



Department of Precision and Microsystems Engineering

Vibroacoustic Optimisation of 3D Printed Loudspeaker Cabinets

H.S. van der Sluis

Report no : 2026.029
Coach : J. Yang
Professor : M. Langelaar
Specialisation : CDM
Type of report : Thesis
Date : 02-04-2026

Vibroacoustic Optimisation of 3D Printed Loudspeaker Cabinets

Optimising Acoustic Performance Through
Hybrid FEM/BEM Simulation of Structural
Cabinet Dynamics and Sound Radiation

by

H.S. van der Sluis

to obtain the degree of Master of Science
at the Delft University of Technology

| | |
|----------------------|--|
| Student number: | 4785533 |
| Instructors: | M. Langelaar & J. Yang |
| External Supervisor: | R. Thomson |
| Project Duration: | March, 2025 - April, 2026 |
| Faculty: | Faculty of Mechanical Engineering |
| Department: | Precision and Microsystems Engineering [PME] |

Contents

| | |
|---|------------|
| Nomenclature | iii |
| 1 Introduction | 1 |
| 1.1 Background and Motivation | 1 |
| 1.2 Challenges in Loudspeaker Cabinet Design | 2 |
| 1.3 Research Gap and Objectives | 2 |
| 1.4 Scope and Contributions | 3 |
| 1.5 Structure of the Thesis | 3 |
| 2 Modelling Approaches: Comparative Exploration | 4 |
| 2.1 Overview of Vibroacoustic Simulation Methods | 5 |
| 2.1.1 FEM-PML | 5 |
| 2.1.2 FEM/BEM | 6 |
| 2.2 Motivation for Hybrid FEM/BEM Approach | 6 |
| 3 Final Modelling Framework | 7 |
| 3.1 Overview of Hybrid FEM/BEM Method | 7 |
| 3.2 Geometry Definition and Parametric Cabinet Representation | 8 |
| 3.3 Driver and Excitation Modelling | 9 |
| 3.4 Material Modelling of 3D-Printed Polymers | 11 |
| 3.5 Structural-Acoustic Coupling and Domain Setup | 12 |
| 3.5.1 Domain Definitions and Boundary Conditions | 12 |
| 3.5.2 Multiphysics Couplings | 12 |
| 3.5.3 Considerations for Vented Enclosures | 13 |
| 3.6 Meshing Strategy and Study Setup | 13 |
| 3.6.1 Frequency-Domain Analysis | 13 |
| 3.6.2 Acoustic Mesh | 13 |
| 3.6.3 Structural Mesh | 14 |
| 3.6.4 Solver Configuration | 14 |
| 3.6.5 FEM Cushion Exploration | 15 |
| 3.7 Verification and Numerical Convergence | 15 |
| 3.7.1 Mesh Convergence | 15 |
| 3.7.2 Quarter-Geometry Validation | 15 |
| 3.7.3 Frequency Sampling Interval | 16 |
| 4 Performance Metrics and Analysis Methods | 18 |
| 4.1 SPL Metrics | 18 |
| 4.2 Structural Energy Metrics & Q-factor Acquisition | 19 |
| 4.3 Directivity Metrics | 20 |
| 4.3.1 Using Directivity in BEM Shape Optimisation | 22 |
| 4.4 Sensitivity Considerations for Optimisation | 22 |
| 4.4.1 Geometric Parameter Sensitivity Study | 23 |
| 4.4.2 Comparison of Sensitivity Coefficients | 25 |
| 4.4.3 Implications for Optimisation Strategy | 25 |
| 5 Experimental Validation of the Simulation Model | 27 |
| 5.1 Description of the Prototype Cabinet | 27 |
| 5.2 Experimental Setup and Measurement Procedure | 28 |
| 5.3 Simulation-to-Measurement Comparison | 29 |
| 5.4 Conclusions on Model Credibility | 32 |

| | | |
|----------|---|-----------|
| 6 | Optimisation Framework and Results | 33 |
| 6.1 | Optimisation Objectives and Constraints | 33 |
| 6.1.1 | Objective Function | 33 |
| 6.1.2 | Manufacturability Constraints | 33 |
| 6.1.3 | Computational Limitations | 34 |
| 6.2 | Cabinet Parameterisation Strategy | 34 |
| 6.2.1 | Parametric Baffle Edge Optimisation | 35 |
| 6.2.2 | Shape Optimisation Panel Parameterisation | 35 |
| 6.2.3 | Mesh Smoothing and Thickness Preservation | 35 |
| 6.3 | Optimisation Algorithm Selection and Implementation | 36 |
| 6.4 | Optimisation Results and Interpretation | 37 |
| 6.4.1 | Preliminary Insights on Relative Improvement Evaluation | 37 |
| 6.4.2 | Baseline Performance and Parametric Optimisation | 37 |
| 6.4.3 | Shape Optimisation: Convergence and Performance | 38 |
| 6.4.4 | Physical Mechanism Analysis | 41 |
| 6.4.5 | Robustness Analysis | 43 |
| 6.4.6 | Framework Sensitivity to Material Printing Direction | 43 |
| 6.4.7 | Alternative Optimisation Approach: Structural Eigenfrequency | 45 |
| 7 | Discussion | 47 |
| 7.1 | Summary of Key Trade-offs Observed | 47 |
| 7.2 | Insights from the Modelling and Optimisation Process | 48 |
| 7.3 | Limitations of the Current Approach | 49 |
| 7.4 | Future Work and Experimental Validation Path | 50 |
| 8 | Conclusions | 51 |
| 8.1 | Summary of Key Findings | 51 |
| 8.2 | Contributions to Vibroacoustic Design | 51 |
| 8.3 | Implications and Outlook | 52 |
| | References | 53 |
| A | Alternative Methods for Modelling Vibroacoustic Exterior Radiation | 55 |
| B | Simulation Results using the FEM-PML Method | 57 |
| B.1 | Key Findings and Observed Limitations | 58 |
| C | Exploratory Simulation Method using an Acoustic Buffer Domain | 60 |
| C.1 | Key Findings and Observed Limitations | 61 |

Nomenclature

Abbreviations

| Abbreviation | Definition |
|--------------|---|
| ASI | Acoustic-Structure Interaction |
| BEM | Boundary Element Modelling |
| CFD | Computational Fluid Dynamics |
| CV | Coefficient of Variation |
| dBFS | Decibels Relative to Full Scale |
| DI | Directivity Index |
| DOF | Degree of Freedom |
| ESE | Elastic Strain Energy |
| FEM | Finite Element Modelling |
| FR | Frequency Response |
| GCMMA | Globally Convergent Method of Moving Asymptotes |
| HiFi | High Fidelity |
| IMU | Inertial Measurement Unit |
| IPOPT | Interior Point Optimiser |
| LFAM | Large Format Additive Manufacturing |
| LMS | Lumped Mechanical System |
| LPM | Lumped Parameter Modelling |
| MAE | Mean Absolute Error |
| MMA | Method of Moving Asymptotes |
| NRA | Narrow Region Acoustics |
| PML | Perfectly Matched Layers |
| SO | Shape Optimisation |
| SPL | Sound Pressure Level |
| STFT | Short-Time Fourier Transform |
| TO | Topology Optimisation |
| TS | Thiele/Small |

Symbols and Expressions

| Symbol | Definition | Unit |
|-----------------------|------------------------|----------------------|
| a | Piston Radius | [m] |
| c | Speed of sound | [m/s] |
| C_{ms} | Mechanical Compliance | [m/N] |
| E | Young's Modulus | [GPa] |
| \mathbf{f} | Body Force Vector | [N/m ³] |
| f | Frequency | [Hz] |
| f_0 | Study Frequency | [Hz] |
| f_r / f_s | Resonant Frequency | [Hz] |
| I | Acoustic Intensity | [Pa ²] |
| k | Acoustic wave number | [m ⁻¹] |
| M_{ms} | Moving Mass | [kg] |
| \mathbf{n} | Outward Normal | [-] |
| p | Acoustic pressure | [Pa] |
| Q | Q-factor | [-] |
| Q_{ms} | Mechanical Q-factor | [-] |
| r_0 | 1 m Reference Point | [m] |
| S | Absolute Sensitivity | [-] |
| S_d | Diaphragm Surface Area | [m ²] |
| t | Time | [s] |
| \mathbf{u} | Displacement Field | [-] |
| v_n | Normal Velocity | [m/s] |
| v_v | Volume Velocity | [m ³ /s] |
| W | Directivity Factor | [-] |
| α | Cone Correction Factor | [-] |
| η_{mat} | Material Loss Factor | [-] |
| λ | Wavelength | [m] |
| λ_n | Eigenfrequency | [Hz] |
| ρ | Material Density | [kg/m ³] |
| ρ_s | Structural Density | [kg/m ³] |
| ρ_0 | Air Density | [kg/m ³] |
| $\boldsymbol{\sigma}$ | Cauchy Stress Tensor | [Pa] |
| ϕ_{ac} | Acoustic Phase | [°] |

Introduction

1.1. Background and Motivation

Loudspeaker design affects most people in their day-to-day lives, often more than they realise. Applications of loudspeakers that may come to mind are in movie theatres, festivals, concerts, amusement parks, but also in more practical applications such as PA systems. The world of audio engineering and loudspeaker design has seen much development in the past decades. Focus has shifted from innovation of complete system-level design, such as speaker-arrays and room acoustics, towards the refinement of individual components, such as cabinets, drivers, and signal processing. This shift has been enabled by developments in material science, fabrication techniques, and computational simulation.

The design philosophy that dominates current loudspeaker cabinet fabrication is based on decades-old research and intuition passed down from generation to generation. There is little incentive to deviate from what has been proven to work. The primary reason as to why loudspeaker cabinets are typically rectangular lies in the fabrication methods. Most cabinets are made through panel construction in varying levels of complexity. This is a cheap and repeatable fabrication method, but it lacks the potential for shape freedom. More sophisticated methods, such as fibre glass bonding and aluminium casting, gain that shape freedom but come with high cost, longer production times, and less design flexibility. The emergence of better digital signal processing has further diminished the need for optimised cabinet geometries, leaving the potential of alternative shapes mostly unexplored.

Addit Audio takes a new approach to cabinet design and fabrication. By utilising the technique of **Large Format Additive Manufacturing [LFAM]**, they are able to create cabinets in organic shapes without the high piece-to-piece production costs of the aforementioned methods. One such cabinet, Addit's Monihorn, can be seen in Figure 1.1. LFAM, like small-scale 3D printing, allows for a more iterative design process and slight tweaks to a design without having to manufacture new moulds or alter part of the fabrication process, which is beneficial from a sustainability standpoint. The more organic cabinet designs enabled by LFAM can be beneficial for the reproduction and dispersion of acoustics, since distortive effects such as diffraction and resonances can be mitigated through the omission of sharp geometrical features. The use of LFAM does pose problems to cabinet design as well, especially when considering that the material properties of 3D printed objects are anisotropic, or transversally orthotropic at best.



Figure 1.1: Addit Audio's 3D printed loudspeaker system 'Monihorn'.

Vibroacoustics encapsulates the interaction between pressure waves interpreted as sound and a solid structure. In traditional loudspeakers, typically consisting of a cabinet and a driving unit or '*driver*', this interaction is perpetual. Vibroacoustic effects act in a bi-directionally coupled system; Pressure acts on a solid structure causing it to move or vibrate and said vibrations in turn generate new pressure waves, affecting the original pressure field. This bi-directionality makes the study of vibroacoustics an inherently coupled problem, spanning multiple physical domains and typically requiring numerical computation to solve.

Solving vibroacoustic problems to aid the design of sound dampers and mufflers has seen recent attention [1] [2] but widespread application in the context of loudspeaker design remains limited. Multiple areas of loudspeaker design can benefit from vibroacoustic analysis. The present research focuses on the design of the cabinet, which houses one or multiple drivers, any acoustic modifiers, and electronics. The combination of new geometric features and structures, new materials and properties, and a general lack of research on vibroacoustic effects seen in loudspeakers form the basis and motivation for this research.

1.2. Challenges in Loudspeaker Cabinet Design

The challenges within loudspeaker cabinet design lie in what vibroacoustics tries to solve for: an interconnected and interdependent relationship between the physics involved. Besides the multiphysics-nature of cabinet design, another difficulty stems from the search for the ideal loudspeaker being a multi-objective endeavour. Disregarding subjectivity and taste, loudspeaker design objectively contends with several acoustic phenomena such as resonances, diffraction, non-linear responses, and viscous losses when the performance of a particular design is to be predicted and tuned. Depending on the type of cabinet, acoustic modifiers may also be incorporated to boost particular frequencies or improve the listening experience when the listener is sitting off-axis, adding to the complexity. Inherent to multi-objective problems, improving the performance in one area is very likely to worsen it elsewhere, especially when considering acoustic modifiers since their optimal design space is very narrow. Loudspeakers also never operate at a single frequency. Their range may be limited to a particular band only, but any optimisation of design has to be considered for tens, hundreds, or thousands of frequencies.

In the current consumer industry, cabinet design is largely driven by intuition, extensive prototyping, and aesthetics. The niche market aimed at **High-Fidelity** [HiFi] audio enthusiast sees more effort being put into the acoustic performance of cabinets, but even there the design process relies heavily on building physical prototypes, testing them, and making changes accordingly. Yet, while such measurements of performance give a good indication of the current state of a cabinet, they do not reveal what geometric changes should be made for the acoustics to be improved. Iterative prototyping nets effective results but can be wasteful of resources and does not provide information on modal performance that can be transferred to different cabinet geometry.

1.3. Research Gap and Objectives

The research gap is identified as the potential benefit of LFAM as a production method of loudspeaker cabinets, evaluated through multiphysics simulation and subsequent optimisation. An efficient method is developed to accurately predict the structural and acoustic behaviour of cabinet geometries when subjected to acoustic and mechanical excitation. This vibroacoustic analysis is used to predict cabinet resonances and distortions, such that they can be filtered out preceding prototype production. This filtering of distortions is to be performed through **Shape Optimisation** [SO], where the definitions of the relevant objective functions follow from the multi-objective problems found in traditional cabinet design.

Several methods exist for solving vibroacoustic problems, more commonly referred to as an **Acoustic-Structure Interaction** [ASI] problem. These methods typically rely on purely **Finite Element Modelling** [FEM] or a combination of FEM with **Boundary Element Modelling** [BEM]. Vibroacoustic shape optimisation of loudspeaker components has seen recent development [3], but its application to LFAM-enabled cabinet geometry that targets resonance reduction remains unexplored. Bokhari et al. [4] present a method in which a loudspeaker's output power is maximised through material distribution-based **Topology Optimisation** [TO], yet no structural dynamics are considered. Andersen et al. [5]

present a method to perform gradient-based SO to promote uniform directivity of a loudspeaker cabinet, again with no structural dynamics being considered. Cabinet dynamics are often disregarded and more focus is put on the optimisation of acoustic horns and lenses, such as in the research by Udawalpola and Berggren [6] and Andersen et al. [7], respectively. This omission is mostly attributed to the high computation cost involved with combining the multiple modelling methods required to form a hybrid that is suitable for a full vibroacoustic problem, including structural cabinet dynamics and sound radiation.

Research by Grande [8] shows that cabinet vibrations produce an audible response, especially when the spectral (or *decay*) behaviour is considered. The existence of this phenomenon is widely acknowledged in the world of Audio Engineering and has resulted in most loudspeaker manufacturers implementing bracing and damping to their cabinets, on top of preferring dense materials for the fabrication. The present research expands on Grande's work, going beyond the sole confirmation of the existence of panel resonances and introducing their classification, along with directivity and **Sound Pressure Level** [SPL] metrics. The objective is to develop and validate a simulation workflow to identify cabinet resonances and their relative contribution to the total audible response, expressed through the **Quality-factor**, or **Q-factor**. The Q-factor is utilised as the primary objective function to formulate the optimisation problem that is solved using SO. The fabrication of these optimised geometries is enabled through LFAM, thus their practical feasibility is evaluated within this context. In short, this research applies numerical simulation and shape optimisation to explore vibroacoustic effects in 3D-printed loudspeaker cabinets.

1.4. Scope and Contributions

This research includes the exploration of both a FEM and a hybrid FEM/BEM simulation method for solving the ASI problem. Contributions are made to exploring the applicability of the stand-alone FEM method for acoustic analysis. The hybrid FEM/BEM method is concluded to be the most suitable for solving the ASI problem and is used to develop a method for full vibroacoustic simulations of loudspeaker cabinets. In particular, cabinets that operate in a low frequency regime using a single narrowband driver. The FEM/BEM simulation model is developed in COMSOL Multiphysics® and is used to quantify loudspeaker performance using SPL metrics, directivity, and the Q-factor. This simulation method is then used to predict performance of a fabricated prototype of Addit Audio that falls within the applicability bounds. These predictions are subsequently compared to real-life measurements performed under anechoic conditions as a means of validation of the method. The method is used within a shape optimisation framework as a means of explorative research on the implementation of shape optimisation on hybrid FEM/BEM simulations that solve the ASI problem. As such, a framework is developed for computationally analysing and optimising the structural and acoustic performance of a low-frequency (20-200 Hz) sealed loudspeaker cabinet. Insights are provided on the transferability and manufacturability of designs optimised through the developed framework.

1.5. Structure of the Thesis

Chapter 1 serves as an introduction to the research and problem statement. Chapter 2 compares two simulation methods that are explored for solving the ASI problem, concluding with a motivation for the chosen method. Chapter 3 presents the details of the simulation model and their implementation in COMSOL Multiphysics®. Chapter 4 presents the performance metrics considered within the optimisation framework. Chapter 5 serves as a first exploration of these metrics in how they are associated with real-life measurements and performance. Chapter 6 presents the optimisation framework, its results, and interpretation, including brief analysis of an alternative optimisation method. In Chapter 7, the results of the research including insights and limitations are discussed. Chapter 8 concludes the thesis by summarising the key findings, highlighting its contributions to vibroacoustic design, and implications and an outlook.

2

Modelling Approaches: Comparative Exploration

The vibroacoustic problem considered in the present research involves coupled pressure acoustics and a solid structure. The structural domain is governed by Navier's equation of motion [9], seen in Equation 2.1. Here, ρ_s is the structural density, \mathbf{u} the displacement field, $\boldsymbol{\sigma}$ the Cauchy stress tensor, and \mathbf{f} the body force vector.

$$\rho_s \ddot{\mathbf{u}} - \nabla \cdot \boldsymbol{\sigma}(\mathbf{u}) = \mathbf{f} \quad (2.1)$$

the acoustic domain is governed by the Helmholtz equation, seen in Equation 2.2. Here, x is the position vector, p is the pressure field (Δp its Laplacian), k the acoustic wavenumber, and $g(x)$ represents the source (e.g. a vibrating surface or monopole source).

$$\Delta p + k^2 p = -g(x) \quad (2.2)$$

At the interfaces of the domains, where the two interact, both are coupled through continuity of normal acceleration and pressure traction, as per equation Equation 2.3 and Equation 2.4, where ρ_f is the fluid density and \mathbf{n} the outward normal to the interface. The choice of method for handling the exterior acoustic domain is the primary distinction between the available simulation methods for solving the ASI problem.

$$\nabla p \cdot \mathbf{n} = -\rho_f \ddot{\mathbf{u}} \cdot \mathbf{n} \quad (2.3)$$

$$\boldsymbol{\sigma} \cdot \mathbf{n} = -p\mathbf{n} \quad (2.4)$$

The core of the ASI problem consists of an acoustic and structural domain that are solved simultaneously. Considered in isolation, either domain is typically modelled and solved for using Finite Element Modelling. When the physics are coupled, FEM can still be used to resolve all physics, but the need arises for some way to resolve acoustic waves propagating into open space. To assess the suitability of each available method, the problem is distilled further to obtain specific boundary conditions and modelling requirements. However, even though the requirements typically dictate quite clearly what method may be best suited, not all problems offer an equally clear answer and preliminary exploration of multiple methods may be advantageous in characterisation of the exact problem and its variables.

The present research concerns such a problem with ambiguous requirements. The vibroacoustic model consists of a solid-domain loudspeaker cabinet including an illustrative driver with single-surface diaphragm, an internal acoustic domain, and an external acoustic domain represented by the void. A digital rendering of this model can be seen in Figure 2.1. Ambiguity stems from the frequency range of interest and the size of the loudspeaker. The cabinet response is evaluated over 20-200 Hz, the

cabinet dimensions are $400 \times 500 \times 458$ mm (W \times H \times D). The primary goals of the present research is to perform shape optimisation, motivating the need for high efficiency and advocating the use of a FEM-only method. This argument however does not hold when the frequency range of interest is low and the domain of interest modest to large in size, as is evident in the discussion of meshing strategy in Section 3.6. In such case, the required element-count, or **Degrees of Freedom** [DOFs], for accurate simulation is substantially high that an alternative method may be favourable. BEM omits the need for discretisation of the external acoustic domain, eliminating high element-count. This chapter describes the differences between a FEM-only and FEM/BEM hybrid method and elaborates on the ultimate decision for the presented ASI problem.

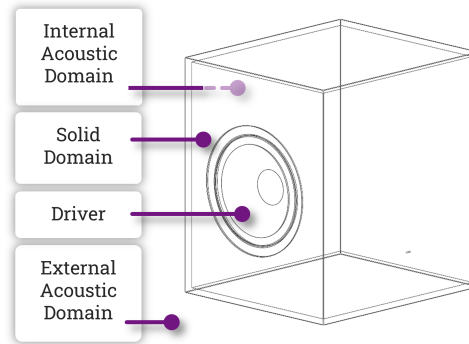


Figure 2.1: Schematic rendering of the 3D model used for vibroacoustic simulation in COMSOL, including domain definitions.

