create(
GetRootPart().Bodies[4])
mirrorPlane = Selection.Create(
GetRootPart().DatumPlanes[6])
options = MirrorOptions()
options.MergeObjects = False
options.CreateRelationships = False
result = Mirror.Execute(selection,
mirrorPlane, options, None)

# Mirroring nozzles' boi
mirror_bodies_in_component(1, 2*
number_of_nozzles_plane1, 6)

# Mirroring impingements' boi
mirror_bodies_in_component(3, 2*
number_of_nozzles_plane1, 6)

# Move mirrored boi into respective
components
move_bodies_between_components(1,2*
number_of_nozzles_plane1,0)
move_bodies_between_components(3,2*
number_of_nozzles_plane1,2)

# Delete mirroring plane
selection = Selection.Create(
GetRootPart().DatumPlanes[6])
result = Delete.Execute(selection)
# EndBlock

# Delete old duct-1
selection = BodySelection.Create(
GetRootPart().Bodies[5])
result = Delete.Execute(selection)
# EndBlock

# Rename 'Duct-r' to 'Duct-1'
selection = BodySelection.Create(
GetRootPart().Bodies[20])
result = RenameObject.Execute(
selection, "Duct-1")

```

```

# EndBlock

#sharing the new duct 1 with the
enclosure
# hide duct-r
selection = BodySelection.Create(
GetRootPart().Bodies[4])
visibility = VisibilityType.Hide
inSelectedView = False
faceLevel = False
ViewHelper.SetObjectVisibility(
selection, visibility,
inSelectedView, faceLevel)
# EndBlock

#hide all bois
#show_hide_bodies_in_component(1,
2*number_of_nozzles_plane1,'hide')
show_hide_bodies_in_component(2, 2*
number_of_nozzles_plane1,'hide')
#show_hide_bodies_in_component(3,
2*number_of_nozzles_plane1,'hide')
show_hide_bodies_in_component(0, 2*
number_of_nozzles_plane1,'hide')

# Show enclosure
selection = BodySelection.Create(
GetRootPart().Bodies[0])
visibility = VisibilityType.Show
inSelectedView = False
faceLevel = False
ViewHelper.SetObjectVisibility(
selection, visibility,
inSelectedView, faceLevel)
# EndBlock


# Show inlet-1
selection = BodySelection.Create(
GetRootPart().Bodies[6])
visibility = VisibilityType.Show
inSelectedView = False
faceLevel = False
ViewHelper.SetObjectVisibility(
selection, visibility,
inSelectedView, faceLevel)
# EndBlock

# Share Topology
options = ShareTopologyOptions()
options.Tolerance = MM(0.2)
result = ShareTopology.FindAndFix(
options)
# EndBlock

```


Appendix C

**ER35C
RH35C**



Level unit using

SUPPLY AIR - sound power level	ETA 40 / 1.5	5000 m³/h / 799 Pa tot.
f [Hz]:	63.5 125 250 500 1000 2000 4000 8000	Σ Hz
LW Fan:	58 67 83 84 86 81 78 67	90 dB(A)
- IL Casing	7 13 17 20 24 24 28 28	
tot.:	51 54 66 64 62 57 50 39	70 dB(A)

Total sound power level (fan section): 70 dB(A)

Total sound pressure level (1.0 m distance): 62 dB(A)


Level connector

Sound power level fresh air connector	(LW (adj.) = -1,1 dB)	
f [Hz]:	63.5 125 250 500 1000 2000 4000 8000	Σ Hz
FA:	58 65 77 85 84 82 76 71	89 dB(A)
- SIL:	2 2 1 1 2 3 4 2	
FA:	56 63 76 84 82 79 72 69	87 dB(A)

Sound power level supply air connector	(LW (adj.) = -1,1 dB)
f [Hz]:	63.5 125 250 500 1000 2000 4000 8000 Σ Hz
SA:	60 68 82 89 88 87 79 71 93 dB(A)

FA=fresh air / SA=supply air / SIL=silencer / FN=flow noise

Tolerance = + 4 dB (without consideration of absorption and reflexion / free field conditions)
LP = LW - LW (adj.) / absorption of all components is considered!!
guarantee only for total level, octavelevel is just for information purpose



view = access side
damper made from aluminium
damper shaft outside
corners made from aluminium
system is built thermal separated (T3/TB3)
frequency converter on side

V_h = 0 - 5000 m³/h

Pos.: 001
Anl.: UD 62 - UAVB - 00622

FA: 58 67 83 84 86 81 78 67
SIL: 2 2 1 1 2 3 4 2
FA: 56 63 76 84 82 79 72 69

SA: 60 68 82 89 88 87 79 71

Project: 0024467 / CERN UD 68
Position: 101/Anlage 1-UAVB-00622

Unit frame:	Aluminium	AlMgSi	05	Insulation:	Mineralfiber
Wall thickness:	40	mm		Surface:	Galvanized

Supply ETA 40/1.5 according to RLT 01 Inside

FV1J
V in m³/h 5000 wL in m/s 2,3 dP ext. 500 Pa

Fan Z-ETA 355 1 Piece HOWATHERM Gesamt 799 Pa
n in U/min 2803 LW in dB 92 N in kW 1,7 ETA fa in % 69,0 dP dyn. in Pa 39
m in kg 36,2 f in Hz 46,7 v in mm/s 2,8 ETA t in % 76,4 Pm in kW 2,2
dP Air 39 Pa

Motor SIEMENS 1LA9 100 L 1 piece
N in kW 3 I in A 5,7 U in V 230/400 ETA in % 87 n in 1/min 2890
Inlet Ait guard

Dampers (1) HOWATHERM
Number 1/2 1 M in Nm 5 Number 1/1 0 M inNm 8
dP Air 41 Pa

panel Filter Kalthoff
typ: VSK - F8 - HT B/H/L 592/592/292 287/592/292
1 piece 1 piece
dP Air 219 Pa

Extras
Unit Connector 1/1 1,0 Potential Wiring 2,0
Mineral Fiber Isol. 1,0

General Data
Therm. Insulation U = 1.11 W/m² K KZG 40 (TB4-TB3) kb = 0.37-0.51
Total weight 414 kg per Unit

General Data
Therm. Insulation U = 1.11 W/m² K KZG 40 (TB4-TB3) kb = 0.37-0.51
Total weight 414 kg per Unit

General Data
Therm. Insulation U = 1.11 W/m² K KZG 40 (TB4-TB3) kb = 0.37-0.51
Total weight 414 kg per Unit

General Data
Therm. Insulation U = 1.11 W/m² K KZG 40 (TB4-TB3) kb = 0.37-0.51
Total weight 414 kg per Unit

General Data
Therm. Insulation U = 1.11 W/m² K KZG 40 (TB4-TB3) kb = 0.37-0.51
Total weight 414 kg per Unit

General Data
Therm. Insulation U = 1.11 W/m² K KZG 40 (TB4-TB3) kb = 0.37-0.51
Total weight 414 kg per Unit

General Data
Therm. Insulation U = 1.11 W/m² K KZG 40 (TB4-TB3) kb = 0.37-0.51
Total weight 414 kg per Unit

General Data
Therm. Insulation U = 1.11 W/m² K KZG 40 (TB4-TB3) kb = 0.37-0.51
Total weight 414 kg per Unit

General Data
Therm. Insulation U = 1.11 W/m² K KZG 40 (TB4-TB3) kb = 0.37-0.51
Total weight 414 kg per Unit

General Data
Therm. Insulation U = 1.11 W/m² K KZG 40 (TB4-TB3) kb = 0.37-0.51