2.1. Overview of Vibroacoustic Simulation Methods

FEM-only and FEM/BEM methods share the discretisation method for the solid domain but differ in how they numerically treat radiated outgoing waves within external acoustic domains. In a FEM-only approach, outgoing acoustic waves are truncated through a volumetric mesh of finite size and absorbed in an artificial outer layer, known as a **Perfectly Matched Layer** [PML]. In a FEM/BEM approach, the external domain is not meshed at all and instead the external acoustic radiation is represented through boundary integral equations that satisfy the *Sommerfeld radiation condition* [2], ensuring that acoustic energy is radiated outwards and no energy may radiate back from infinity.

These two methods are considered the most applicable for the ASI problem at hand. Several other methods for modelling exterior radiation in ASI problems are presented in Appendix A. FEM-PML and FEM/BEM are considered in more detail.

2.1.1. FEM-PML

The FEM-PML method remains a FEM-only method by relying on a volumetric mesh that surrounds the radiator of interest. Correct exterior radiation is enforced numerically by the outgoing waves being attenuated within the PML before they reach its outer boundary. It can be recognized that this method only works correctly if the waves entering the PML are indeed attenuated to negligible amplitudes. If not, these waves may reflect back into the exterior FEM domain, causing PML reflections and artifacts, denoted '*PML Modes*'. Numerical accuracy depends heavily on the correct configuration of the PML, which depends on the frequencies of interest, size of the solid structure, level of detail of the solid structure, PML thickness, and absorption profile [10].

The FEM-PML approach is advantageous for keeping the setup of the simulation simple and is straightforward to couple to an existing FEM structure, providing good near-field detail around complex geometries and easy multiphysics integration. In terms of numerical solver configurations, standard FEM solvers and pre-conditioners can be applied, as well as domain decomposition and parallelization for parallel computing. Furthermore, FEM-PML is compatible with non-linear and transient analysis, though this is not of interest in the present research.

In a FEM-PML strategy, the element count of the volumetric mesh scales with the cube of frequency. An acoustic wave of 20 Hz has a wavelength of $\lambda = \frac{c}{f} = \frac{343 \text{ m/s}}{20 \text{ Hz}} = 17.15 \text{ m}$ in air. Correctly resolving

such a wave using PML requires a volumetric mesh (consider it a sphere that surrounds the radiator) with a radius of $0.5 - 1\lambda$, or $8.08 - 17.15$ m. As such, element-count grows rapidly when evaluating lower frequencies, increasing computational cost to prohibitive levels. An acoustic mesh can be coarser for lower frequencies, as is discussed in Section 3.6, but this is a linear relation with wavelength as opposed to the cubic relation between PML size and element count and provides an insignificant advantage.

2.1.2. FEM/BEM

In a FEM/BEM method, only the structure and internal acoustic domain are discretised with FEM. Instead of relying on an artificial outer boundary, BEM satisfies the Sommerfeld radiation condition for the outgoing acoustic field. The Sommerfeld radiation condition is a constraint that 'selects' only the outgoing solutions of the Helmholtz equation, thus ruling out solutions of the Helmholtz equation [11] that correspond to energy coming in from infinity. The Helmholtz equation is presented in Equation 2.2. The Sommerfeld Radiation condition can be seen in Equation 2.5, where $r = |x|$.

$$\lim_{r \rightarrow \infty} \left(\frac{\delta u}{\delta r} - iku = 0 \right) \quad (2.5)$$

Due to this construction only containing outgoing radiation, reflections back into the exterior acoustic domain are avoided. This analytical approach promotes highly accurate radiation modelling and improves far-field predictions. In terms of solver configuration, BEM does require dense matrices to be solved and is computationally more expensive than FEM calculations per frequency. However, due to only needing surface mesh elements, element-count remains much lower compared to an equivalent FEM-PML model, especially when the solid structure is small compared to the wavelengths of interest. This implies that a FEM/BEM method is not suitable for higher frequencies. These advantages reveal a method that is well-suited for extraction of far-field quantities such as directivity patterns and radiated power, both highly relevant to loudspeaker cabinet design. BEM becomes the preferred modelling strategy when a PML domain becomes prohibitively large, though care has to be taken to keep the computational cost manageable, especially in the context of an optimisation framework.

2.2. Motivation for Hybrid FEM/BEM Approach

The FEM-PML method is explored in detail to observe the applicability within the present research. Details of this implementation and the subsequent results are found in Appendix B. Asymmetric responses and numerical inaccuracies found in the FEM-PML method motivates the use of a hybrid FEM/BEM approach, summarised as follows:

- The FEM-PML method proves inadequate for providing accurate near-field, far-field, and structural responses. The representation of exterior radiation requires extensive meshing of the external acoustic domain and increases element-count to prohibitive levels when compared to a FEM/BEM method.
- The primary metric, the Q-factor, cannot be obtained reliably using a FEM-PML method when PML modes cause inconsistencies in the structural response (see Chapter 4 for details). These inconsistencies manifested as omission of Q-factor peaks and a mean absolute error with the FEM/BEM method of 0.91. An appropriate and reflection-free PML configuration proved infeasible, hindering the use of the Q-factor in shape optimisation.
- The FEM/BEM method offers a good compromise between accuracy, computational efficiency, and suitability for optimisation. The structural-acoustic coupling is strong and the method is proven to be robust and reliable in existing literature [12] [13]

Thus, the FEM/BEM method is selected over the FEM-PML method. Its discretisation of only exterior boundaries avoids the large computational cost associated with an appropriately-sized volumetric PML domain for reliable broadband numerical accuracy.

Final Modelling Framework

The FEM/BEM method is established as the most suitable method for solving the ASI problem in the present research, based on the numerical robustness and the method being well-suited for use in a shape optimisation framework. Computational efficiency is another key advantage of FEM/BEM.

The implementation of the method is presented in this chapter, touching upon key aspects including excitation modelling, material modelling, the structural-acoustic coupling, meshing and solver strategy, and concluding with numerical convergence and validation of the method. This method serves as the backbone for the optimisation framework presented in Chapter 6, all subsequent results are therefore obtained using this implementation.

3.1. Overview of Hybrid FEM/BEM Method

The hybrid FEM/BEM method partitions the ASI problem into three domains: The *internal acoustic domain*, *external acoustic domain*, and the *structural domain*.

The internal and structural domain are discretised using FEM. BEM is applied to the external domain, interacting with the structural domain through shared boundaries. The external domain does not have to be discretised. As established in Section 2.1.2, BEM satisfies the Sommerfeld radiation condition for the outgoing acoustic field, allowing for far-field pressure to be extracted without additional post-processing. Together, this hybrid formulation enables extraction of radiation impedance, panel resonances, frequency-domain Q-factors, and directivity.

The acoustic-structure interaction is as follows: Movement of the driver's diaphragm generates pressure waves on its front- and backside, named front- and rear-radiation. The front-radiation radiates into the external BEM domain, the rear-radiation radiates into the internal FEM domain. The external radiation is propagated through the radiation condition. The internal radiation is propagated through the FEM domain and interacts with interfacing boundaries. These boundaries are part of the structural domain. The internal radiation causes pressure fluctuations that generate vibrations and deformations in the structure. Said vibrations and deformations are carried through the structure and in turn generate new pressure waves in both acoustic domains, describing the acoustic-structure interaction. Note that the BEM domain does not feed back upon the FEM domain directly, but it does enforce radiation loading at the interface [14].

Limitations of Hybrid FEM/BEM Method

Limitations of the hybrid FEM/BEM method in the present research are particularly relevant to low-frequency studies. These limitations are highlighted here.

- In terms of solver configuration, BEM requires preconditioning and computation of dense and often ill-conditioned matrices [14], adding to computational cost. The implications of this high cost are discussed in Section 3.6.
- In BEM modelling, the exterior air domain is assumed homogeneous, thus any inclusion of viscothermal effects or acoustic non-linearities requires extra modelling effort, e.g. through the extension of the FEM domain or impedance boundary conditions [14]. This limitation is prevalent in

low-frequency loudspeakers as these often contain acoustic modifiers such as waveguides and damping. These effects are observed and discussed in Chapter 5.

- Another limitation relevant to the present research is the strong vibroacoustic interior-exterior coupling of the FEM/BEM method. The structural behaviour, even small cabinet modes, significantly impacts impedance and radiation conditions [12]. This phenomenon is observed and discussed in Section 6.4.

3.2. Geometry Definition and Parametric Cabinet Representation

The geometry definition consists of two distinct parts: The *driver* and the *cabinet*. Both are modelled separately and combined into an assembly that describes the full geometry. The shape of the driver and cabinet is arbitrary to the simulation model but not to each other as they have to interface properly.

Driver Geometry

The driver model is simplified for efficiency. The key components include the diaphragm, the surround, and the mount. Each component is modelled as a single surface with no thickness since it only serves as a boundary that separates the acoustic domains and allows for the application of excitation. The dimensions are taken from technical drawings provided with the driver. For the presented simulation model, a low-frequency and low-cost 10" subwoofer driver has been selected: The Dayton Audio SD270A-88 [15]. Further detail on the usage of this driver in the simulation model is described in Section 3.3. Figure 3.1a shows an image of the selected driver, Figure 3.1b shows a digital model of the same driver of medium detail, and Figure 3.1c shows the simplified digital model used for the simulations.




Figure 3.1: Dayton Audio SD270A-88, represented as real item (a), detailed CAD model (b), and simplified CAD model used for vibroacoustic simulation (c).

Cabinet Geometry

The cabinet is modelled with the minimal geometry that effectively describes modal performance. Omitting small details is beneficial to improving simulation cost and robustness, but the removed details must be non-critical or non-dominant in the overall cabinet behaviour.

The simplified cabinet is rectangular, as seen in Figure 2.1, which is *not* a geometry that is interesting in the context of LFAM, but is easier to interpret. The cabinet is modelled as a solid with walls of consistent thickness. The cut-out for the driver is made to fit the chosen driver and small details like mounting ridges or bolt holes are omitted. The ratio of the outer dimensions is arbitrary but the internal volume is made to match what is recommend for optimum efficiency of the driver.

Two variations of the cabinet are made: one with a Helmholtz resonator and port and one without. The resonator is modelled as a tube and tuned to a frequency of 36 Hz. This tuning frequency is calculated using an online calculator [16], and the walls of the resonator are made equally thick to the cabinet walls. The port, or the '*mouth*' of the resonator, is flared with a constant radius. This is often done in real speakers to aid dispersion and prevent turbulence of the outgoing radiation [17]. The location of the resonator is considered arbitrary, here placed on the width-wise centre plane to keep the geometry symmetric over said plane. The placement of the resonator *does* affect simulation results,

but these resonator effects are not within scope of the present research. The resonator variation of the cabinet is used only in Chapter 5 for comparison. The rectangular cabinet has several options for parameterisation that are explored through sensitivity analysis in Section 4.4, such as the panel thickness and curvature. The resonator is not considered in the optimisation. The baffle cut-out is also driver-dependent and equally excluded from any optimisation.

A final note on cabinet design is made in the context of LFAM. While LFAM offers form-freedom, several constraints do still have to be followed. For loudspeaker cabinets in particular, the printed structure needs to be air-tight and the structure should be built up using a **continuous contour**. Moreover, the current LFAM process does not allow for intermittent variations of printing thickness, thus imposing that a printed cabinet structure has a consistent wall thickness. This process parameter is considered within the optimisation framework. Figure 3.2 shows the rectangular cabinet model as described, with (a) and without (b) the Helmholtz resonator. Figure 3.2c shows the cabinet model that is used in Chapter 5 for experimental validation of the simulation model.

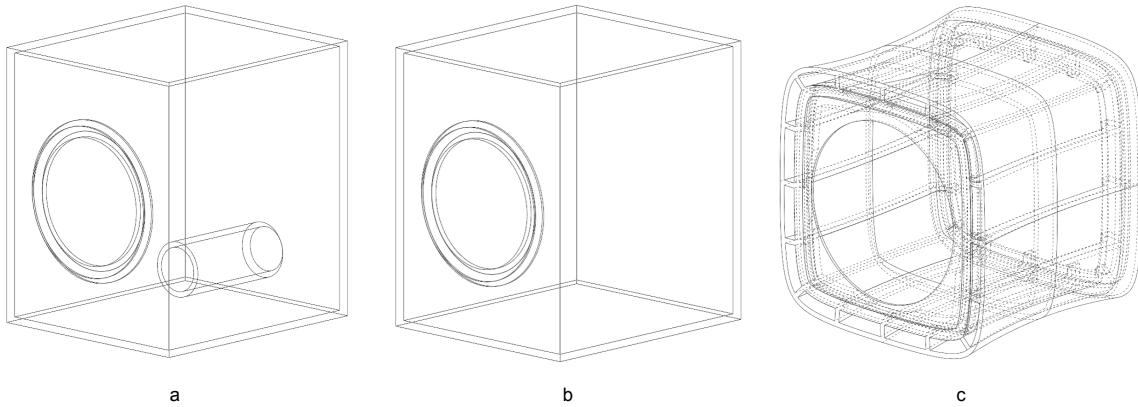


Figure 3.2: Cabinet geometry used within the developed simulation model, rectangular with Helmholtz resonator (a), rectangular without Helmholtz resonator (b), and prototype cabinet used for experimental validation (c).

3.3. Driver and Excitation Modelling

The driver provides acoustic excitation by movement of its diaphragm. **Thiele/Small [TS]** parameters [18] are commonly used to construct an equivalent electromechanical circuit, a practice better known as **Lumped Parameter Modelling [LPM]**. The TS Parameters are a set of electromechanical properties that describe the frequency-dependent response of a driver. This response is often sufficient for simulations and cabinet design, unless a detailed structural response of the driver is required. However, in the presented implementation an explicit excitation is applied to the diaphragm instead.

Drivers also typically come with testing data that typically includes a **Frequency Response [FR]** curve and impedance data of the electromechanical system. The FR curve shows SPL for each frequency, along with phase data of the diaphragm. In the present research, these parameters are transformed to an equivalent normal velocity of the simulated diaphragm that replicates the original FR curve. Using this normal velocity as the acoustic excitation provides a more accurate representation of the driver contribution than LPM does, including all mechanical, electrical, and acoustic effects.

Equation 3.1 shows the conversion from FR data $SPL_{input}(f)$ as a function of frequency (f) to equivalent normal velocity v_n . p_{mag} is the acoustic pressure magnitude at 1 m from the driver, the reference point for the FR measurements as per the manufacturer. p_{ref} is the reference pressure (20×10^{-5} Pa), p_{1m} the complex pressure at 1 m, $\phi_{ac}(f)$ the acoustic phase, v_v the complex volume velocity, r_0 the reference measure of 1 m, ρ_0 the air density (here 1.225 kg/m^3), S_d the surface area of the diaphragm, and v_n the complex normal velocity averaged over S_d . Figure 3.3 shows a comparison between the imported FR data and the simulated FR response obtained using the equivalent normal velocity, evaluated at the 1 m reference point.

$$\begin{aligned}
p_{mag} &= p_{ref} \left(10^{\frac{SPL_{input}(f)}{20}} \right) \\
p_{1m} &= p_{mag} \cdot e^{i \cdot \phi_{ac}(f)} \\
v_v &= p_{1m} \left(\frac{4\pi r_0}{i\omega \rho_0} \right) \\
v_n &= \frac{v_v}{S_d}
\end{aligned} \tag{3.1}$$

This transformation holds under the assumption of *piston* diaphragm motion, i.e. that it moves as a rigid body. This assumption is often valid at low frequencies but becomes less valid with higher frequencies and larger diaphragms, as the chance of *cone breakup* increases [19]. A practical rule of thumb is the Rayleigh criterion [20] (Equation 3.2), where k is the wavenumber $\frac{2\pi}{\lambda}$ with λ being the wavelength, a piston radius, f frequency, and c the speed of sound in air (= 343 m/s). For the selected 10" driver and a maximum frequency of 200 Hz, $ka = \frac{2\pi \cdot 200 \cdot 0.127}{343} = 0.47 \ll 1$, thus the piston assumption holds. For the cabinet prototype used in the experimental validation of the simulation model, piston assumption does not hold. A compensation for *non-piston behaviour* is presented and implemented in Section 5.3.

$$ka = \frac{2\pi f a}{c} \ll 1 \tag{3.2}$$

Besides acoustic excitation, a driver also provides mechanical excitation through its moving mass. Using the TS parameters of mechanical compliance, stiffness, and mass, a rudimentary **Lumped Mechanical System** [LMS] can be modelled in the form of a mass-spring-damper system. The complex frequency dependent resultant force that follows from this system is applied to the boundaries that represent the driver mount on the baffle. Figure 3.4 shows the simulated SPL for each excitation applied in isolation. The acoustic excitation follows from the equivalent normal velocity of the diaphragm, the mechanical excitation from application of the reaction force to the driver mount boundaries. The acoustic excitation dominates the response. However, the mechanical excitation drives structural resonances in the cabinet panels that are not captured by stand-alone acoustic FR analysis. The non-zero mechanical FR curve indicates contribution to the total radiated sound. Only considering acoustic excitation would under-predict the total radiated SPL at structural resonance frequencies.




Figure 3.3: Comparison of imported and simulated FR data, obtained through transformation of the imported data to an equivalent normal velocity of the diaphragm boundaries.




Figure 3.4: Comparison of the contribution of isolated acoustic and mechanical excitation on the reference SPL, showing that acoustic excitation dominates over mechanical excitation.

3.4. Material Modelling of 3D-Printed Polymers

Simulating the structural response of a loudspeaker cabinet requires proper material characterisation. Conventional cabinets are fabricated from isotropic materials such as MDF, plywood, or aluminium, which behave predictably. LFAM introduces a new challenge, as the layer-by-layer process produces anisotropic parts. Considering that highly accurate modelling of 3D-printed polymers is outside of scope, the material is approximated as *transversely orthotropic*. This section describes the characterisation of this approximation.

The chosen polymer is **Dahltram C-250CF** by *Airtech Advanced Materials Group* [21], the same material used for the prototype cabinet in Chapter 5. This is a modified polycarbonate with 20% carbon fibre loading for added stiffness. A top-level comparison to conventional materials is presented in Table 3.1. Note that the reported Young's Modulus for C-250CF is in the X-direction, along the printing direction; in the transverse direction it drops to 3 GPa.

| Property | C-250CF | MDF | Birch Plywood | Aluminium 6061-T6 |
|--------------------------------|------------------|------------|---------------------|-------------------|
| Density [kg/m ³] | 1210 | 700–800 | 650–700 | 2700 |
| Young's Modulus [GPa] | 11 (X-direction) | 3.5–4.5 | 10–14 (along grain) | 69 |
| Damping Coefficient [η] | 0.015* | 0.01–0.015 | 0.005–0.01 | 0.0001 |

*[22]

Table 3.1: Material properties of common loudspeaker cabinet materials.

The reported density of C-250CF is higher than MDF and plywood, yet lower than aluminium. The Young's Modulus along the printing direction is higher than MDF and similar to plywood. The internal damping coefficient is assumed similar, as no validated value is provided. The conclusion is that C-250CF performs comparably to MDF and plywood when modelled as isotropic.

Transverse orthotropic modelling assumes that material properties are equal within a single printed layer. This holds best when printing lines align with the in-plane global axes. The transverse properties are applied in the remaining direction, representing the build direction. The damping coefficient $\eta_{mat} = 0.015$ is applied by defining a complex Young's Modulus: $E(1 + i\eta_{mat})$.

The material is applied in two different directions to illustrate the transverse orthotropic behaviour on the structural response. Figure 3.5a and b are representations of the 3D printing layers for both tested directions. Direction 1 is used in development of the modelling and optimisation framework. Figure 3.5c shows the structural response of both simulations, expressed using the Q-factor. The Q-factor is extracted using the **Elastic Strain Energy** [ESE] according to Equation 3.3, explained in detail in Section 4.2. The Q-factor response reveals unique modal behaviour for the different transverse orthotropic definitions, confirming that non-isotropic material property definitions alter structural behaviour.

$$Q = \frac{|\Re(ESE)|}{|\Im(ESE)|} \quad (3.3)$$

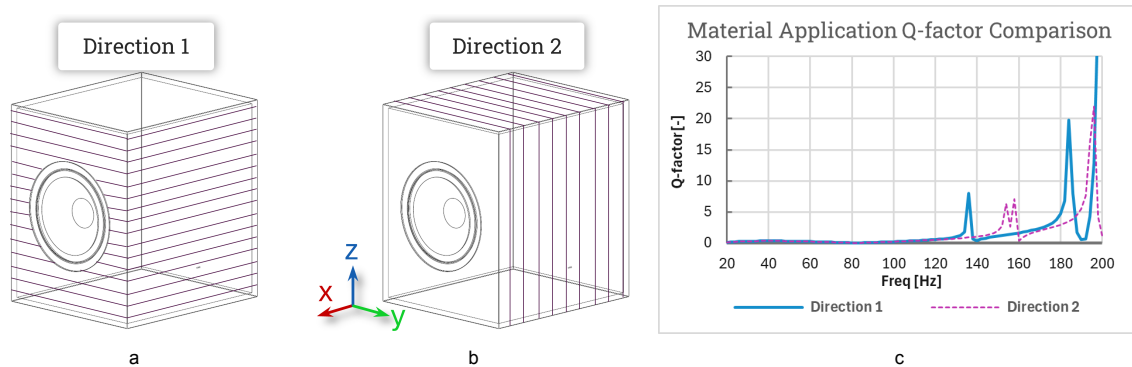


Figure 3.5: Representations of transverse orthotropic material application in global z-direction (a) and global x-direction (b) and their respective Q-factor response (c).

3.5. Structural-Acoustic Coupling and Domain Setup

Section 3.1 conceptually describes the domain setup, including the explicit couplings that are made to form the ASI problem. This section provides details on exact domain definitions, interface conditions, and how the coupling is handled.

3.5.1. Domain Definitions and Boundary Conditions

The internal acoustic domain is defined using COMSOL's FEM Pressure Acoustics module. The cabinet walls that surround the domain are initially assigned a sound-hard boundary condition, which is subsequently overridden by the acoustic-structure coupling, enabling vibration of the walls. In practice, only the coupling boundary condition is active. The normal velocity described in Section 3.3 is applied to the diaphragm boundaries and is set to be applied inwards to the domain. This is analogous to the rear-radiation.

The external acoustic domain is defined using COMSOL's BEM Pressure Acoustics module. Geometric and acoustic symmetry is applied along the width and height planes of the rectangular cabinet (see Section 3.7.2 for more details). The same normal velocity is again applied to the boundaries of the diaphragm, now analogous to the front-radiation. The normal velocity from the FEM acoustic environment is not inherently transferred to the BEM environment and thus has to be explicitly applied here as well.

The structural domain is defined using COMSOL's Solid Mechanics module. The material as described in Section 3.4 is applied to the structure. The mechanical excitation as described in Section 3.3 is applied as a total force to the boundaries that form the driver mount. These boundaries are part of the structural domain and thus rigidly connected to the structure. Symmetry is again applied and the aforementioned boundary load is quartered to accommodate for quarter-symmetry.

3.5.2. Multiphysics Couplings

Two multiphysics interfaces are used to describe the couplings. The first one connects the acoustic and solid FEM domains by enforcing displacement continuity and pressure-force equilibrium, allowing the internal pressure to act upon the structure and cause vibrations and deformations. Said vibrations in turn act back upon the internal domain, resulting in a two-way interaction.

The second interface connects the solid FEM and acoustic BEM domains by enforcing velocity continuity, transferring the structural vibration velocity to the external domain, where it generates the outgoing radiation. This radiation is combined with the active radiation emanating from the diaphragm. The acoustic BEM domain provides sound radiation damping but does not act back upon the structure. Figure 3.6a shows the definitions of the domains as described and Figure 3.6b shows the multiphysics coupling interfaces.




Figure 3.6: Schematic overview of domain definitions used to describe the vibroacoustic model (a) and the definition of the multiphysics couplings (b).

3.5.3. Considerations for Vented Enclosures

To accommodate simulation of the prototype used in Chapter 5, ported designs are briefly considered as well. The domain definition and boundary conditions change slightly when a resonator is included. The same three domains are retained by closing the port, or capping it off, keeping the internal and acoustic domain physically separated. A third multiphysics coupling is then added to the cap, combining the `acpr` and `pabe` environments to numerically couple them. This allows the internal pressure to propagate from the internal to the external domain and contribute to the external radiation.

This setup is only considered to be illustrative and comparative to the measured prototype. Accurate port representation is a known limitation of the FEM/BEM method [23] [24]. Correctly modelling such physics typically requires **Computational Fluid Dynamics** [CFD] or corrections for acoustic non-linearities and viscothermal losses beyond the scope of this thesis [25]. The simulation in Chapter 5 is performed with and without this resonator contribution to visualise and quantify the discrepancy. The resonator is disregarded otherwise. Figure 3.7 shows the multiphysics coupling required to include a Helmholtz resonator in the simulation model.




Figure 3.7: Schematic overview of multiphysics coupling required to include a Helmholtz resonator in the simulation model.

3.6. Meshing Strategy and Study Setup

Proper consideration of all physics interfaces is taken regarding meshing and solver configuration. Compatibility and robustness between each interface is essential for an optimisation framework, thus the requirements are considered per interface but also in coupled context. A strategy is required that, besides numerical robustness, focuses on computational efficiency and accuracy. First, justification for the use of frequency-domain studies is given.

3.6.1. Frequency-Domain Analysis

Frequency-domain analysis is chosen to capture steady-state response across the bandwidth of 20-200 Hz. These steady-state responses are used to obtain the SPL, the directivity, and the Q-factor (see Chapter 4 for details). Transient analysis requires long computation runs and further post-processing to extract these parameters. Transient analysis does allow for the inclusion of non-linear phenomena, such as large-amplitude cone excursion, but the simulation in the present research is linear, so this advantage is considered irrelevant. Analysis of transient behaviour, or *decay behaviour*, using time-domain analysis is interesting to quantify cabinet ringing, as demonstrated by Grande [8]. Frequency-domain analysis can only provide an identification of resonant behaviour and not exact quantification. However, identification of the phenomenon is considered suitable for the present research and optimisation framework.

3.6.2. Acoustic Mesh

The acoustic mesh uses quadratic elements. A guideline for adequate mesh resolution is to use 5-6 elements per wavelength [26]. This can be implicitly defined in the model by stating that `Max Element Size` = $343/f_0/6$, where f_0 is the frequency of interest. This would give an adaptive mesh that changes with frequency. This is not desirable for consistent structural calculations as the stiffness matrices become inconsistent which prevents solution reuse. A fixed mesh is defined using the shortest wavelength, i.e. the highest frequency of 200 Hz.

This size expression is applied to both the FEM and BEM acoustic domain using free tetrahedral elements, which are recommended for solving numerical acoustic problems [26]. A slight mesh refinement is made for the diaphragm to ensure accuracy of the excitation and to properly resolve the curved geometry.

3.6.3. Structural Mesh

The structural mesh is initially made to accommodate the acoustic requirements. Free tetrahedral elements are again used for consistency and numerical robustness. A different type of element is built up of different vertices that do not necessarily align with the interfacing tetrahedral mesh. The vertices of both meshes have to re-mapped by the solver to align, increasing computational effort.

Recommendations for the size expression for the structural mesh depend on application and are less consistent. Thus, a convergence study is performed to establish adequate resolution. A common method for structural-acoustic models is to evaluate the location of resonance peaks in the structural response, with a rule-of-thumb being $|\Delta f_r| \leq 1 - 2\%$, Δf_r being the relative movement in frequency of resonant peaks [27].

This condition was tested by varying the mesh between COMSOL's built-in size presets and observing Q-factor behaviour. Details on this convergence study are presented in Section 3.7. The `normal` size preset (Maximum/Minimum element size of 0.05/0.009 m) was found to satisfy the convergence criterion, with a relative peak deviation between the `normal` and `fine` preset of **1-2%**.

This mesh resolution is finer than what is acoustically required, thus the structural mesh is prioritized in the mesh hierarchy. Thus, the acoustic domains inherit the higher mesh resolution at the interfaces. For the internal FEM domain, it can be seen in Figure 3.8 that the mesh becomes coarser further away from the interface, adhering to the 5-6 elements-per-wavelength rule. Mesh sizes are to match at the interfaces to aid with mapping of the continuity conditions satisfied by the multiphysics couplings.




Figure 3.8: Visualisation of the mesh used in the ASI simulation model, showing coarsening of the acoustic mesh further away from the structural mesh.

3.6.4. Solver Configuration

The optimal solver configuration depends on system setup and preference, the goal here being to balance between robustness and efficiency. Empirical studies reveal that, based on the available computing resources and physics involved, a segregated frequency-domain study utilizing iterative solvers is most robust and adaptable. Within the segregated steps, all interacting physics variables are kept together to improve conditioning, the lumped mechanical system is considered in a separate step.

GMRES proves the most reliable solver, with a residual tolerance of 0.01 to 0.001 providing a good balance between speed and accuracy. A single direct preconditioner on the iterative solver is used to prevent ill-conditioning from hybridization. `cuDSS` (CUDA Direct Sparse Solver) gives the fastest solves over other preconditioners, utilizing available GPU resources that are otherwise ignored by preconditioners such as `MUMPS` or `PARADISO`. This setup results in an average computation time per frequency step of 42 seconds on a workstation with a 12-core AMD Ryzen™ 9 HX 370 processor, 16 GB of LPDDR5X RAM, and an NVIDIA Geforce RTX™ 4070 Laptop GPU.

Shape optimisation adds an extra entry in the segregated steps that contains the optimisation variables.

The same setup of GMRES with a single direct preconditioner is applied to all segregated steps, though the optimisation is rather unstable with cuDSS and is better preconditioned using PARADISO. Solution reuse across frequencies and iterations is enabled as well to increase efficiency further. If memory allows, a fully coupled solver made be utilised instead of segregated steps to lower computation times even further.

3.6.5. FEM Cushion Exploration

Statistics of the ASI study show that most computational effort is spent on the BEM-solid coupling. This coupling is inefficient to solve for as BEM generates fully populated dense matrices. Coupling these to a sparse structural FEM system destroys sparsity and increases factorisation cost [28]. To combat this, an experimental method was explored using an acoustic FEM buffer domain between the BEM and structural domain. This method proved numerically inaccurate and was thus discarded. The results of the exploration can be found in Appendix C.

3.7. Verification and Numerical Convergence

To verify the model setup as described thus far, implementation and numerical convergence are considered. This section presents the results of analysis of mesh convergence, the application of quarter-geometry, and the frequency sampling interval required to adequately resolve resonance behaviour for Q-factor analysis.

3.7.1. Mesh Convergence

Structural mesh convergence is verified by comparing Q-factor behaviour across different mesh densities corresponding to COMSOL's `coarser`, `coarse`, `normal`, `fine`, and `finer` presets. Figure 3.9a shows the Q-factor response for each mesh resolution that is used to evaluate the aforementioned convergence criterion $|\Delta f_r| \leq 1 - 2\%$. Figure 3.9b shows a detail view of the same plot for clarity. The sampling interval is set to be equidistant with a step size of 2 Hz to omit sampling inaccuracies. More detail on sampling intervals is provided in Section 3.7.3. The results show several distinct resonant peaks, for each of which the convergence criterion is evaluated. Note that absolute values of the Q-factor are found to be mesh-dependent. The presented convergence criterion ensures that the *location* of Q-factor peaks remain consistent and that trend behaviour is preserved.

From this analysis, it follows that the structural mesh meets the convergence criterion when the `normal` preset is applied. Since this mesh size is much finer than the prescribed mesh size for acoustics of 6 elements per wavelength, no acoustic convergence is carried out. The `finer` mesh provides slight refinement of resonant frequency location but is considered of prohibitively high resolution for development of the framework.

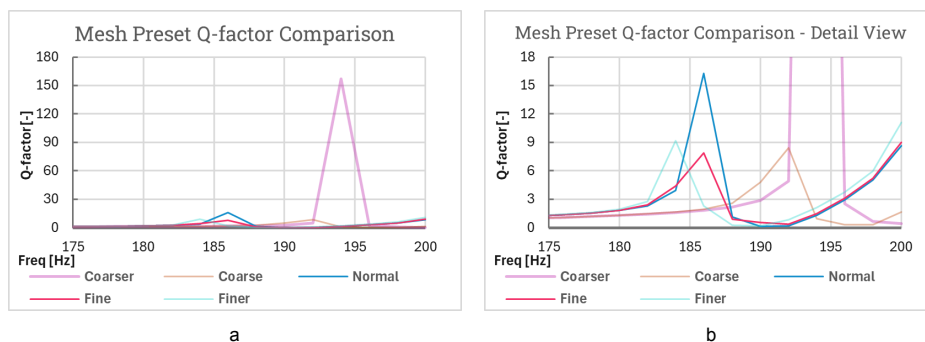


Figure 3.9: Q-factors observed in structural mesh convergence test using varying mesh size presets, demonstrating that a `normal` and `fine` preset exhibit convergence of resonant frequencies.

3.7.2. Quarter-Geometry Validation

The geometric and acoustic symmetry implementation is validated by comparing full-geometry and quarter-geometry simulation results. Figure 3.10a shows the SPL curve obtained at the 1 m reference point for both geometries and Figure 3.10b shows the Q-factor. Both the SPL and Q-factor results show

agreement between the two configurations, with the SPL showing a mean absolute error of 4.05×10^{-3} and the Q-factor showing a mean absolute error of 6.75×10^{-2} . This result confirms correct application of symmetry. The quarter-geometry model is adopted to reduce computational cost by reducing element count by approximately 75%. Note that the Q-factor extracted for the quarter-symmetric model is multiplied by 4 to accommodate for the volume integral being $4 \times$ smaller than that of the full geometry.

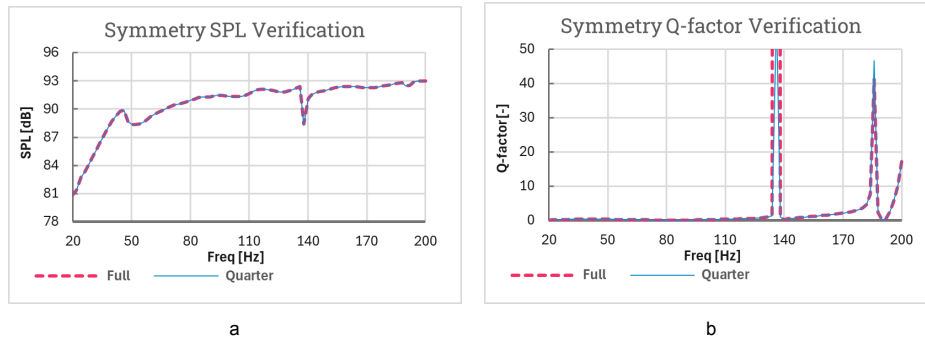


Figure 3.10: Comparison of SPL (a) and Q-factor (b) for full- and quarter-geometry, showing agreement in SPL through a mean absolute error of 4.05×10^{-3} and in Q-factor through a mean absolute error of 6.75×10^{-2} .

3.7.3. Frequency Sampling Interval

Accurate Q-factor extraction requires sufficient frequency resolution to capture resonant peak locations and trend behaviour. Coarse sampling risks missing narrow peaks or mischaracterising trends, leading to artificially low Q-values. Yet, coarse sampling is beneficial for efficiency as it keeps the number of evaluated frequencies lower.

Two sampling methods are explored: using ISO preferred frequencies in different intervals (1/3-octave, 1/6-octave, 1/12-octave, and 1/24-octave sampling) and using equidistant sampling with frequency steps of 1, 2, and 3 Hz.

Figure 3.11 shows that ISO preferred frequencies tend to miss resonant peaks, mostly in the higher frequencies. This is because the gaps between evaluated frequencies become larger on the higher end of the range, leading to undersampling in this region. Equidistant sampling provides better resolution across the range and reveals resonant trend behaviour more accurately. An interval of 2 Hz is chosen to provide a good balance between computational effort and identification of trends in resonant behaviour. A range of 175-200 Hz is shown to highlight the impact of sampling strategy on the higher frequencies, with ISO 1/12 demonstrating undersampling.




Figure 3.11: Q-factors extracted using sampling interval variations, showing that ISO-preferred frequencies undersample higher frequencies while equidistant sampling captures high-frequency Q-factors more reliably.

To assess higher-resolution sampling results, a single resonant peak earlier identified at **186 Hz** is evaluated using yet higher resolution equidistant sampling of 0.2 Hz and 0.05 Hz steps, the result of which can be seen in Figure 3.12. This analysis reveals that uncovering the true magnitude of Q-factor

peaks requires sampling much finer than 2 Hz intervals. The implication of this finding is discussed in Section 4.2.

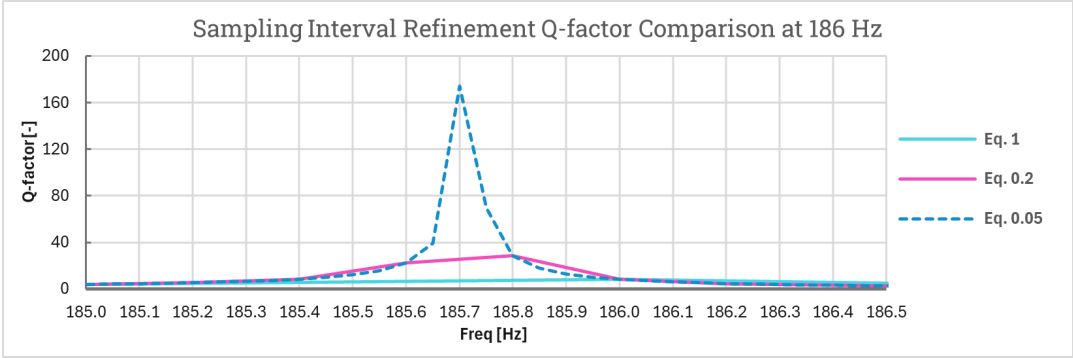


Figure 3.12: Q-factor plot showing that refinement of the sampling interval around Q-factor peak at 186 Hz allows for more accurate extraction of the absolute magnitude, confirming that Q-factor peaks can fall between points of the evaluation grid.

Based on these verification studies, the model setup presented in this chapter is confirmed to be numerically robust and suitable for the optimisation framework presented in Chapter 6.

4

Performance Metrics and Analysis Methods

The simulation method described in Chapter 3 enables prediction of performance metrics within coupled acoustic-structure interactions, such as SPL, **Directivity Index [DI]**, and structural energy metrics like the ESE. The latter enables the extraction of the Q-factor.

This chapter formalises these performance metrics in sections 4.1 through 4.4 to enable use in the optimisation framework established in Chapter 6. The Q-factor is emphasized to be the primary objective function upon which the presented shape optimisation framework is built, enabling the attenuation of audible cabinet resonances. Section 4.4 presents parametric sensitivity studies used to quantify how *baffle edge radius*, *panel thickness*, and *panel curvature* influence each performance metric. This sensitivity analysis shows that directivity is most sensitive to baffle edge geometry, while the Q-factor is influenced more by panel surface features and alterations in spatial panel stiffness.

4.1. SPL Metrics

The Sound Pressure Level, observed at a reference point, is the performance metric most commonly found in characterisation of loudspeaker performance. This metric is often expressed using the frequency response, plotting SPL in dB versus frequency in Hz. In the FEM/BEM framework, the SPL is computed from the radiated complex pressure field at evaluation points. This evaluation point is taken 1 m in front of the loudspeaker in the centre of the main driver. This is analogous to a standard near-field measurement [29].

The pressure magnitude used for calculation of the SPL is extracted from the BEM solution at the reference point. This can be a singular point but an infinitesimally small point cannot capture slight spatial variation and perturbations. Thus, surface integration over a small spherical domain at the reference point is preferred. The SPL is computed using Equation 4.1, where $|p|$ is the pressure magnitude at the reference point and p_{ref} is the reference pressure (here, 20×10^{-5} Pa).

$$\text{SPL} = 20 \log_{10} \left(\frac{|p|}{p_{ref}} \right) \quad (4.1)$$

The FR is objectively considered ideal when it is a flat curve, showing no variation in pressure level over the entire bandwidth of interest. The cabinet influences the FR through panel resonances as demonstrated in Section 3.3, and by guiding propagating wavefronts. In the context of optimisation, a cabinet ideally promotes a flat FR curve and smooth response, discouraging audible colouration. However, the cabinet contribution to a near-field FR is minimal compared to the active radiation from the driver. Thus, SPL is rather used as a metric to validate that shape modifications do not degrade the output, not as a guiding objective function.

Figure 4.1a shows the 1 m reference point representation and Figure 4.1b the FR of the rectangular cabinet simulated using the method presented in Chapter 3.

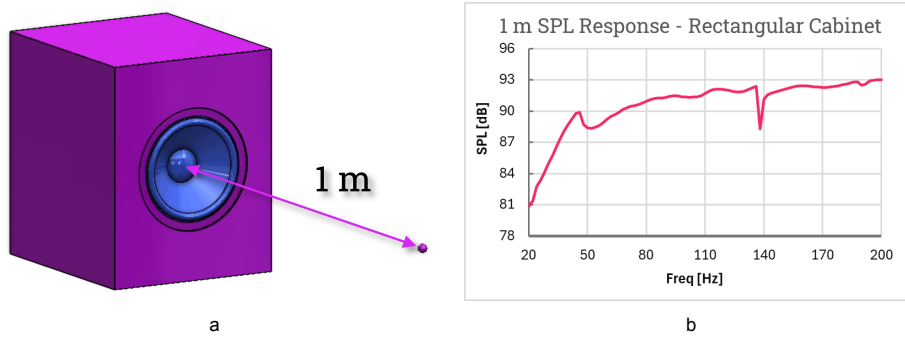


Figure 4.1: Representation of the 1 m reference point used for FR analysis (a) and the FR plot of the rectangular reference cabinet (b) evaluated at said point.

4.2. Structural Energy Metrics & Q-factor Acquisition

The primary structural energy metric is the ESE. This metric quantifies spatial distribution of strain energy within the vibrating structure of interest, providing a frequency-domain response that enables Q-factor extraction and gives insight into resonance damping of a cabinet.