Total weight 414 kg per Unit

General Data
Therm. Insulation U = 1.11 W/m² K KZG 40 (TB4-TB3) kb = 0.37-0.51
Total weight 414 kg per Unit

General Data
Therm. Insulation U = 1.11 W/m² K KZG 40 (TB4-TB3) kb = 0.37-0.51
Total weight 414 kg per Unit

General Data
Therm. Insulation U = 1.11 W/m² K KZG 40 (TB4-TB3) kb = 0.37-0.51
Total weight 414 kg per Unit

General Data
Therm. Insulation U = 1.11 W/m² K KZG 40 (TB4-TB3) kb = 0.37-0.51
Total weight 414 kg per Unit

General Data
Therm. Insulation U = 1.11 W/m² K KZG 40 (TB4-TB3) kb = 0.37-0.51
Total weight 414 kg per Unit

General Data
Therm. Insulation U = 1.11 W/m² K KZG 40 (TB4-TB3) kb = 0.37-0.51
Total weight 414 kg per Unit

General Data
Therm. Insulation U = 1.11 W/m² K KZG 40 (TB4-TB3) kb = 0.37-0.51
Total weight 414 kg per Unit

General Data
Therm. Insulation U = 1.11 W/m² K KZG 40 (TB4-TB3) kb = 0.37-0.51
Total weight 414 kg per Unit

General Data
Therm. Insulation U = 1.11 W/m² K KZG 40 (TB4-TB3) kb = 0.37-0.51
Total weight 414 kg per Unit

General Data
Therm. Insulation U = 1.11 W/m² K KZG 40 (TB4-TB3) kb = 0.37-0.51
Total weight 414 kg per Unit

General Data
Therm. Insulation U = 1.11 W/m² K KZG 40 (TB4-TB3) kb = 0.37-0.51
Total weight 414 kg per Unit

General Data
Therm. Insulation U = 1.11 W/m² K KZG 40 (TB4-TB3) kb = 0.37-0.51
Total weight 414 kg per Unit

General Data
Therm. Insulation U = 1.11 W/m² K KZG 40 (TB4-TB3) kb = 0.37-0.51
Total weight 414 kg per Unit

General Data
Therm. Insulation U = 1.11 W/m² K KZG 40 (TB4-TB3) kb = 0.37-0.51
Total weight 414 kg per Unit

General Data
Therm. Insulation U = 1.11 W/m² K KZG 40 (TB4-TB3) kb = 0.37-0.51
Total weight 414 kg per Unit

General Data
Therm. Insulation U = 1.11 W/m² K KZG 40 (TB4-TB3) kb = 0.37-0.51
Total weight 414 kg per Unit

General Data
Therm. Insulation U = 1.11 W/m² K KZG 40 (TB4-TB3) kb = 0.37-0.51
Total weight 414 kg per Unit

General Data
Therm. Insulation U = 1.11 W/m² K KZG 40 (TB4-TB3) kb = 0.37-0.51
Total weight 414 kg per Unit

General Data
Therm. Insulation U = 1.11 W/m² K KZG 40 (TB4-TB3) kb = 0.37-0.51
Total weight 414 kg per Unit

General Data
Therm. Insulation U = 1.11 W/m² K KZG 40 (TB4-TB3) kb = 0.37-0.51
Total weight 414 kg per Unit

General Data
Therm. Insulation U = 1.11 W/m² K KZG 40 (TB4-TB3) kb = 0.37-0.51
Total weight 414 kg per Unit

General Data
Therm. Insulation U = 1.11 W/m² K KZG 40 (TB4-TB3) kb = 0.37-0.51
Total weight 414 kg per Unit

General Data
Therm. Insulation U = 1.11 W/m² K KZG 40 (TB4-TB3) kb = 0.37-0.51
Total weight 414 kg per Unit


General Data
Therm. Insulation U = 1.11 W/m² K KZG 40 (TB4-TB3) kb = 0.37-0.51
Total weight 414 kg per Unit


General Data
Therm. Insulation U = 1.11 W/m² K KZG 40 (TB4-TB3) kb = 0.37-0.51
Total weight 414 kg per Unit

General Data
Therm. Insulation U = 1.11 W/m² K KZG 40 (TB4-TB3) kb = 0.37-0.51
Total weight 414 kg per Unit

General Data
Therm. Insulation U = 1.11 W/m² K KZG 40 (TB4-TB3) kb = 0.37-0.51
Total weight 414 kg per Unit

General Data
Therm. Insulation U = 1.11 W/m² K KZG 40 (TB4-TB3) kb = 0.37-0.51
Total weight 414 kg per Unit





HOWATHERM Klimatechnik GmbH				Version : 0108705
Keiperweg 11-15	55767 Brücken	Tel.: 06782/9999-0	Fax: 06782/9999-10	
Internet: www.howatherm.de		email: info@howatherm.de		
Date: 16.05.2006		Specialist: Schroeck_H		
AHU ACOUSTIC				Page : A 1
Project: 0024467 / CERN UD 68				