ESE is computed by taking the volume integral of the elastic strain energy density over a structural domain. Within COMSOL, the relevant expression is `solid.Ws`. This expression returns a complex-valued result with a real and imaginary component. The real component represent elastic energy storage while the imaginary component represents dissipative losses from material damping and acoustic radiation. Peaks in the ESE response are identified to correspond to cabinet resonance modes. Figure 4.2a shows the ESE response of the rectangular cabinet, displayed on log scale for clarity. Figure 4.2b visualises the ESE peak at 138 Hz through the stress concentration within the structure at that same frequency.

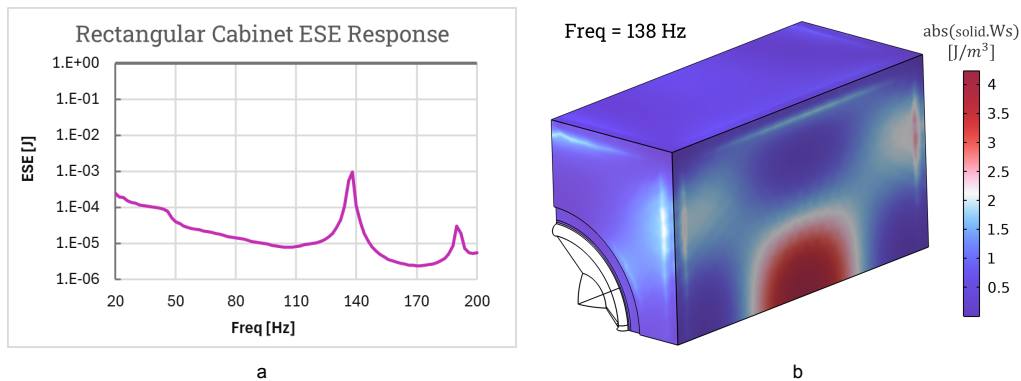


Figure 4.2: ESE response of the rectangular cabinet (a) and visualisation of the stress concentration within the structure at the ESE peak identified at 138 Hz (b).

The Q-factor is extracted using the elastic strain energy by taking the ratio between the real and imaginary component, as per Equation 4.2 (repeated from Equation 3.3 to aid the reader). Note that absolute values are taken to ensure positive Q-factors in cases where imaginary components exhibit sign variations. By formulating the Q-factor as such, the material damping and acoustic radiation damping are both inherently accounted for. Q-factors up to a value of 3-4 are indicative of resonances being adequately damped, while the Q-factor can rise to >10 for frequencies that exhibit large or sharp resonances, indicating audible cabinet ringing and longer decay. The particular Q-factor values presented here are representative of responses observed in the present research.

$$Q = \frac{|\Re(ESE)|}{|\Im(ESE)|} \quad (4.2)$$

Comparing the ESE in Figure 4.3a and the extracted Q-factor in Figure 4.3b shows a physical consistency between both metrics that corresponds to the previously presented concentration of strain energy within the structure. The baseline rectangular cabinet shows Q-factor peaks ranging from **46.6-60.3** at dominant frequencies in the **138-200** Hz band. These high Q-factors indicate significant modal energy concentrations that transform into audible cabinet colourations and ringing.

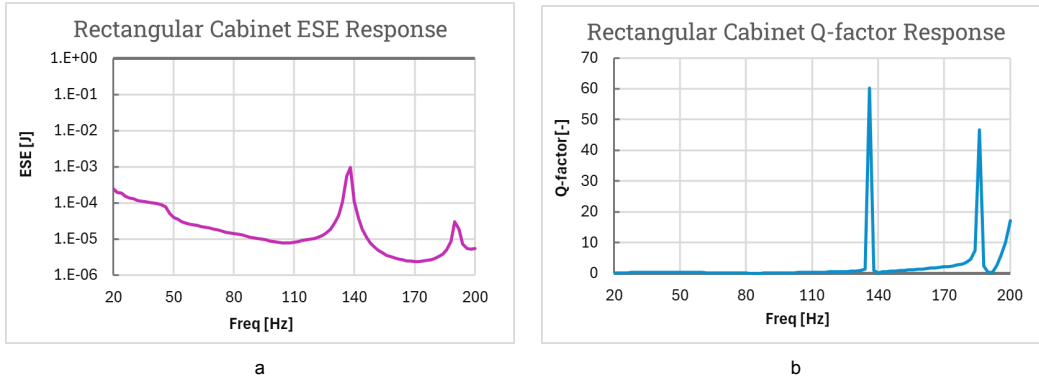


Figure 4.3: ESE response of the rectangular cabinet (a) and the corresponding extracted Q-factor plot (b), showing physical consistency between the metrics.

To minimise cabinet colouration, Q-factor minimisation across the frequency range of interest serves as the primary objective function for the optimisation framework presented in Chapter 6. Physically, the optimisation is targeting reduction of resonant energy accumulation, thereby reducing audible colouration. More insights on the physical mechanism behind the optimisation is presented in Section 6.4.4. A note is also made that absolute magnitudes of ESE and the extracted Q-factor are *mesh-dependent*, while the location of resonant frequencies are not as per the convergence presented in Figure 3.9.

4.3. Directivity Metrics

Directivity characterises the spatial variation of SPL around a loudspeaker and is used to quantify the on-axis vs off-axis performance. A loudspeaker with *lower directivity* exhibits *less spatial variation*, implying that the perceived response for a listener is invariant to their listening position.

The directivity is quantified as a metric through the Directivity Index, relating the SPL variation around a reference point to the on-axis response. The SPL is evaluated at discrete equidistant locations over the azimuth angle. The number of discrete evaluation points is arbitrary, with more points providing increased accuracy. Here, 10 points are used spanning 0-90° with an angular spacing of 10°. The DI metric can be mirrored over the width-wise symmetry plane, having verified symmetric response of the quarter-geometry in Section 3.7. Thus, the spatial response is obtained for -90° to +90°.

Equation 4.3 shows the formulae used to extract the DI from spatial evaluations. Here, I is the intensity (proportional to the squared pressure magnitude), N denotes both the total number of evaluation points and serves as the index for individual points, I_{mean} is the arithmetic mean, W is the directivity factor, I_1 is the reference intensity at 0° azimuth angle, and DI is the Directivity Index. Note that the formula for intensity is $I = \frac{|p|^2}{2\rho c}$ where ρ is fluid density. Since $\frac{1}{2\rho c}$ is a constant that cancels out in the ratio $W = \frac{I_1}{I_{mean}}$, it is omitted for brevity.

$$\begin{aligned}
 I_N &\propto |p_N^2| \\
 I_{mean} &= \frac{1}{N} \sum_1^N I_N \\
 W &= \frac{I_1}{I_{mean}} \\
 DI &= 10 \log_{10}(W)
 \end{aligned} \tag{4.3}$$

Directivity is most sensitive to the shape of the baffle edge due to being dominated by diffraction effects, as shown by Olson [30]. His research shows that a spherical cabinet exhibits the best theoretical performance in minimising diffraction effects and lowering directivity, resulting in a uniform pressure field. This behaviour is confirmed through sensitivity analysis performed in Section 4.4.

The DI is shown for the rectangular cabinet to observe the spatial response of the baseline geometry preceding optimisation. Figure 4.4a shows, for each frequency, the minimum and maximum SPL across all evaluation angles to illustrate the degree of spatial variation in the response and Figure 4.4b shows the DI. The deviation between minimum and maximum SPL increases with frequency. An anomaly is observed around 140 Hz, manifesting in Figure 4.4b as a sharp dip in DI, corresponding to the mean off-axis SPL being higher than the SPL on the reference axis. This is consistent with the high Q-factor seen in Figure 4.3b, indicating that the related panel resonance colours the spatial response.

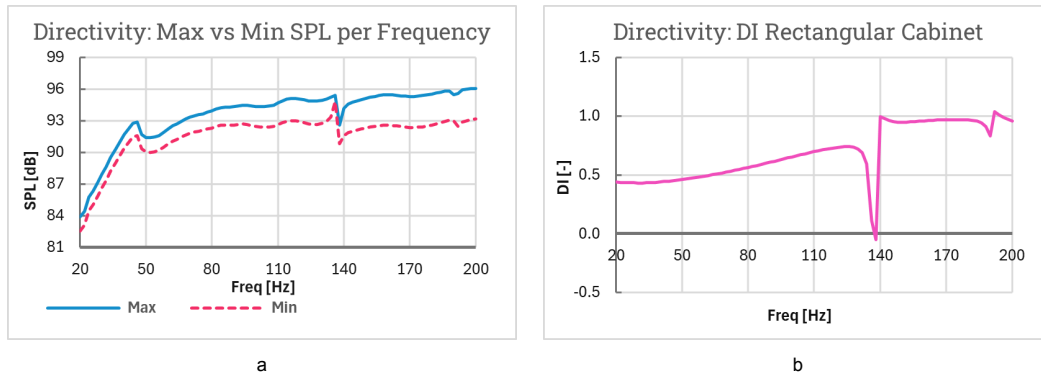


Figure 4.4: SPL response of the rectangular cabinet showing maximum and minimum SPL values per frequency for each evaluation angle (a) and the corresponding DI response (b).

The behaviour of the cabinet around the anomaly is evaluated using the mode shapes and a polar plot. Figure 4.5a and b show the mode shape for 140 Hz and 142 Hz, where it can be seen that the side panel of the cabinet shows a global resonance by deforming from concave to convex. The polar plot seen in Figure 4.5c shows the SPL over the -90° to $+90^\circ$ azimuth angle evaluation range and confirms that the off-axis SPL is indeed higher than the on-axis SPL. These findings confirm that resonance of the side panel exerts off-axis pressure and causes the observed DI anomaly.




Figure 4.5: Analysis of the DI anomaly around 140 Hz through mode shape visualisation at 140 Hz (a) and 142 Hz (b) and a polar plot of the SPL around the azimuth angle (c)

Figure 4.6 presents the full range directivity data in heat map format, a standard visualisation in audio engineering to show spatial SPL distribution. It can be seen that the high directivity at 140 Hz also shows up in this visualisation, here as a vertical band that shows lower SPL on-axis and higher SPL further away from the reference axis.




Figure 4.6: Heat map representation of the directivity index response of the rectangular cabinet.

4.3.1. Using Directivity in BEM Shape Optimisation

A gradient-based shape optimisation strategy requires that sensitivities are determined for evaluating objective functions. This is rudimentary for many physics interfaces, but not for the BEM interface. BEM only solves for boundary conditions during a solve and any metrics derived from the radiation conditions are obtained through post-processing. Within the current SO framework in COMSOL, *no post-processed quantities can be used as objective functions*. Furthermore, *any evaluation points or boundaries may not be physically disjointed from the geometry over which the SO is being performed*.

These prerequisites make that the SPL cannot be evaluated over a near-field point or surface in free space during shape optimisation. As such, the DI can also not be obtained directly for use as objective function. A workaround has been found for this in the form of geometry extensions. These extensions are included in the solid mechanics node and meshing strategy such that sensitivities may be determined, but are otherwise made structurally and acoustically insignificant by keeping them small and not allowing displacements nor deformations in the structure. Figure 4.7 shows a CAD representation of what this setup may look like.




Figure 4.7: Visualisation of domain setup required to extract adjoint sensitivities for use in shape optimisation targeting DI improvements.

This workaround is confirmed to provide a workable solution for including DI as an objective function within shape optimisation. Due to the low sensitivity of DI to panel deformations akin to shape optimisation, this method is omitted in the final optimisation framework, opting instead for parametric baffle edge optimisation.

4.4. Sensitivity Considerations for Optimisation

Sensitivities of the Q-factor and directivity are determined prior to optimisation to evaluate how each metric responds to specific geometric changes. Doing so enables more effective shape optimisation, improving convergence rate and solver robustness and efficiency. The sensitivity analysis aids in defining the objective functions, what geometric features are most responsive to shape optimisation, and what order-of-magnitude should be targeted considering the displacements and feature sizes.

The performance metric versus geometric change behaviour found through sensitivity analysis also reveals robustness trends, which are essential for setting up optimisation studies that converge and do not fall into local minima or that are prone to geometric issues such as mesh inversion [13].

First, tests are performed and presented in which the Q-factor and directivity are evaluated for geometric alterations of three kinds: *Baffle edge filleting* (radius 0-60 mm, Figure 4.8a), *panel thickness variations* (4-20 mm), and *panel curvature variations* (bulge depth 0-20 mm, Figure 4.8b). Next, sensitivity coefficients are formulated and evaluated for each performance metric. Lastly, the coefficients are compared and discussed within the context of optimisation, concluding with a final justification of the chosen objective functions and optimisation strategy presented in Chapter 6.




Figure 4.8: Geometric variations of the cabinet for testing parameter sensitivities, showing variation of baffle edge fillet from 0 to 60 mm (a) and panel bulge from 0 to 20 mm depth (b).

To guide the discussion of the results, varying illustrative comparisons are made. The sums of the DI and the Q-factor over 20-200 Hz serve as singular scalar values that are trivial to compare analytically between iterations, but omit details on performance trends and outliers. The effect of baffle edge variation on the DI and Q-factor is discussed in more detail to demonstrate orthogonality between the metrics. Table 4.1 provides more details to complement the discussion.

4.4.1. Geometric Parameter Sensitivity Study

The baffle edge geometry is varied parametrically by applying a fillet of increasing radius, from 0 to 60 mm in 5 mm increments. Figure 4.9a shows the DI response for each variation around the 140 Hz anomaly while highlighting the first and last variation (0 and 60 mm). It can be seen that the 0 mm variation exhibits a sharp dip. This dip persists up to a fillet radius of 45 mm, demonstrating that diffraction effects can be mitigated with increasing baffle edge radius. The baffle edge radius also alters the stiffness of the side panel of the cabinet, likewise contributing to altered directivity.

The Q-factor is also altered by said change of panel stiffness, which can be seen in the Q-factor response for the 0 and 60 mm variation in Figure 4.9b. The Q-factor decreases for higher frequencies but can be seen to decrease from **3.36** to **1.73** around 140 Hz, indicating diminished panel resonances as a result of altered panel stiffness.

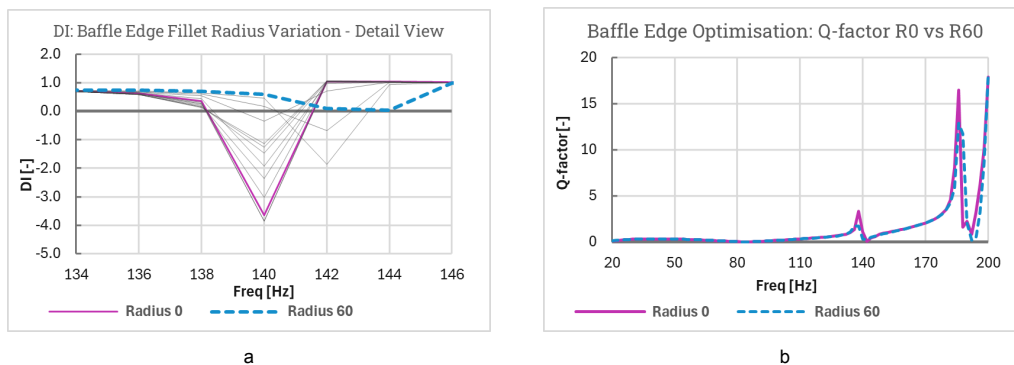


Figure 4.9: DI response for the rectangular cabinet with fillet radius 0 and 60 mm (a) showing increased uniformity and the corresponding Q-factor response (b) showing altered panel stiffness for either variation.

Panel thickness is varied from 8 to 18 mm in 2 mm increments, which is representative of the thicknesses often used in the LFAM process [31]. The Q-factor is assessed and summed over 20-200 Hz, shown in Figure 4.10. It can be seen that there is no direct correlation between the Q-factor and panel thickness, a variance of **46%** is observed between maximum and minimum value. Global panel stiffness increases with panel thickness, but such uniform changes appear insufficient for significant Q-factor improvement. This motivates the exploration of spatial variations in stiffness distribution through shape optimisation.

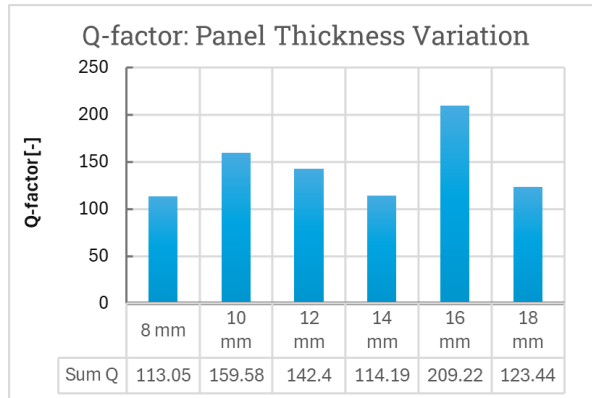


Figure 4.10: Summed Q over 20-200 Hz for several cabinet panel thickness variations between 8-18 mm, varied in 2 mm increments, showing no direct correlation between panel thickness and the Q-factor

Panel curvature sensitivity is evaluated by introducing a bulge to the cabinet panels with a depth of 2 to 20 mm in increments of 2 mm. Figure 4.11a shows the results of this variation for the summed DI and Figure 4.11b shows the summed Q-factor. Both evaluations are made to assess the interconnected nature of these performance metrics. The directivity shows invariance to the panel curvature, with summed DI over 20-200 varying by **19.7%** when comparing maximum and minimum value. The summed Q-factor instead shows large variance, with a **78.4%** variance between maximum and minimum value. This is also a larger variance than seen with varying panel thickness. This result confirms that the Q-factor is more sensitive to spatial variations in stiffness that are able to detune resonance modes. It is concluded that shape perturbations in an order-of-magnitude that can be reproduced using LFAM meaningfully alter modal damping.

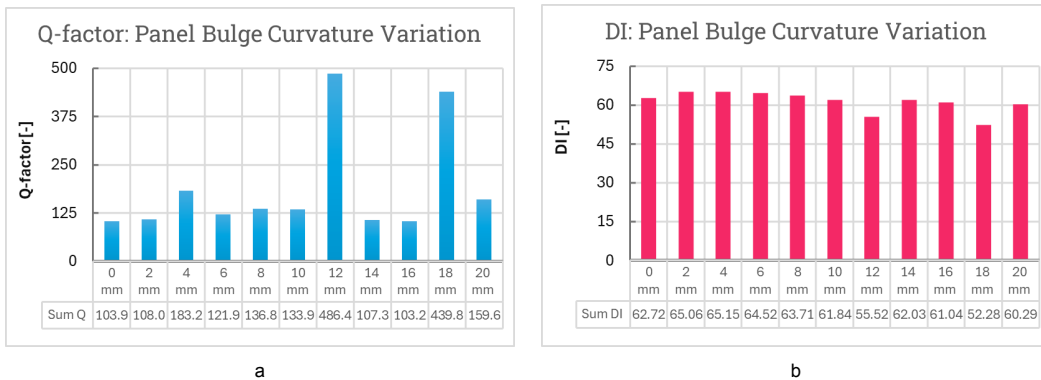


Figure 4.11: Summed Q (a) and DI (b) values for panel bulge curvature between 2-20 mm, showing a high sensitivity of the Q-factor and invariance of the DI to panel curvature.

4.4.2. Comparison of Sensitivity Coefficients

Several sensitivity coefficients are formulated and analysed using the data obtained from the previous studies. Since the primary focus is with the effect of geometry changes on the performance metrics, a general definition of sensitivity can be formulated as $S = \frac{\Delta \text{Objective}}{\Delta \text{Geometry}}$ where Δ represents the difference between minimum and maximum value. For example $S_{DI} = \frac{\Delta DI}{\Delta r_{fillet}}$, where r_{fillet} is the baffle edge fillet radius, quantifies directivity sensitivity to baffle edge radius (change/mm). Table 4.1 presents computed sensitivities for each combination of metric and parameter by computing both the relative variation and the absolute sensitivity.

| Sensitivity Coefficient | Geometry Variation | Minimum Value [-] | Maximum Value [-] | Relative Variation % | Absolute Sensitivity [$\frac{\Delta Obj}{mm}$] |
|-------------------------|--------------------|-----------------------|-----------------------|-------------------------|---|
| Summed DI ² | Fillet Radius | 45.9 | 64.5 | 40.4 | 0.31 |
| Summed DI ² | Panel Thickness | 54.5 | 101 | 186 | 4.69 |
| Summed DI ² | Panel Curvature | 52.3 | 147 | 281 | 5.27 |
| max DI | Fillet Radius | 1.01 | 1.06 | 4.95 | 8.33×10^{-4} |
| max DI | Panel Thickness | 1.10 | 2.56 | 233 | 0.15 |
| max DI | Panel Curvature | 1.06 | 1.54 | 145 | 2.65×10^{-2} |
| DI variance (log) | Fillet Radius | 1.85×10^{-2} | 4.53×10^{-2} | 145 | 4.47×10^{-4} |
| DI variance (log) | Panel Thickness | 0.02 | 0.05 | 318 | 3.67×10^{-3} |
| DI variance (log) | Panel Curvature | 0.02 | 0.08 | 457 | 3.27×10^{-3} |
| Q-target (p=2) | Fillet Radius | 27.0 | 243 | 901.7 | 3.60 |
| Q-target (p=2) | Panel Thickness | 21.2 | 80.1 | 378 | 5.89 |
| Q-target (p=2) | Panel Curvature | 20.3 | 344 | 1.69×10^3 | 18.0 |
| Summed Q ² | Fillet Radius | 727 | 5.91×10^4 | 8.13×10^3 | 973 |
| Summed Q ² | Panel Thickness | 448 | 6.41×10^3 | 1.43×10^3 | 596 |
| Summed Q ² | Panel Curvature | 413 | 1.19×10^5 | 2.87×10^4 | 6.56×10^3 |
| max Q | Fillet Radius | 15.0 | 241 | 1.61×10^3 | 3.77 |
| max Q | Panel Thickness | 11.0 | 60.3 | 548 | 4.93 |
| max Q | Panel Curvature | 9.88 | 344 | 3.48×10^3 | 18.6 |

Table 4.1: Comparison of sensitivity function formulations for shape optimization.

The relative variation is used for comparison as it inherently accounts for differences in scale between design variables, offering a fairer comparison of their influence on the objective function. The data show that summed DI² is most sensitive to panel curvature (281%) while less responsive to fillet changes (40.4%). However, note that the minimum attained value is lowest for varying fillet radius. Max DI and DI variance show varying sensitivity to geometry variations, relative sensitivity ranging between 4.95% and 457%. The higher sensitivity to panel changes stems from a relatively large negative impact on the DI and cabinet uniformity. This conclusion highlights that the sensitivity does not dictate what geometry variation nets the biggest *improvements*, the large change in objective value that increases sensitivity can be both positive and negative.

The Q-factor shows moderate to high sensitivity to any geometry variation. Summed Q² shows by far the biggest relative variations (1.43×10^3 to 2.87×10^4 %) and is clearly most affected by panel curvature (relative sensitivity 2.87×10^4 %). Q-target and Max Q show moderate sensitivity, high enough to guide optimisation. This finding confirms that panel curvature variations have a big impact on the Q-factor. The discussed results confirm that directivity and the Q-factor are governed by different physical mechanisms. The directivity is best improved by minimising edge diffraction effects and the Q-factor is lowered most effectively by spatial variations in panel stiffness.

4.4.3. Implications for Optimisation Strategy

The sensitivity analysis reveals a distinct difference in the approach required to optimise directivity and the Q-factor. Directivity *improves* most with local edge optimisation, which cannot be realised through the local perturbations of shape optimisation in COMSOL. More specifically, shape optimisation cannot transform a sharp edge into a filleted one and alter its shape, which is acknowledged as a constraint imposed by COMSOL's SO implementation.

The Q-factor can however be improved through shape optimisation perturbations in mm to cm scale. The sensitivity magnitudes presented provide guidance on required displacement budgets and targeted feature sizes, showing that an order-of-magnitude of millimetres to centimetres can be targeted for effective Q-factor reduction while adhering to the characteristics of LFAM.

The implication of these conclusions is a hybrid optimisation approach. The directivity is first optimised through parameter variation and the resulting geometry is subsequently optimised through shape optimisation to lower cabinet resonances by minimisation of the Q-factor. The sequential nature of this setup requires extra care to ensure that the optimisation of the Q-factor does not significantly worsen the previously optimised directivity. However, Table 4.1 shows that the directivity is minimally sensitive to panel deformations, suggesting that parametric DI improvements will be preserved during the subsequent Q-factor optimisation.

Objective Functions

The directivity objective function is formulated as the **minimisation of DI^2** . DI^2 is used since it pronounces sharp peaks, enabling the penalisation of large peaks so as to promote a more smooth response while also minimising directivity as a whole.

The Q-factor objective function is formulated as **minimisation of Q^2** , with the Q-factor extracted through Equation 4.2. The same motivation for taking the square of the metric applies, promoting a smooth response through penalisation of large peaks.

5

Experimental Validation of the Simulation Model

To establish confidence in the simulation model presented in Chapter 3 prior to its use in optimisation, experimental validation is conducted to assess the predictive accuracy of the model. This chapter presents measurements conducted under anechoic conditions of a prototype subwoofer fabricated using LFAM, the results of which are compared to the predicted behaviour following from simulations of a computational model of the same prototype. The prototype subwoofer, named the **Reflex Sub**, features a sophisticated double-walled structure, with the volume between the double walls forming ducts that function as integrated Helmholtz resonators. The comparison made here serves to test the model's capability beyond the simplified rectangular cabinet geometry and to analyse translation between simulation and real-world measurements.

Section 5.1 describes the Reflex Sub, providing details on its construction. Section 5.2 provides details on the measurement procedure and data acquisition. Section 5.3 presents the results following the measurements, simulation, and post-processing of data, used to formulate the conclusions found in Section 5.4 regarding model credibility and suitability for the subsequent optimisation.

5.1. Description of the Prototype Cabinet

The Reflex Sub is a large front-firing subwoofer with outer dimensions of $608 \times 608 \times 520$ mm ($W \times H \times D$), featuring an 18-inch 18NW100 4Ω driver from B&C Speakers [32]. Internal wrapped ducts serve as an air cushion that dampen vibrations from the inner shell being propagated towards the outer shell.

This subwoofer targets reproduction between 20-150 Hz with port tuning at approximately 34 Hz. The main cabinet shell is printed out of the C-250CF polycarbonate material presented in Section 3.4 using a strategy that deposits 2 parallel lines for each wall, resulting in a 2×6 mm configuration to obtain walls that are 12 mm thick in total. The front and rear panel are constructed from aluminium, 10 mm and 6 mm thick respectively. These panels are mounted using evenly-spaced mechanical fasteners around the perimeters. Figure 5.1a shows the prototype Reflex Sub as tested in the Anechoic chamber.




Figure 5.1: The prototype Reflex Sub in the Anechoic at the TU Delft (a) and a render of its CAD counterpart (b).

The CAD model of the cabinet is imported into SolidWorks for initial clean-up, where negligible features are removed to ensure a clean domain setup. This digital representation can be seen in Figure 5.1b. The simplified model is then imported into COMSOL for further processing, where the clean-up wizard is applied to remove remaining details <0.1 mm. The shell is modelled according to the method as described in Chapter 3. The C-250CF transversely orthotropic material is applied per the printing direction (aligned front-to-back as seen in Figure 5.1b). The aluminium panels are modelled to be isotropic using aluminium alloy (EN AW 1050A). The panels are modelled to be rigidly connected to the shell through continuous bonded surfaces, rather than using discrete points representing each bolt.

The critical material properties of C-250CF are $E_1 = 11$ GPa and $\eta = 0.015$, while the aluminium alloy is modelled using $E = 69$ GPa and with damping $\eta = 0.001$ [33]. The Lumped Mechanical System is governed by the TS parameters provided by B&C speakers, namely $C_{ms} = 0.1153$ mm/N, $M_{ms} = 0.244$ kg, $Q_{ms} = 6.2$, and $F_s = 30$ Hz.

These simplifications are justified by their minor influence on 20-200 Hz structural-acoustic coupling, the primary validation focus. A summary of the changes made within geometry and simulation modelling compared to the physical prototype is as follows:

- Omission of fasteners, replaced by continuous bonded surfaces
- Omission of integrated electronics (amplifier, digital signal processing, etc.)
- Removal of small geometric features (<0.1 mm)
- Idealisation of digital geometry versus printed geometry, e.g. the omission of printing defects and layer lines resulting in the rippled outer surface
- Simplification of driver modelling using the LMS

5.2. Experimental Setup and Measurement Procedure

The measurements were conducted in the Anechoic chamber of Delft University of Technology, which has dimensions of $8 \times 8 \times 8$ m and minimises reflections, providing near free-field conditions. The subwoofer was positioned on the suspended wire mesh with a thick piece of foam separating the subwoofer and the mesh. The subwoofer was placed slightly off centre in the room to have enough space on the baffle side to perform the measurements. The background noise floor was not formally quantified but also not detected during measurements.

The microphone was positioned 1 m in front of the driver. The microphone was not put in line with the centre of the driver but rather laid down on the mesh. This was deemed not critical at the time since any point in space could be extracted from the simulation data, but does deviate from standard practice and is acknowledged as a limitation. Directivity measurements were performed by positioning the microphone around the azimuth angle at 45° increments over $0-180^\circ$. A log swept sine at different **Decibels relative to Full Scale** [dBFS] levels was used as excitation signal, enabling amplitude control. Figure 5.2a shows the setup within the Anechoic chamber.



Figure 5.2: Setup of the Reflex Sub within the Anechoic chamber (a) and the IMU mounted to the rear panel of the subwoofer (b).

The measurements were repeated three times to average out measurement uncertainties. Besides SPL measurements, vibration measurements were also performed. A small and lightweight microcontroller (nRF52840 Sense) with integrated LSM6DS3TR-C 6-axis **Inertial Measurement Unit [IMU]**, was utilised. The controller was mounted to the rear, top, and both side panels using Museum wax to minimise coupling damping. Figure 5.2b shows the IMU being mounted to the rear panel of the subwoofer. The IMU data was recorded at 500 Hz sampling rate normal to each panel surface, capturing out-of-plane vibrations. Per the Nyquist sampling criterion [34], the 500 Hz sampling rate allows vibration measurements up to 250 Hz. The excitation sweep was applied at 20-800 Hz to enable inspection of behaviour at higher frequencies.

The recorded vibration data provides a time-domain spectrum showing acceleration [$\frac{m}{s^2}$] versus time [s]. Such a spectrum is shown in Figure 5.3a for illustrative purposes. The time-domain data is processed using MATLAB through a **Short-Time Fourier Transform [STFT]** using a Hann window of length 512 to obtain the frequency-domain response, shown in Figure 5.3b. The short-time transform is used in combination with tight time windowing to avoid aliasing of the higher-frequency response into the 20-200 Hz range of interest.

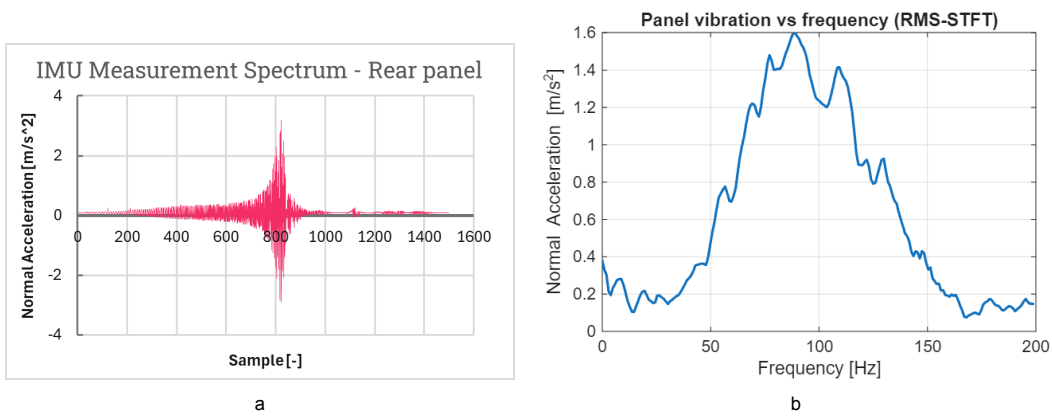


Figure 5.3: Time-domain acceleration spectrum obtained through vibration measurements (a) and the corresponding frequency-domain acceleration spectrum obtained through STFT (b).

5.3. Simulation-to-Measurement Comparison

Before a full simulation of the Reflex Sub is conducted, a verification run is performed for the driver modelling. As mentioned in Section 3.3, the assumption of piston motion does not hold for the 18" driver, thus a correction factor is applied. Implementation of said compensation is achieved by defining the diaphragm surface to move at a different velocity at its edge compared to its centre. The corrected normal velocity is then defined by Equation 5.1, where v_{corr} is the corrected normal velocity, $\alpha_f(f)$ is the correction factor as a function of frequency, r is the distance from the centre of the diaphragm to its edge, and S_d is the area of the diaphragm. $\alpha_f(f)$ is a monotonically increasing function that increases from 0 at the breakup frequency of 100 Hz to α_{max} at 200 Hz. Figure 5.4a shows the distribution of velocity on the diaphragm above the breakup frequency and Figure 5.4b shows the SPL response obtained with the original simulated driver excitation and the corrected excitation.

$$v_{corr} = v_n \left(1 - \alpha_f \frac{r^2}{S_d} \right) \quad (5.1)$$

With the correction factor applied to the driver, a baseline simulation is performed of the Reflex Sub. All simulations and measurements are considered for the same frequency range of interest: **20-200 Hz**. An initial SPL response is obtained and shifted along the y-axis to line up visually with the measurements for *normalised* trend comparison. For this comparison, three SPL measurements performed with identical excitation signals are averaged. Both the simulated and average measured response can be seen in Figure 5.5.

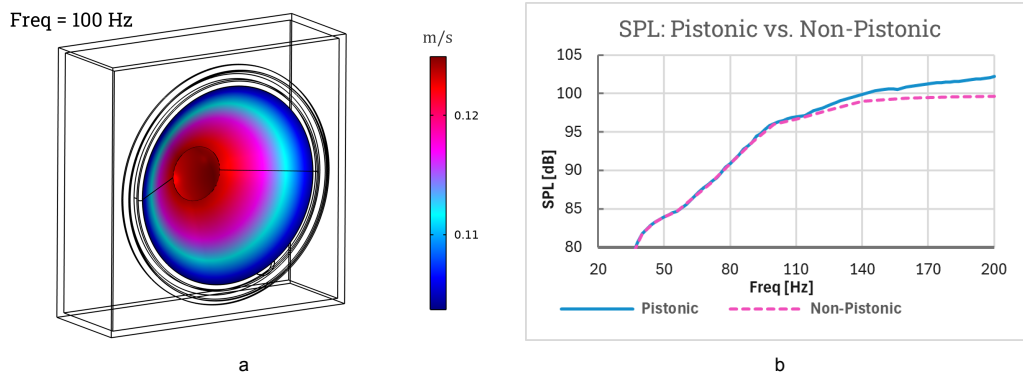


Figure 5.4: Non-uniform velocity distribution applied to the 18NW100 driver diaphragm (a) and comparison of the original and corrected SPL response obtained through non-pistonic correction (b).

A first comparison between the SPL responses shows a discrepancy in the 20-80 Hz range. The simulated response shows a peak in SPL at **42 Hz** of **88.6 dB**, followed by a dip at **58 Hz** of **29.9 dB** before gradually tapering back to a flat response. The measurements, in contrast, show a more continuous response without this peak and dip. This discrepancy is identified to be a typical Helmholtz resonance over-prediction in FEM/BEM simulations and is a known limitation in modelling flow effects within ported ducts [35]. Literature shows that accurate port modelling requires other methods such as CFD [25], which is beyond the scope of the present research.

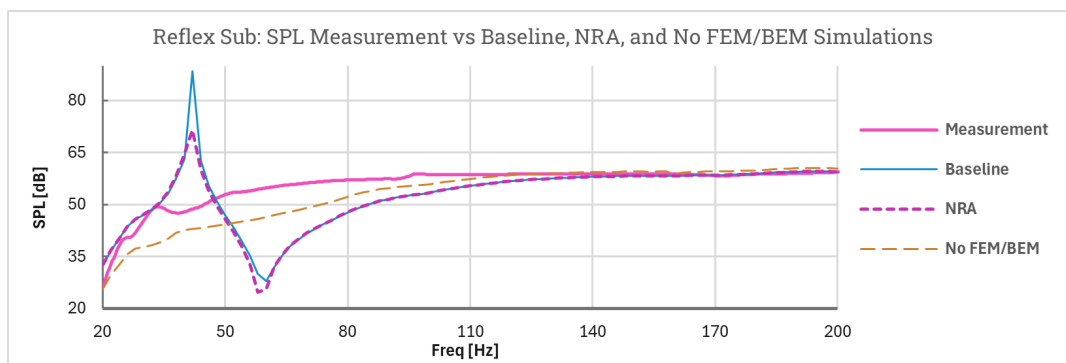


Figure 5.5: Comparison of measured and simulated SPL response of the Reflex Sub, with the measured response obtained by averaging 3 identical measurements and the simulated response shown for Baseline, including NRA, and excluding the Helmholtz resonators (isolated).

Above 80 Hz, the simulated and measured response show good agreement, showing a **Mean Absolute Error [MAE]** of **2.12 dB** between 80-200 Hz compared to a MAE of **11.7 dB** between 20-80 Hz. This confirms that the observed discrepancy is contained to the lower bandwidth. The material property characterisation, application, and structural boundary conditions appear sufficient for accurate simulation results. A modification is made to the interior acoustic domain by utilising **Narrow Region Acoustics [NRA]** to include some viscous damping, which lightly dampens the Helmholtz resonance. The effective hydraulic diameter is calculated to be 0.04 m for the integrated ducts. The simulated results including NRA can be seen in Figure 5.5, where the Helmholtz resonance is attenuated by **17.0 dB** but not fully removed. The MAE between 20-80 Hz is **11.54 dB**, and between 80-200 Hz **2.11 dB**.

The simulation is re-run while omitting the FEM/BEM coupling at the port openings by disabling this multiphysics coupling in COMSOL to omit the Helmholtz resonance. Doing so eliminates the resonator contribution to the response and brings the simulated and measurements response closer together still. All subsequent response analyses demonstrate the simulated response with and without the Helmholtz resonator, the latter denoted as 'No FEM/BEM'. Figure 5.5 also includes the No FEM/BEM response. The MAE between measurement and simulation between 20-80 Hz is reduced by **40.5%** to **6.94 dB**, and between 80-200 Hz by **38.2%** to **1.31 dB**. The overall MAE between 20-200 Hz is **3.16 dB**.

The structural response is compared between simulation and measurement to identify trends, not absolute values. Trend agreement is sufficient for understanding modal cabinet behaviour at a level suitable for effective shape optimisation. The primary goal is to identify the resonance frequencies of the cabinet in order to select effective frequency bands for Q-factor optimisation. The frequency-domain data from the measurements are averaged per panel and a smoothed trend line is plotted for clarity. Here, the data from the *right-side panel* of the subwoofer is compared since most measurements were performed on said panel. The simulated response is obtained by taking the point on the digital model of the cabinet where the IMU was placed in real-life and taking the effective normal acceleration (`solid.accY`). The simulation excluding the Helmholtz resonators is considered to forego introducing known uncertainty.

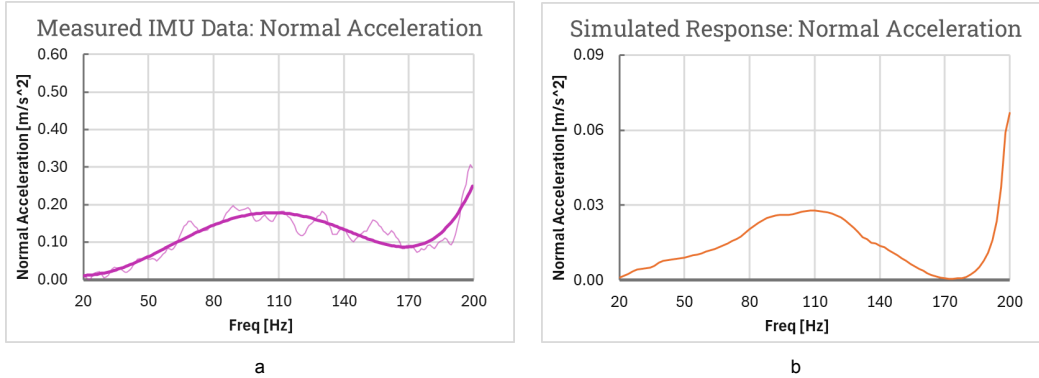


Figure 5.6: Normal acceleration response of the Reflex Sub panel reference point used in both the measurements (a) and simulation (b), showing good agreement in trend behaviour with a comparable peak at 110 Hz, a subsequent valley, and a rise near 200 Hz.

Figure 5.6a shows the measured and Figure 5.6b the simulated response, both of which show a similar trend with a peak in acceleration around 110 Hz, a valley between 110-190 Hz, and a rise in acceleration near 200 Hz. The amplitude of the measured response shows a non-zero response between 110-190 Hz whereas the simulation does. This behaviour is attributed to decay behaviour of the cabinet, showing ringing as the swept swine moves up in frequency. This response is not observed in the simulation since it assumes a steady-state harmonic response, omitting transient behaviour. The acceleration amplitudes are a factor **6.5** apart between measurement and simulation. This difference is attributed to the driving voltage not being calibrated. In the simulation, a RMS voltage of 2.83 V is used. During measurements, the driving voltage was likely much higher, manifesting in the observed discrepancy.

The directivity response is compared in a similar manner as the SPL, now for 0° , 45° , 90° , 135° , and 180° around the azimuth angle. The DI is extracted according to Equation 4.3. Figure 5.7 shows the simulated and measured DI response, which show decent agreement, with a MAE of **0.049** between 20-200 Hz. The difference in shape of the two responses is hypothesised to be the result of measurement uncertainties. Both curves show a low DI at <0.4 , from which it can be concluded that the response has good spatial uniformity.

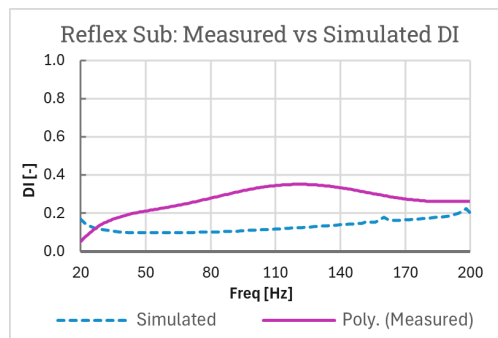


Figure 5.7: DI responses for the Reflex Sub as measured and simulated, showing decent agreement between both curves, with a MAE of 0.049.