Position: 101/Anlage 1-UAVB-00682				

Level unit using

SUPPLY AIR - sound power level				ETA 40 / 1.5		5000 m³/h / 799 Pa tot.					
f [Hz]:	63.5	125	250	500	1000	2000	4000	8000	Σ		
LW Fan:	58	67	83	84	86	81	78	67	90	dB(A)	
- IL Casing	7	13	17	20	24	24	28	28			
tot.:	51	54	66	64	62	57	50	39	70	dB(A)	
Total sound power level (fan section):										70 dB(A)	
Total sound pressure level (1,0 m distance):										62 dB(A)	

Level connector


Sound power level fresh air connector										(LW (adj.) = -1,1 dB)	
f [Hz]	63.5	125	250	500	1000	2000	4000	8000	Σ	Hz	
FA:	58	65	77	85	84	82	76	71	89	dB(A)	
- SIL:	2	2	1	1	2	3	4	2			
FA:	56	63	76	84	82	79	72	69	87	dB(A)	
Sound power level supply air connector										(LW (adj.) = -1,1 dB)	
f [Hz]	63.5	125	250	500	1000	2000	4000	8000	Σ	Hz	
SA:	60	68	82	89	88	87	79	71	93	dB(A)	

FA=fresh air / SA=supply air / SIL=silencer / FN=flow noise

Tolerance = + 4 dB (without consideration of absorption and reflexion / free field conditions)
LP = LW - LW (adj.) / absorption of all components is considered!!
guarantee only for total level, octavelevel is just for information purpose

HOWATHERM Klimatechnik GmbH				Version : 0108705
Keiperweg 11-15	55767 Brücken	Tel.: 06782/9999-0	Fax: 06782/9999-10	
Internet: www.howatherm.de		email: info@howatherm.de		
Date: 16.05.2006		Specialist: Schroeck_H		
AHU ACOUSTIC				Page : A 2
Project: 0024467 / CERN UD 68				
Position: 101/Anlage 1-UAVB-00682				

Level connector



Declaration of originality

The signed declaration of originality is a component of every written paper or thesis authored during the course of studies. **In consultation with the supervisor**, one of the following two options must be selected:

- ☐ I hereby declare that I authored the work in question independently, i.e. that no one helped me to author it. Suggestions from the supervisor regarding language and content are excepted. I used no generative artificial intelligence technologies¹.
- ☒ I hereby declare that I authored the work in question independently. In doing so I only used the authorised aids, which included suggestions from the supervisor regarding language and content and generative artificial intelligence technologies. The use of the latter and the respective source declarations proceeded in consultation with the supervisor.

Title of paper or thesis:

Design and Validation of an Enhanced Cooling System Model for the High-Luminosity LHC Beam Dump at CERN

Authored by:

If the work was compiled in a group, the names of all authors are required.

Last name(s):

Binda Notarianni

First name(s):

Giovanni Angelo Roberto

With my signature I confirm the following:

- I have adhered to the rules set out in the [Citation Guidelines](#).
- I have documented all methods, data and processes truthfully and fully.
- I have mentioned all persons who were significant facilitators of the work.

I am aware that the work may be screened electronically for originality.

Place, date

Geneva 30/05/2025

Signature(s)



If the work was compiled in a group, the names of all authors are required. Through their signatures they vouch jointly for the entire content of the written work.

¹ For further information please consult the ETH Zurich websites, e.g. <https://ethz.ch/en/the-eth-zurich/education/ai-in-education.html> and <https://library.ethz.ch/en/researching-and-publishing/scientific-writing-at-eth-zurich.html> (subject to change).