Concluding, the simulation captures the dominant trend in structural behaviour with sufficient accuracy to guide the optimisation framework. The SPL and directivity response show larger deviations due to improper modelling of the Helmholtz resonator. Choosing to omit the resonator excludes a known limitation from the simulation and improves robustness and predictability, but under-represents its added benefit in real-life scenarios. The implications of these results and conclusions on model credibility are further discussed in Section 5.4, including several potential sources of the deviations observed between simulation and measurements.

5.4. Conclusions on Model Credibility

The experimental validation presented in the previous sections indicates that the FEM/BEM model presented in Chapter 3 produces reliable predictions of modal cabinet behaviour in the 80-200 Hz band, suitable for guiding optimisation parameters. The simulation results in the 20-80 Hz band prove less accurate when the Helmholtz resonance is included, but still provide information on trends in behaviour that are likewise suitable for guiding optimisation parameters. Omitting the BEM coupling at port openings removes the Helmholtz resonance, bringing the measurements and simulations in closer agreement in the 20-80 Hz band.

With NRA applied to dampen contribution of the Helmholtz resonance, the SPL response in the 80-200 Hz band exhibits a MAE of **2.11 dB**. Structural resonance frequencies as observed through the panel accelerations align within 1-2%. Directivity analysis shows a MAE of **0.049**. These results show sufficient accuracy in trend behaviour to guide the optimisation process, primarily in the 80-200 Hz band, with reduced accuracy in the 20-80 Hz band. Absolute value prediction is proven to be limited. SPL evaluation requires normalisation and the acceleration amplitudes are a factor 6.5 apart, indicating quantitative limitations despite accurate trend predictions.

The low-frequency Helmholtz resonance is acknowledged as an inherent limitation to the FEM/BEM method. As such, the port-cavity coupling responsible for this response is typically omitted in the present research, being included here to quantify the discrepancies generated by its inclusion. Within the optimisation framework, the port-cavity coupling is disabled, eliminating port modelling inaccuracies that can disrupt guidance of the optimisation due to the Helmholtz resonance dominating over panel resonances in the 20-80 Hz band. This implies that the optimisation framework is developed for use with **sealed** loudspeaker cabinets.

Potential sources of deviations between the simulation and measurements are identified:

- Material property uncertainty, such as the loss factor and Young's Modulus depending on LFAM process variations.
- Boundary condition idealisations, such as the omission of fasteners in favour of continuous bonded surfaces, omitting localised stiffness and damping.
- Geometric tolerances and defects in the LFAM process, such as dimensional inaccuracies and layer-to-layer variations.
- Driver modelling assumptions, such as using the TS parameters and a Lumped Mechanical System, omit non-linearities and complex behaviour like cone breakup [36].

Despite these uncertainties, the simulation model is able to capture dominant trends in structural behaviour with sufficient accuracy to inform optimisation decisions. Resonance frequency prediction within 1-2% facilitates reliable identification of Q-factor peaks to be minimised. The presented validation suggests physical consistency across the validation metrics, establishing confidence that geometries obtained through optimisation will manifest in physical prototypes that indeed improve real-world performance. Chapter 6 presents the performance of optimised geometry against a baseline rather than absolute target specifications.

6

Optimisation Framework and Results

This chapter presents the implementation and results of the optimisation framework developed to minimise loudspeaker cabinet resonances through geometric modifications at scales compatible with the LFAM process. Exploration of parameter sensitivities in Section 4.4 motivates the use of a hybrid sequential optimisation strategy, starting with a parametric optimisation of the baffle edge to reduce diffraction effects, followed by shape optimisation targeted at Q-factor reduction. Objectives and constraints are discussed in Section 6.1. The parametric optimisation strategy is discussed in Section 6.2.

Two case studies are presented, the rectangular reference cabinet and the prototype Reflex Sub presented in Chapter 5. The optimisation framework is developed using the reference cabinet and later adapted to the Reflex Sub to evaluate scalability and transferability. The optimisation framework is built on gradient-based methods within COMSOL's built-in shape optimisation environment, described in detail in Section 6.3. The results of successful optimisation studies performed with the presented framework are shown and discussed in Section 6.4, along with physical mechanism analyses and robustness studies, concluding with a comparison to an alternative optimisation method.

6.1. Optimisation Objectives and Constraints

6.1.1. Objective Function

The primary objective function for minimising cabinet resonances is presented in Section 4.2 to be the Q-factor, described by Equation 4.2. A restriction of the shape optimisation environment in COMSOL is that any objective function when determined over a range, such as a *frequency* range, is automatically aggregated across all frequencies. Thus, any geometric perturbation is evaluated over the entire band of interest, which does not allow for evaluation of frequency-specific metrics such as response smoothness or spectral variance. Direct terms that define the smoothness of a response, such as the Q-factor response, are desirable for obtaining non-jarring behaviour but cannot be implemented directly due to the aggregation. However, proper scaling of the objective function indirectly promotes smoothness by emphasizing the attenuation of outliers, as is shown in Section 6.4. The scaled Q-factor formulation used in the present research is Q^2 , which serves as the primary objective function. The objective function is described by Equation 6.1.

$$\min \sum_{\omega_i \in [180, 200] \text{ Hz}} Q^2(\omega_i) \quad (6.1)$$

6.1.2. Manufacturability Constraints

For geometric features to be physically transferable and admissible within the LFAM process, several constraints are considered. These constraints stem from the process characteristics, primarily the usable nozzle sizes and printing strategy. Typical nozzle sizes range between 2 and 18 mm, with the practical minimum feature size being 2 times the nozzle diameter, i.e. 4-36 mm [31]. The feature sizes explored in Section 4.4 fall within this range, thus a constraint is formulated for the minimum viable feature sizes. The reference cabinet has a 10 mm thick wall, thus a minimum feature size of **20 mm** is adopted as constraint.

The printing strategy also requires panel thicknesses to be constant and the geometry to form a closed contour when sliced transversely to the printing direction. Figure 6.1a shows a cross-section view of an LFAM 3D print illustrating such a closed contour, which is crucial to ensure continuous deposition and avoiding print defects caused by *discontinuous* deposition. The contour and thickness constraint are handled through mesh smoothing and translation of displacements applied within the optimisation framework, discussed in Section 6.2.3. In the present research, **Yeoh smoothing** is employed, which is based on the *Yeoh hyperelastic material model* [37]. This is a highly non-linear method that allows for the largest deformations out of all options available in COMSOL with the least risk of mesh inversion. Alternative smoothing options, such as Laplacian and Winslow, caused mesh inversion during preliminary optimisation runs. Figure 6.1b shows a cross-section view of a shape-optimised bracket from the COMSOL application library [38] with Yeoh smoothing and displacement translation applied. Here, it can be seen that panel thickness is well maintained, justifying its use in the present research.




Figure 6.1: Illustrative figures of continuous contour strategies used in LFAM, showing the cross-section of an LFAM print by Addit Audio (a) and the cross-section of an eigenfrequency-optimised bracket from the COMSOL application library showing thickness preservation (b).

6.1.3. Computational Limitations

Besides manufacturability, computational effort is considered to pose limitations, since shape optimisation remains inherently costly. The final cost is determined by choice of algorithm, resolution of the mesh, number of optimisation variables, solver tolerances, solver configuration, termination conditions, and the number of evaluated points. These topics are discussed in Section 6.3, except for the number of evaluated points. This count depends directly on the frequency range of interest and sampling interval. For development and exploration of the optimisation framework, baseline simulations are used to reveal high Q-factors, around which a range of interest is defined. This range of interest has a size of **20 Hz** and sampling interval of 2 Hz, which is smaller than the full range of interest of 20-200 Hz. Depending on the exact setup and hardware used, this size of range takes approximately 1 hour per iteration to solve, allowing for reasonably rapid iteration during framework development.

Preliminary shape optimisation trials performed with the Reflex Sub validated in Chapter 5 revealed <5% Q-factor improvement. This minimal improvement is attributable to the existing double-walled and braced structure, along with the curved outer geometry of the design. Spatial variations in stiffness are already attained, confirming that the shape optimisation framework provides largest improvement gains on baseline geometries that are not acoustically refined. Subsequent optimisation therefore focused on the rectangular reference cabinet as primary development platform. The Reflex Sub serves as validation of the framework's scalability across geometries and the underlying physical mechanisms discussed in Section 6.4.4.

6.2. Cabinet Parameterisation Strategy

The parameterisation method adopted in the present optimisation framework accommodates the hybrid sequential strategy discussed in Section 4.4.3. This strategy involves parametric optimisation of the DI through baffle edge geometry variations, followed by gradient-based optimisation through panel geometry perturbations. This strategy is motivated by the sensitivities established in Section 4.4, which

show that directivity responds primarily to edge diffraction, while the Q-factor is governed by spatial distribution of panel stiffness through curvature.

6.2.1. Parametric Baffle Edge Optimisation

The baffle edge geometry is parameterised using a constant-radius fillet. Olson [30] demonstrated that spherical cabinets demonstrate reduced edge diffraction artifacts compared to rectangular ones. The fillet radius is increased from 0-60 mm, which is the maximum radius admissible without distorting the driver mounting surfaces. The radius is increased in 5 mm increments and each configuration is simulated independently to extract DI performance per the method presented in Section 4.3 at 10° increments across 0-90° azimuth angle, yielding 10 evaluation points. Parametric baffle edge optimisation is followed by gradient-based panel shape optimisation.

6.2.2. Shape Optimisation Panel Parameterisation

The panels of the cabinet are parameterised through COMSOL's free-form deformation framework, which combines parametric representations of external boundaries with a moving-mesh approach. In the present research, the cabinet geometry is designed externally and imported to COMSOL, meaning that no native parametric representation is available and the deformations are enabled solely through mesh perturbations. The panels subject to optimisation are the rear, right-side, and top panel of the quarter-geometry model. The baffle panel is excluded to preserve the driver mounting surfaces, the filleted baffle edge is excluded to preserve the parametrically optimised geometry.

Panels that share an edge can potentially intersect when their geometry is perturbed, a failure mode known as **mesh inversion**. This failure mode was observed in preliminary trials when conflicting displacements at shared edges created incompatible geometry. This problem is mitigated by defining an explicit deformation region on each panel that is inset from the outer boundary, excluding edge and corner regions from the optimisation. An offset equal to the minimum feature size of 20 mm is used. These edge exclusion zones alter the base mesh, a refinement can be applied to maintain element quality at an average similar to the mesh without exclusions zones (here, 0.62). This refinement proves unnecessary for 20 mm exclusion zones. Figure 6.2 shows a schematic overview of what the edge exclusion zones look like on the rectangular reference cabinet.




Figure 6.2: Schematic overview of the edge exclusion zones utilised during shape optimisation to prevent mesh inversion failure mode

Another mechanism for preventing mesh inversion is implemented by limiting geometric perturbations to the normal direction of each panel. Maximum displacements are set to ± 20 mm, while the maximum displacements in directions tangent to each panel are limited to ± 0 mm, preventing the panels from sliding and causing edge collapse. The range of ± 20 mm is chosen based on the sensitivity analysis from Section 4.4 and the minimum viable feature size reproducible by LFAM, as discussed in Section 6.1.

6.2.3. Mesh Smoothing and Thickness Preservation

Preservation of panel thickness is achieved by prescribed displacement coupling, coupling the geometric perturbations of the outer surfaces to their respective inner surfaces. This is done in COMSOL using *General Extrusion Coupling* (`genext`). Prescribed deformation nodes are then used to apply the outer surface displacements (`material.dX/dY/dZ`) onto the respective inner surfaces by giving the displace-

ments an offset equal to panel thickness. This implementation is essential for allowing displacements larger than panel thickness while preserving mesh continuity and constant thickness. Additionally, Yeoh smoothing preserves element quality by propagating surface deformations through the volume mesh, maintaining mesh validity across large deformations.

The described prescribed deformation strategy implicitly applies only to deforming domains. Selecting the cabinet domain to be a free shape domain means that *only* the cabinet boundaries will move. To get the interior air domain (the cavity) to move with the cabinet, it has to be explicitly selected as a deforming domain in the `Deformed Geometry` interface. This is counter-intuitive. Since the panels share their inner boundaries with the cavity, it may be incorrectly presumed that the cavity inherently follows the cabinet deformation.

Preliminary tests showed that more extreme displacements (beyond the limit of the present research: >20 mm) can result in inner mesh elements crossing over one another, causing mesh defects. This is not a failure mode to the optimiser since the prescribed displacement coupling is considered a post-processing step that does not impact the optimised outer surfaces. A clean-up of the mesh in CAD can be performed to obtain a result suitable for LFAM.

6.3. Optimisation Algorithm Selection and Implementation

The shape optimisation environment in COMSOL offers gradient-based algorithms only (v6.4). The choice of algorithm is between **Method of Moving Asymptotes** [MMA], **Globally Convergent Method of Moving Asymptotes** [GCMMA], or **Interior Point Optimiser** [IPOPT]. All three algorithms use adjoint-based sensitivities for calculating gradient information needed to determine which geometric perturbations most effectively improve the chosen objective functions. The choice of algorithm depends on the problem being solved. MMA and GCMMA are better suited for handling a large number of control variables, such as in topology or shape optimisation. MMA adapts based on iteration history and is often considered more efficient [39], reaching convergence with relatively few objective function evaluations compared to IPOPT. Due to the highly coupled nature of the problem in the present research, MMA is the most suitable algorithm. GCMMA works in a similar way to MMA, but rejects any iterations that result in an adverse effect on the objective functions. Inclusion of suboptimal designs is considered beneficial for analysis of optimiser behaviour, further justifying the choice for MMA.

MMA allows the use of move limits, limiting the step size per iteration to a fraction of the maximum displacement budget set in the SO parameterisation. Early trials with the move limits at 0.2 and ≤ 10 iterations produced non-converged results with oscillating objective values, occasionally hitting ESE singularities, causing objective explosions (e.g. $\times 100$ spike). Conversely, a conservative move limit value of 0.02 leads to very small perturbations that require excessive iterations to use the full displacement budget, increasing the total runtime. Setting the move limits to 0.05 was found to strike a balance between cost and stability, enabling usage of the full displacement budget within approximately 15-20 iterations. 25 iterations are allocated instead of 15-20 to ensure convergence is found beyond the exhaustion of the displacement budget. Note that the optimiser does not stop perturbation of the geometry after the limits of ± 20 mm are initially reached. The geometry is only prevented from being perturbed *beyond* the limits, but can still be refined within the allocated range after 15-20 iterations.

The solver configuration employs segregated steps, splitting the computational load over three steps: the lumped mechanical system, the acoustic and structural variables, and the optimisation variables. The optimisation step is solved using a GMRES iterative solver with direct preconditioning through PARADISO. The other two steps are solved using a direct PARADISO solver with default tolerance settings. This configuration was found to be most efficient but is limited in its applicability due to associated high memory requirements. The simulation model is set up according to the method described in Chapter 3. This physics and meshing strategy, combined with the optimisation setup described in this chapter, leads to a framework wherein 11 frequencies are evaluated for $\sim 201k$ DOFs, including ~ 717 optimisation variables. Computational cost of each iteration for this particular case averages 1.28 hours on desktop hardware with a 6-core AMD Ryzen™ 5 3600 processor, 32 GB of LPDDR4 RAM, and an NVIDIA Geforce GTX™ 1080 GPU. A 25-iteration convergence run with ± 20 mm displacement exploration averages 1.6 days of total runtime. This implies the use of an iteration-count termination criterion, as opposed to formal MMA convergence. Section 6.4.3 discusses the termination criteria in more detail.

6.4. Optimisation Results and Interpretation

This section presents the results obtained using the optimisation framework on the reference cabinet. Section 6.4.1 provides preliminary insights on the evaluation of improvements being relative in nature. Section 6.4.2 presents baseline performance of the cabinet and the outcomes of parametric baffle edge optimisation. Section 6.4.3 presents convergence and the performance of the shape optimisation. The underlying physical mechanisms behind the shape optimisation are analysed in Section 6.4.4. The robustness of the shape optimisation results is evaluated in Section 6.4.5, and Section 6.4.7 presents an alternative optimisation method based on eigenfrequency analysis.

6.4.1. Preliminary Insights on Relative Improvement Evaluation

Section 3.7.1 established that Q-factor magnitudes exhibit mesh dependency. Convergence of the structural mesh is governed by resonance frequency stability, rather than absolute value prediction. Thus, improvements obtained through parametric and shape optimisation are assessed within consistent discretisation to forego mesh dependency. Performance comparisons are made through relative summed Q^2 reduction and energy redistribution, rather than prediction of absolute magnitudes, consistent with standard practice in gradient-based structural optimisation where comparative metrics are used for convergence assessment [40]. Chapter 5 presented validation of trend prediction accuracy, supporting reliability of comparative optimisation performance.

The base mesh of the cabinet is altered by the edge exclusion zones. Thus, the absolute magnitude of the Q-factor changes from baseline discretisation to that used in SO. The location of resonance frequencies and the trends in behaviour remain consistent with the baseline without edge exclusion zones. Clarity of comparison is improved by considering subsequent optimisation steps within consistent discretisation. This means that two meshes are employed, the first between baseline - parametrically optimised geometry and the second between parametrically optimised - shape optimised geometry. Between the two meshes, average element quality is preserved within 0.5%.

6.4.2. Baseline Performance and Parametric Optimisation

Baseline performance of the rectangular cabinet establishes reference values for evaluating optimisation improvements. High Q-factors guide the frequency range selection for optimisation. Baseline analysis reveals that the reference cabinet exhibits dominant Q-factor peaks at **138 Hz** ($Q = 3.36$), **186 Hz** ($Q = 16.45$), and **200 Hz** ($Q = 17.89$). The Q-factor peak at 200 Hz exhibits the largest Q-factor magnitude. The concentration of two peaks within 180-200 Hz allows for multi-peak attenuation and motivates its selection as range of interest. The baseline objective value, or **baseline summed Q^2** ($\sum_{180}^{200} Q^2$), is **849.7**.

Modal analysis identifies the Q-factor peaks as panel bending modes. Figure 6.3a, b, and c illustrate mode shapes for each respective frequency, revealing spatial distribution of the respective displacements, with maximum panel displacements occurring at 200 Hz on the rear panel. The other mode of interest, 186 Hz, shows maximum panel displacements at the top panel. Note that the deformations shown serve as visualisation and are not to relative scale.




Figure 6.3: Visualisation of baseline structural modes of the rectangular cabinet at 138 Hz (a), 186 Hz (b), and 200 Hz (c) (relative deformation scale not consistent).

Baseline directivity, expressed as the sum of DI^2 (emphasizing uniformity deviations) over the 20-200 Hz range ($\sum_{20}^{200} DI^2$), is **63.11** [-]. The results presented here are discussed in Section 4.4.1, including a

discussion on the anomaly observed around 140 Hz. Figure 6.4a shows the on-axis SPL, Figure 6.4b shows the DI^2 response of the baseline cabinet, and Figure 6.4c shows the accompanying heat map to visualise the directivity spatially. These figures are repeated from Section 4.4.1 to aid the reader.




Figure 6.4: Baseline directivity performance for the rectangular cabinet expressed through the on-axis SPL response (a), DI^2 response (b), and heat map (c) for 20-200 Hz.

These baseline metrics establish the targets for the optimisation; Reduction of the Q-factor in the 180-200 Hz range, while maintaining or reducing $\sum_{20}^{200} DI^2$. $\sum_{20}^{200} Q$ and $\sum_{20}^{200} Q^2$ are occasionally presented for illustrating broadband performance, where $\sum_{20}^{200} Q^2$ serves to illustrate the smoothness of the Q-factor response over the entire frequency range. All metrics can be found in Table 6.1. Parametric baffle edge optimisation addresses directivity first, followed by shape optimisation for Q-factor reduction, concluding with verification of directivity preservation.

Details of the baffle edge optimisation are presented in Section 4.4.1. Varying the fillet radius from 0-60 mm in 5 mm increments, 60 mm is identified as optimal, reducing $\sum_{20}^{200} DI^2$ from **63.11** to **45.91**, a 27.3% improvement. Figure 4.9a demonstrated that the anomaly at 140 Hz is diminished with increasing fillet radius. Figure 6.5a illustrates $\sum_{20}^{200} DI^2$ per fillet radius variation, showing a subtle downward trend with an outlier at 30 and 35 mm. Figure 6.5b shows the Q-factor response of the baseline and parametrically optimised design. The Q-factor is reduced from **16.6** to **12.9** (186 Hz) and increased from **17.9** to **18.2** (200 Hz) without frequency shift. The Q-factor peak at 138 Hz is reduced from **3.36** to **1.73**. $\sum_{180}^{200} Q^2$ is reduced from **849.7** (0 mm fillet) to **789.6** (60 mm fillet) for this particular discretisation. $\sum_{20}^{200} Q$ changes from **125.7** to **119.7**, a 4.8% improvement.

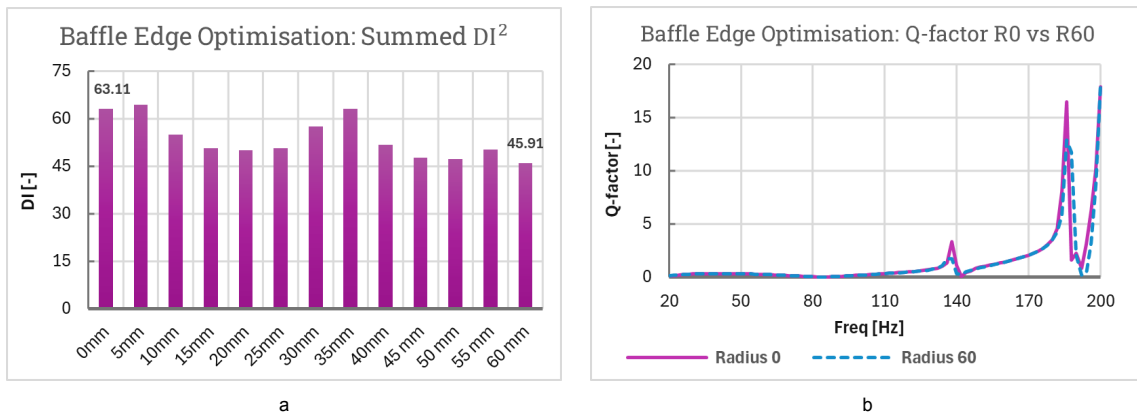


Figure 6.5: Bar graph showing summed DI^2 for each baffle edge fillet variation over 20-200 Hz (a) and the Q-factor response for the first and last iteration of the parametric optimisation (b) (repeated from Section 4.4.1 to aid the reader).

6.4.3. Shape Optimisation: Convergence and Performance

Preceding shape optimisation, the switch is made to the second discretisation to improve relative evaluation of the results. The parametrically optimised design is re-evaluated to obtain new baseline values: $\sum_{20}^{200} DI^2 = 45.72$, $\sum_{20}^{200} Q^2 = 1039.2$, and $\sum_{180}^{200} Q^2 = 974.4$.




Figure 6.6: Optimisation convergence indicators, showing the objective function value per iteration (a), the surface-integrated adjoint sensitivity magnitude per iteration (b), and the mean geometric change per iteration (c).

Shape optimisation applied to the 60 mm fillet radius cabinet reduces $\sum_{20}^{200} Q^2$ to **634.3**, a reduction of 39%. $\sum_{180}^{200} Q^2$ is reduced to **565.6**, a 41.9% reduction. $\sum_{20}^{200} DI^2$ becomes **45.56** and is preserved within 1%. Convergence is demonstrated through several indicators. Figure 6.6a shows that the objective function per iteration plateaus after 14 iterations, where the final 6 iterations exhibit a **coefficient of variation [CV] 3.04%**. Figure 6.6b illustrates the surface-integrated adjoint gradient magnitude¹ showing a CV of 8.16% over the last 6 iterations, indicating that the high-sensitivity design regions have been targeted effectively. Figure 6.6c shows the mean geometric change per iteration with a CV of 1.41% over the last 6 iterations, indicating decreasing and stabilising geometry perturbations. While the formal gradient tolerance of MMA ($<10^{-6}$) is not satisfied, the demonstrated metrics show practical convergence and additional iterations yield minimal improvements over added computational cost.

The mean Q-factor over 20-200 Hz is reduced from **1.38** to **1.21**, a 12.6% reduction. $\sum_{20}^{200} Q$ is reduced from **125.5** to **109.7**, likewise a 12.6% reduction. Figure 6.7 shows the Q-factor responses for the rectangular baseline, 60 mm fillet (mesh 2), and shape-optimised cabinet. Broadband smoothing of the response by the optimisation is seen. The peaks at 186 Hz and 200 Hz are attenuated, being reduced in magnitude from **17.87** to **6.57** and from **20.21** to **8.02**, respectively. Q-factor standard deviation [σ_Q] decreases from **3.09** to **2.35**, a reduction of 23.9%, demonstrating improvement of spectral uniformity. Smoothing of the response indicates modal dispersion of energy, discussed in detail in Section 6.4.4.




Figure 6.7: Q-factor responses between baseline rectangular cabinet, 60 mm fillet parametrically optimised geometry, and final shape optimised geometry.

Final geometric changes compared to iteration 1 range ± 3.5 mm, well below the displacement budget of ± 20 mm. The optimiser converged to a solution requiring only modest changes, a finding that aligns with the high sensitivity presented in Section 4.4. The final displacements are small but still within LFAM fabrication tolerances [41]. Largest changes are made to the side panel. The top and rear panel exhibit more moderate deformations. Figure 6.8a shows the final geometry with displacement colour map. Figure 6.8b presents a cross-sectional view and verifies preservation of panel thickness using prescribed displacement coupling. Mesh quality remains stable throughout, demonstrating an average mesh quality of **0.645** at the last iteration, confirming that element integrity is preserved over the deformations. The distribution of deformation does not correlate with iteration 1 adjoint sensitivities

¹ $\sqrt{\text{abs}(\text{adj}(u))^2 + \text{abs}(\text{adj}(v))^2 + \text{abs}(\text{adj}(w))^2}$ [m]

seen in Figure 6.9. Regions of high sensitivity, where geometric perturbations produce large objective improvements, should align with regions of most displacement. Early iterations indeed show this mechanism, confirming that MMA targets highly-sensitive areas most, exploiting gradient information. Subsequently, the optimiser targets areas more specifically, changing stiffness locally.

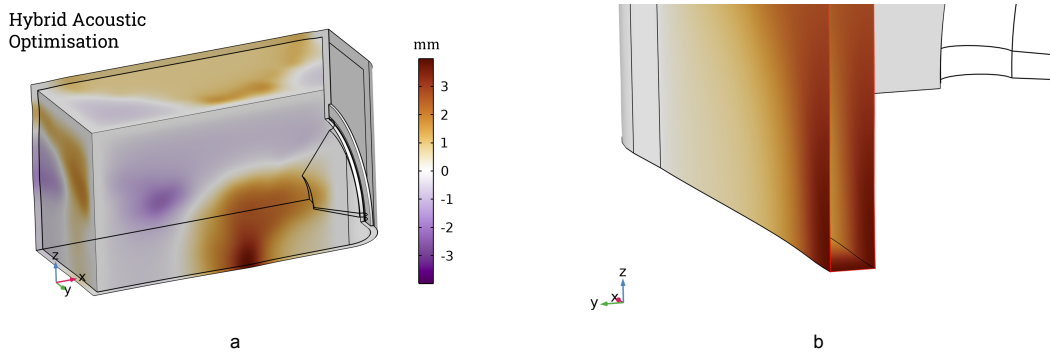


Figure 6.8: Visualisation of the final shape-optimised geometry, showing the total displacements (a) and a cross-sectional view of the final geometry showing that panel thickness is preserved throughout the optimisation process (b).

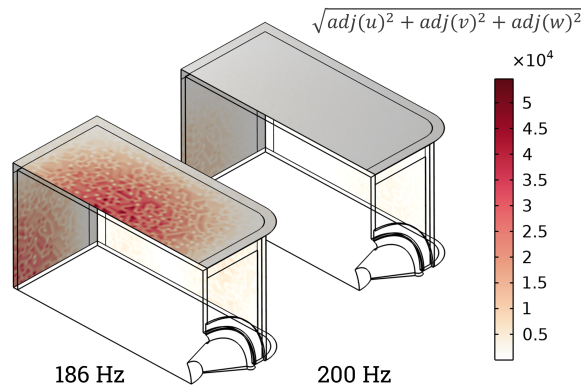


Figure 6.9: Spatial distribution of the adjoint sensitivity of the objective function at iteration 1 for 186 Hz and 200 Hz.

$\sum_{20}^{200} DI^2$ decreases from 45.72 to 45.56 compared to the 60 mm parametric result, which is a deviation of $<1\%$. This confirms that the panel modifications from optimisation do not compromise the uniformity obtained through parametric optimisation of the baffle edge radius. Figure 6.10a compares DI responses for the parametric and hybrid optimisation. Figure 6.10b shows the corresponding heat map of the final optimised geometry.

Both steps in the hybrid sequential optimisation improve their respective objective without compromising the other. Table 6.1 presents several performance metrics for the baseline, parametrically optimised, and hybrid optimised geometry, specifying between the two meshes.

| Configuration | Mean Q [-] | $\sum_{180}^{200} Q$ [-] | $\sum_{20}^{200} Q$ [-] | $\sum_{180}^{200} Q^2$ [-] | $\sum_{20}^{200} Q^2$ [-] | σ_Q [-] | $\sum_{20}^{200} DI^2$ [-] | Time [hrs] |
|-----------------------|------------|--------------------------|-------------------------|----------------------------|---------------------------|----------------|----------------------------|------------|
| Baseline | 1.38 | 74.7 | 125.7 | 849.7 | 925.1 | 2.87 | 63.1 | — |
| 60 mm fillet (Mesh 1) | 1.32 | 70.9 | 119.7 | 789.6 | 855.5 | 2.77 | 45.9 | 10 |
| 60 mm fillet (Mesh 2) | 1.38 | 77.4 | 125.5 | 974.4 | 1039.2 | 3.09 | 45.7 | 10 |
| Hybrid Optimised | 1.21 | 60.1 | 109.7 | 565.6 | 634.3 | 2.35 | 45.6 | 36 |

Table 6.1: Optimisation performance summary for the rectangular cabinet, showing iterative improvements from parametric baffle optimisation and shape optimisation.

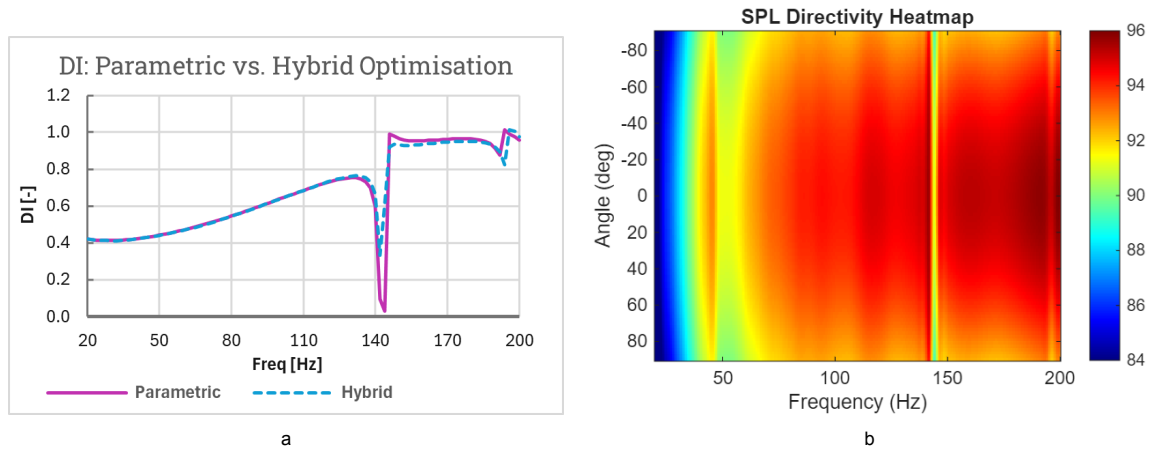


Figure 6.10: Comparison between the directivity index spectra of the parametric and hybrid optimisation (a) and the associated directivity heatmap corresponding to the final optimised geometry (b).

Several early iterations exhibit objective variations and occasional spikes. These spikes correlate with resonance frequency shifts relative to the 2 Hz sampling interval. Altered panel stiffness due to geometry perturbations shifts structural resonance modes within the evaluation range. These modes can fall near or on an evaluation frequency, causing the Q-factor to spike. Evaluation of intermediate iterations reveals that mode shapes and their accompanying Q-factor peaks are redistributed by optimisation, demonstrating that shape optimisation implicitly targets redistribution of resonance modes, namely resonance frequency relocation. Optimised results exhibit resonance frequencies that avoid the evaluation grid or are pushed out of the evaluation range entirely, validating that Q-factor reduction operates through modal frequency redistribution rather than increased panel damping.

6.4.4. Physical Mechanism Analysis

Q-factor improvement is attained through the optimisation by redistribution of stiffness and energy. Resonance frequencies, mode shapes, and spatial distribution of strain energy are analysed to highlight the physical mechanisms behind the Q-factor redistribution.

Resonance frequency analysis of baseline (60 mm fillet) and the optimised geometry reveals systematic shifts, presented in Table 6.2: 138 Hz to 142 Hz (rear and side panel), 186 Hz to 190 Hz (top panel), and 200 Hz to >200 Hz (rear panel, shifted outside of evaluation range). Panel deformations introduce spatial stiffness variations that lead to these frequency shifts.

| Mode | Panel | Baseline [Hz] | Optimised [Hz] | Shift [Hz] | Shift [%] |
|------|-----------|---------------|----------------|------------|-----------|
| 1 | Rear-Side | 138 | 142 | +4 | +2.9 |
| 2 | Top | 186 | 190 | +4 | +2.2 |
| 3 | Rear | 200 | >200 | +? | +? |

Table 6.2: Structural resonance frequency shifts between baseline and optimised geometries, demonstrating modal detuning mechanism.

Figure 6.11 presents baseline and optimised mode shapes for frequencies targeted in the SO, visualising how panel deformations alter modal behaviour. Changes in panel curvature alter stiffness locally, shifting anti-node locations and reducing displacements at the baseline resonance frequencies. The baseline mode shapes are detuned by stiffness variations, dispersing energy concentrations.

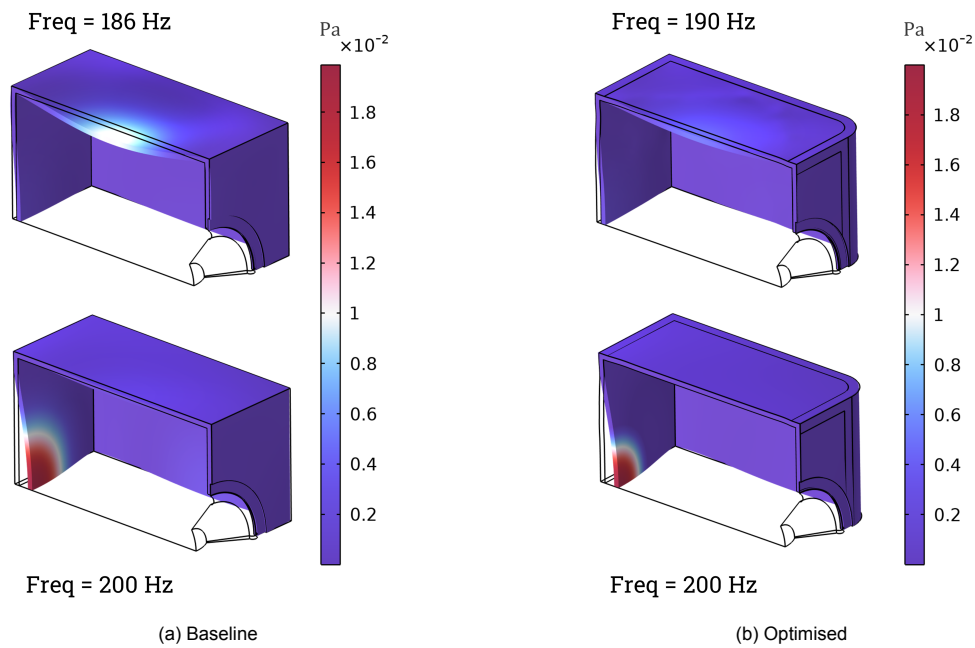


Figure 6.11: Mode shape comparison between baseline (a) and optimised (b) geometry, showing 186-190 Hz and 200-200 Hz, demonstrating mode detuning between baseline and optimised geometries.

Spatial distribution of the ESE reveals strain energy concentration within the solid structure. Redistribution of the ESE is analogous to energy concentrations being averaged out over a larger volume, reducing localised hotspots responsible for high Q-factors. Figure 6.12a visualises the baseline ESE distribution at 186 Hz and the ESE distribution after optimisation. The baseline geometry exhibits sharper localizations, with a peak value of $3.93 \times 10^{-2} \text{ J}$, concentrated on the top panel. The optimised geometry exhibits more distributed ESE, with peak value reduced to $3.29 \times 10^{-2} \text{ J}$, a 16.3% reduction. Figure 6.12b visualises the same baseline ESE at 186 Hz and the ESE distribution after optimisation at 190 Hz, where it can be seen that the ESE concentration on the top panel has indeed moved to this new resonance frequency.

Energy distributes over all panels and the localization on the top panel diminishes. The peak ESE at baseline resonances is reduced by 16.3% (186 Hz) and 9.65% (200 Hz), while mean ESE decrease is implied by the demonstrated reduction of the mean Q-factor. This result demonstrates ESE magnitude attenuation and improvement of spatial uniformity, manifesting as distributed lower-intensity resonances.




Figure 6.12: Comparison of spatial distribution of elastic strain energy at 186 Hz for baseline and final optimised geometry (a) and between 186 Hz and 190 Hz for baseline and optimised geometry (b) to show the relocation of the ESE concentration.

Analysis of iteration 1 adjoint sensitivity, visualised in Figure 6.9, confirms MMA's gradient-following behaviour. Areas of high sensitivity appear on the top and rear panels, coinciding with high localised ESE concentrations. This spatial correlation between sensitivity, ESE, and deformations validates that MMA targets high-impact areas, balanced over all considered frequencies. Modest deformations thus reflect lower sensitivity across multiple frequencies rather than low sensitivity at a single frequency. These mechanisms, resonance frequency shifts, mode shape detuning, and ESE redistribution, demonstrate that Q-factor reduction is achieved through stiffness modifications that relocate resonances within or outside the evaluation range. As local stiffness perturbations alter the modal response across all frequencies simultaneously, attenuation becomes more complex when multiple peaks are considered within a single optimisation.

6.4.5. Robustness Analysis

Robustness is evaluated through material property sensitivity, which is a dominant manufacturing uncertainty for structures produced using LFAM. Geometric robustness, adhering to LFAM constraints, is inherently addressed through the manufacturability constraints discussed in Section 6.1.2. The material loss factor exhibits the greatest uncertainty as LFAM process parameters, such as deposition rate, temperature, and compression ratio, affect the damping properties.

The shape optimised geometry is obtained using a nominal material loss factor of $\eta_{mat} = 0.015$. This geometry is re-simulated using the ASI model over 20-200 Hz while η_{mat} is varied from **0.005 - 0.025**, representing approximately $\pm 67\%$ uncertainty. Figure 6.13 shows the mean Q-factor for baseline and optimised geometry across 20-200 Hz for $\eta_{mat} = 0.005, 0.010, 0.015, 0.020, 0.025$. Both configurations show linear sensitivity to the loss factor. Baseline mean Q-factor ranges from **1.74** ($\eta_{mat} = 0.005$) to **1.11** ($\eta_{mat} = 0.025$), showing a CV of 15.4%. The optimised mean Q-factor ranges from **1.51** ($\eta_{mat} = 0.005$) to **0.96** ($\eta_{mat} = 0.025$), a CV of 15.3%. This finding demonstrates that the optimisation benefits are robust to loss factor uncertainty, providing confidence that the 39% $\sum_{20}^{200} Q^2$ improvement found for the nominal case ($\eta_{mat}=0.015$) is retained under LFAM material property variations.




Figure 6.13: Mean Q-factor for baseline and the final optimised geometry over material loss factor variations.

6.4.6. Framework Sensitivity to Material Printing Direction

Throughout the present research and optimisation, a transversely orthotropic material model is applied to the cabinet in a consistent direction, described in Section 3.4 as *direction 1*. The sensitivity of the optimisation framework to this material application is evaluated by re-running the shape optimisation with the material applied in *direction 2*. The baseline Q-factor response differs between material directions, as shown in Figure 3.5, and are thus not directly comparable. Rather, the performance of the optimisation is compared within a consistent material application.

| Configuration | Mean Q [-] | $\sum_{180}^{200} Q$ [-] | $\sum_{20}^{200} Q$ [-] | $\sum_{180}^{200} Q^2$ [-] | $\sum_{20}^{200} Q^2$ [-] | σ_Q [-] | $\sum_{20}^{200} DI^2$ [-] | Time [hrs] |
|---------------------------|------------|--------------------------|-------------------------|----------------------------|---------------------------|----------------|----------------------------|------------|
| Baseline (Dir. 2) | 1.51 | 74.4 | 137.4 | 914.0 | 1074.0 | 3.09 | 63.1 | — |
| Hybrid Optimised (Dir. 2) | 1.04 | 41.7 | 94.2 | 196.2 | 292.5 | 1.46 | 49.7 | 36 |

Table 6.3: Optimisation performance summary for the rectangular cabinet with the transversely orthotropic material applied in direction 2 (Dir. 2), showing the results of shape optimisation.

Table 6.3 presents the same metrics used for evaluation of the direction 1 results for the direction 2 baseline and hybrid optimised result. The objective function $\sum_{180}^{200} Q^2$ decreases from **914.0** to **196.2**, a 78.5% reduction. $\sum_{20}^{200} Q^2$ reduces by 72.8%, indicating that the framework achieves even greater relative improvement under material direction 2 compared to direction 1. $\sum_{20}^{200} DI^2$ is reduced from **63.1** to **49.7**, a 21.2% improvement, consistent with the DI preservation seen for the direction 1 results.

Table 6.4 shows the relative improvement that shape optimisation made with either direction compared to the rectangular baseline. It is evident that the optimisation in direction 2 shows much larger improvements. This finding can be attributed to direction 2 having only a single Q-factor peak in the 180-200 Hz evaluation range, which is more straightforward for the framework to attenuate than two peaks that need simultaneous attenuation, like in the case of direction 1.

| Configuration | Δ Mean Q [-] | $\Delta \sum_{180}^{200} Q$ [-] | $\Delta \sum_{20}^{200} Q$ [-] | $\Delta \sum_{180}^{200} Q^2$ [-] | $\Delta \sum_{20}^{200} Q^2$ [-] | $\Delta \sigma_Q$ [-] | $\Delta \sum_{20}^{200} DI^2$ [-] |
|---------------|---------------------|---------------------------------|--------------------------------|-----------------------------------|----------------------------------|-----------------------|-----------------------------------|
| SO Dir. 1 | -12.3% | -19.5% | -12.6% | -33.4% | -31.4% | -18.1% | -27.7% |
| SO Dir. 2 | -31.1% | -43.9% | -31.4% | -78.5% | -72.8% | -52.7% | -21.2% |

Table 6.4: Relative improvement of hybrid optimisation for each printing direction, with respect to the rectangular baseline.

Figure 6.14 shows the Q-factor response for the rectangular baseline and optimised geometry for both material directions. The modal detuning mechanism identified Section 6.4.4 is again observed. The single resonance peak in the 180-200 Hz band is attenuated and appears to be moved to >200 Hz. The lower resonance modes, near 150-160 Hz, are shifted to 174 Hz.




Figure 6.14: Q-factor responses between baseline rectangular cabinet and final shape optimised geometry for both material direction 1 and direction 2.

Figure 6.15a and b show the final optimised geometry for material directions 1 and 2, respectively. Direction 2 exhibits significantly larger deformations of ± 6 mm, with the side panel showing a larger affected area and the rear panel targeted more centrally. This confirms that while the framework converges on different optimal geometries depending on material orientation, the underlying mechanism remains consistent. These findings highlight the value of the framework in evaluating printing direction, revealing that acoustic gain potential varies per printing strategy and that material orientation warrants consideration as a design variable in future work.

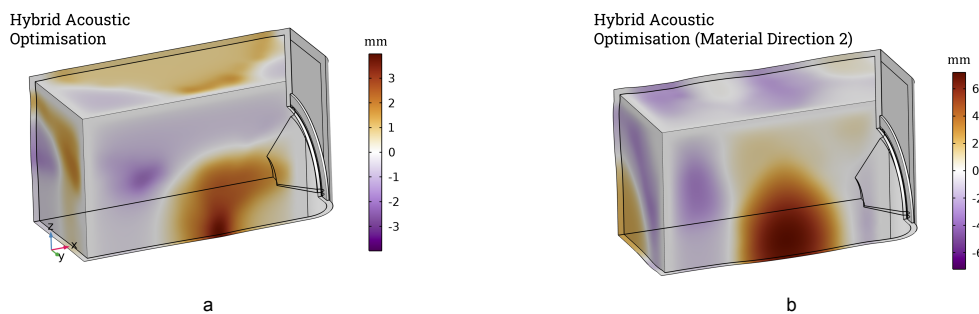


Figure 6.15: Visualisation of the final geometry of the hybrid acoustic for material direction 1 (a) and material direction 2 (b).

6.4.7. Alternative Optimisation Approach: Structural Eigenfrequency

An alternative optimisation strategy is evaluated to assess computational efficiency and mechanism generality. Based on the resonance frequency mechanism identified in Section 6.4.4, the objective function is redefined to maximisation of the structural eigenfrequencies of the cabinet within the range of interest. Baseline analysis identifies two eigenfrequencies between 100-200 Hz: **139.41 Hz** (λ_1) and **194.27 Hz** (λ_2). Interestingly, both eigenfrequencies are close to the resonance modes found for the baseline rectangular cabinet in the ASI simulation. The sum of λ_1 and λ_2 is maximised, pushing both modes *above* 200 Hz. The formulation of eigenfrequency maximisation is chosen to contrast the Q-factor-based approach, allowing for more in-depth exploration of the physical mechanism behind each approach and their similarities. Eigenfrequency analysis solely involves solid mechanics, eliminating acoustics, loss factors, and multiphysics couplings. This strategy tests whether similar structural energy redistribution and Q-factor improvements can be attained at lower computational cost.

Eigenfrequency shape optimisation converges in ~ 20 iterations, with a full 25-iteration run taking 10 minutes on average to compute, $192\times$ faster than hybrid acoustic optimisation. This speed-up results from eliminating acoustic-structure interaction coupling, primarily the BEM-solid interaction that dominates assembly time. The same geometric constraints are applied to the eigenfrequency study, ± 20 mm displacement, 20 mm features, and identical panel selection. The two initial modes are moved up to **262.46 Hz** and **>300 Hz**, respectively. λ_2 is pushed above 300 Hz and falls outside of the eigenfrequency search range. Crucially, no new eigenmodes are introduced between 20-200 Hz, highlighting again the transferability of the physical mechanism to broadband analysis.

The optimised geometry exhibits panel deformations larger than the hybrid acoustic optimisation does, ranging ± 18 mm, close to the maximum displacement limit. Figure 6.16a shows the final geometry of the hybrid acoustic optimisation and Figure 6.16b that of the eigenfrequency optimisation. The eigenfrequency optimisation shows larger deformation magnitudes but a similar distribution on the top and rear panel. These similarities are identified as follows: A v-shape on the top panel and a larger deformation in the top-right corner of the rear panel. Figure 6.18a and b show the final optimised cabinets for the Hybrid Acoustic and Eigenfrequency optimisation, respectively. Both eigenfrequencies being shifted to >200 Hz confirms that structural eigenfrequency optimisation shifts the eigenfrequencies through stiffness redistribution, similar to the hybrid acoustic optimisation.

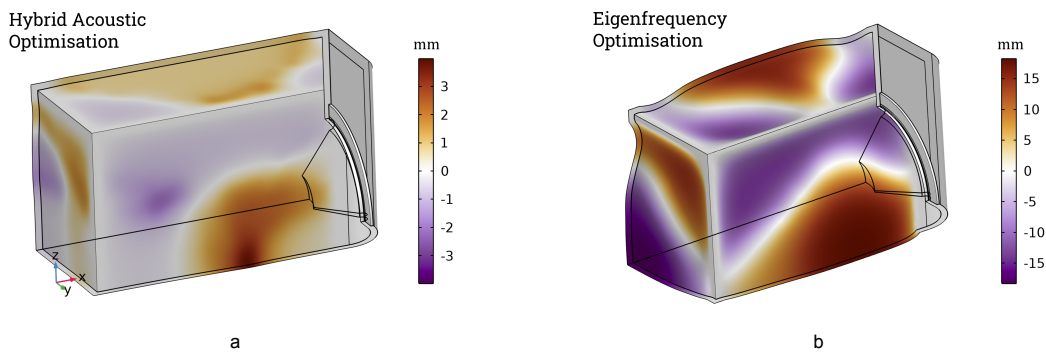


Figure 6.16: Visualisation of the final geometry of the hybrid acoustic (a) and eigenfrequency optimisation (b).

Simulating the eigenfrequency-optimised geometry in the full ASI model reveals a mean Q-factor of **3.30**, a **139.2%** increase over baseline (mesh 2), compared to 12.6% *reduction* obtained through the hybrid acoustic optimisation. $\sum_{20}^{200} Q^2$ becomes **1.29×10^4** , a **12.4** \times increase over baseline, compared to 39% decrease obtained by the hybrid acoustic optimisation. The results are collected for comparison in Table 6.5. Figure 6.17 shows the Q-factor response for the eigenfrequency- and hybrid-optimised design over the 20-200 Hz range, revealing a nearly flat and 0 Q-factor response for the eigenfrequency optimised design, except for a large peak at 194 Hz ($Q=96.4$).

| Method | Objective | Iterations | Time [hrs] | Mean Q [-] | $\sum_{180}^{200} Q$ [-] | $\sum_{20}^{200} Q$ [-] | $\sum_{180}^{200} Q^2$ [-] | $\sum_{20}^{200} Q^2$ [-] |
|----------------------------|----------------|------------|------------|------------|--------------------------|-------------------------|----------------------------|---------------------------|
| Baseline (Mesh 2) | — | — | — | 1.38 | 77.4 | 125.5 | 974.4 | 1039.2 |
| Acoustic (Q^2) | Min Q-factor | 25 | 36 | 1.21 | 60.1 | 109.7 | 565.6 | 634.3 |
| Structural (λ_1) | Max eigenfreq. | 25 | 0.1 | 3.30 | 243.1 | 300.1 | 1.28×10^4 | 1.29×10^4 |

Table 6.5: Comparison of eigenfrequency and hybrid acoustic optimisation.




Figure 6.17: Comparison of the Q-factor response over 20-200 Hz of the eigenfrequency- and hybrid acoustic-optimised final geometry.

This comparison provides critical strategy insights. Eigenfrequency optimisation manages to move structural eigenfrequencies through a physical mechanism similar to that utilised in the hybrid acoustic optimisation, but fails to consider air domains and multiphysics couplings that prove crucial to the ultimate ASI performance.

The eigenfrequency optimisation does prove valuable in its computational cost being $192\times$ lower than that of hybrid acoustic optimisation, or **0.52%** of the cost. Eigenfrequency optimisation offers a highly efficient preliminary design study and manages to shift structural resonance modes outside of the range of interest. However, the emergence of a strong resonance mode when combined with acoustics demonstrates that an eigenfrequency approach alone cannot capture the multiphysics nature of the full ASI problem. Eigenfrequency optimisation proves effective for rapid design space exploration, but the hybrid acoustic optimisation is necessary for attaining optimal acoustic performance.



Figure 6.18: Render of the final cabinets of the hybrid acoustic (a) and eigenfrequency optimisation (b).

7

Discussion

7.1. Summary of Key Trade-offs Observed

Development of the vibroacoustic framework required navigation of several trade-offs, concerning modelling fidelity, computational effort, and sufficiency of validation. The core trade-off made in the modelling framework selected a FEM/BEM strategy over FEM-PML to eliminate numerical inaccuracies such as directivity asymmetry and Q-factor location inconsistencies resulting from inadequately sized PML regions. An equidistant sampling interval of 2 Hz was adopted over ISO preferred frequencies to prevent undersampling of higher frequencies, at the cost of an increase of evaluation points (91 equidistant versus 41 ISO preferred frequencies at 1/12th octave interval). Chapter 5 demonstrated agreement in trends between simulation and measurement of a physical prototype, showing an approximate 1-2% deviation of resonance frequency. This result revealed sufficient accuracy for optimisation trend prediction, despite mesh dependency of absolute values of the Q-factor.

A trade-off following from the choice of FEM/BEM strategy was its accurate prediction of structural dynamics and far-field radiation, paired with an inability to accurately capture the Helmholtz resonance at ports. This requires empirical correction, inclusion of another method such as CFD, or the exclusion of the port contribution from the simulation. FEM-PML does not suffer the same inability but comes at prohibitive computational cost within the context of an optimisation framework. A trade-off between acoustic completeness and computational feasibility was made and shaped the scope of the present research.

Section 6.4.7 demonstrated that maximisation of eigenfrequencies through shape optimisation shifts structural resonance modes outside of the range of interest, but fails to capture full ASI behaviour, solidifying the need for hybrid acoustic optimisation. The eigenfrequency optimisation does come at 0.52% the computational cost of hybrid acoustic optimisation and proves valuable for preliminary design space exploration, such as establishing deformation limits. The gain in computational performance stems from omission of acoustic radiation damping, BEM-structure interfaces, and structural damping.

A trade-off in displacement budget balanced computational practicality, numerical convergence, and transferability to LFAM. Prescribed deformation coupling enables ± 20 mm panel deformations while preserving consistent panel thickness and minimum LFAM feature size ($2 \times$ nozzle diameter). Move limits within MMA of 0.05 enables practical convergence within 25 iterations, with the displacement limits reached in 15-20 iterations, followed by geometry refinement. Formal gradient tolerance ($< 10^{-6}$) was not satisfied. Practical convergence of the optimisation was instead identified through objective function CV (3.04%), declining adjoint gradient value, and minimisation of the mean geometric change per iteration.

Lastly, the bandwidth of interest balances computational effort versus identification of physical mechanisms transferable to full-range optimisation. Section 6.4.3 demonstrated that narrowband optimisation of summed Q^2 over 180-200 Hz reduced full-range (20-200 Hz) Q-factor by 12.6%, confirming that mechanisms identified in high-Q regions transfer to broadband performance. Optimisation over the entire 20-200 Hz bandwidth would scale runtime by approximately $9 \times$, prohibitive for iterative development of the framework.

7.2. Insights from the Modelling and Optimisation Process

Validation of the FEM/BEM simulation against the Reflex Sub presented in Chapter 5 showed a mean absolute error of SPL of 3.16 dB and resonance frequency prediction within 1-2% of anechoic measurements over 20-200 Hz. This comparative accuracy proved sufficient for trend prediction and optimisation, while absolute prediction required normalisation to the measured response. Q-factor trends and directivity patterns aligned within measurement uncertainty. Robustness of the optimisation results was validated across $\pm 67\%$ loss factor variation, representing material property uncertainty from LFAM process variations. The practically identical linear sensitivity confirms robustness.

Port modelling limitations observed in Chapter 5 dictated limitation of scope to sealed cabinets. Over-prediction of Helmholtz resonances, observed as a 39.8 dB peak and 24.4 dB dip between 20-80 Hz, stemmed from omission of non-linear acoustics and viscothermal losses in narrow port geometries. NRA attenuated the peaks marginally, providing a net reduction of 17 dB, but could not eliminate the over-prediction entirely. Accurate prediction requires empirical port loss corrections or CFD [42].

The FEM/BEM simulation and optimisation model demonstrated reliability in comparative optimisation despite mesh dependency of absolute Q-factor magnitudes. Changes in discretisation shifted baseline values by approximately 5% while preserving resonance frequency locations within 1-2%. Preserving element quality between discretisation proves effective in minimising Q-factor magnitude variations.

Prescribed deformation coupling and Yeoh smoothing proved critical for adhering to LFAM constraints. Unconstrained mesh perturbations do not maintain consistent panel thickness within ± 1 mm from 10 mm nominal and risk inversion errors when perturbations exceed panel thickness. Translation of deformations represents a novel regularisation method for large geometry perturbations within LFAM constraints, enabling designs accessible to free-form additive manufacturing but not to conventional construction. Yeoh hyperelastic smoothing ensures smooth and non-discontinuous mesh features.

Analysis of the physical mechanisms behind the optimisation framework revealed a mechanism of energy redistribution rather than damping enhancement. Table 6.2 presents resonance frequency shifts (+2.9%, +2.2%, +>2%) that represent relocation of vibration modes relative to the 2 Hz evaluation grid. Figure 6.12 reveals spatial redistribution of ESE, changing localised hotspots to diffuse energy sharing over all panels, decreasing peak ESE magnitude by 16.3%. Q-factors decreased via modifications made to localised stiffness while material loss factor remained constant ($\eta_{mat} = 0.015$), validating that non-uniform panel geometry and curvatures enable targeted detuning of resonance modes. Shifted resonance frequencies falling between the sample points with 2 Hz interval exhibited lower Q^2 contributions, while shifted resonances that aligned with the evaluation grid produced objective spikes.

Evaluation of the framework under alternative transverse material application (direction 2, as per Section 6.4.6) further revealed that printing direction significantly affects both the baseline Q-factor response and the magnitude of optimisation gains. Despite differing baseline characteristics and optimal geometries, the modal detuning mechanism remained consistent, confirming that the framework's effectiveness is not contingent on a specific material orientation.

Practical convergence was confirmed through plateau and perturbation analysis, rather than formal gradient tolerance of MMA. BEM-solid interactions were identified as inherently ill-conditioned, preventing MMA gradient tolerance termination criteria ($< 10^{-6}$) from being met. Practical convergence manifested as objective value CV of 3.04%, gradient value CV of 8.16%, and saturation of geometric change at a CV of 1.41% over the last 6 iterations. BEM-solid coupling ill-conditioning, inherent to dense boundary element matrices in vibroacoustic optimisation, impedes mathematical rigour.

The physical mechanisms enabling the presented framework extend beyond loudspeaker cabinet design to other vibroacoustic domains where curved geometries are attainable. Examples include soundboard vibrations in musical instruments, road noise in automotive panels, low-frequency rumble in HVAC ductwork, and electromechanical hum in electrical enclosures. All involve thin-walled structures excited near resonances. Redistribution of modal energy through panel curvatures is attainable through not only LFAM but also other methods such as stamping or composite layup. The redistribution mechanism can enable resonance attenuation that outperforms added mass or damping.

7.3. Limitations of the Current Approach

The primary limitation of the optimisation framework is the inability to predict absolute Q-factor magnitudes due to mesh dependency, rather relying on comparative performance assessment. Baseline summed Q^2 varied as discretisation changes altered strain energy integration accuracy. Resonance frequency peak locations remained stable within 1-2% per the convergence criterion presented in Section 3.7.1. Mesh dependency is inherent to FEM energy metrics, since these are computed as ratios of volume integrals. The optimisation performance improvement requires relative analysis, focusing on percentages, energy redistribution patterns, and mode detuning, rather than absolute values. Experimental validation carried out in Chapter 5 confirmed accuracy in trend comparison, yet physical prototyping remains necessary to quantify whether the simulated relative improvements carry over to manufactured prototypes. No optimised geometry has been manufactured using LFAM, leaving performance improvement claims experimentally unverified.

Port modelling discrepancies (Helmholtz resonance over-prediction from omitted viscothermal losses and acoustic non-linearities) limit the applicability of the developed framework to sealed cabinets. Quarter-symmetry adopted for the optimisation framework assumes symmetry of cabinet and driver in both width and height planes. Thus, applicability of the framework to asymmetric driver placement and cabinet geometry was not validated. Computational efficiency and feasibility were balanced against physical mechanism identification by narrowband development (180-200 Hz). Section 6.4.3 demonstrated that narrowband improvements can transfer to the full evaluation range (12.6% improvement over 20-200 Hz). However, low-frequency and high-frequency modal distributions may conflict during simultaneous optimisation, which remains unexplored. Full bandwidth optimisation would require approximately $9\times$ computational cost (324 hours per study), prohibitive for iterative framework development.

Uncertainties in material properties and characteristics from LFAM process variations remain experimentally unverified. Material property robustness was validated numerically through $\pm 67\%$ variation of the material loss factor η_{mat} , producing 15.3% performance variation, but other process variations persist. These include geometric variations through bead width fluctuations, inconsistencies in layer adhesion, and anisotropy from print path orientation varying relative to model coordinate axes. These factors can potentially shift optimised resonance frequencies. Batch-to-batch Q-factor variations remain experimentally unvalidated.

Constraints of shape optimisation limited the available design space and convergence behaviour. Edge exclusion zones were introduced to prevent mesh-inversion errors at panel junctions, reducing the average optimisable surface area per panel by 17.4%. Panel selection and exploration of asymmetric modes was also limited due to utilisation of quarter geometry to enable iterative framework development within feasible computational effort. BEM-solid coupling ill-conditioning prevented formal MMA gradient tolerance from being met, with practical convergence instead characterised through objective value CV of 3.04%, gradient value CV of 8.16%, and saturation of geometric change at a CV of 1.41% over the last 6 iterations. Global optimality without local minima cannot be guaranteed with gradient-based methods. Furthermore, COMSOL's optimisation module restricts objective function definitions to aggregate operators over the solution set, or frequency range in this case, such as sum, maximum, or minimum. There is no access to indexed individual solutions within the frequency range, which results in an inability to properly define smoothness and variation functions.

Development of the framework on a single geometry limits demonstrated applicability. The reference cabinet served as a platform lacking acoustic refinement for development of the optimisation framework. The Reflex Sub presented in Chapter 5 however, indicated preliminary improvements of $<5\%$ (Section 6.1.3) due to the effective existing double-wall braced construction already providing distributed stiffness. The framework proves effective on baseline geometries prior to acoustic refinement but presents diminishing returns on pre-optimised cabinets. Scalability to cabinet configurations featuring multiple drivers, horns, higher-frequency drivers, and other acoustic modifiers such as transmission lines remains unvalidated. Generality of the resonance frequency redistribution mechanism presented in Section 6.4.7 suggests transferability, yet characterisation of complex acoustic loading coupling phenomena remains unexplored.

7.4. Future Work and Experimental Validation Path


The next critical step is represented by physical prototyping and experimental validation in order to quantify whether simulated relative improvements translate to manufactured cabinets. Printing the optimised cabinet using C-250CF and testing its performance under anechoic conditions as per the method presented in Chapter 5 would validate the predicted 12.6% full-band reduction in Q-factor through accelerometer measurements and acoustic decay analysis, with preservation of directivity verified. Robustness can be assessed across 5-10 printed specimens to quantify batch-to-batch Q-factor variation and validate panel thickness preservation under bead-width fluctuations. This approach would transform the presented comparative optimisation improvements into absolute performance validation and establish whether simulated Q^2 predictions correlate with physical decay behaviour.

Computational efficiency and scalability improvements would allow for full-bandwidth optimisation. Extension from 11 frequencies (180-200 Hz) to 91 frequencies (20-200 Hz) requires reductions of per-frequency computational cost to maintain feasibility. Suggestions include adaptive frequency sampling near high-Q regions to reduce evaluation points, parallel computation across cluster nodes via domain decomposition, and use of surrogate models trained on high-fidelity iterations. Alternatively, the suggested workflow of preliminary eigenfrequency optimisation followed by acoustic refinement may be explored. These methods could enable a reduction in computation times that facilitates modal detuning over a wider frequency range. More robust solver configurations could be used to solve the BEM-solid coupling, such as multigrid preconditioning for dense BEM matrices or alternative BEM formulations. These configurations may enable formal MMA convergence.

Scope extensions would validate transferability of the framework and address configuration limitations. Inclusion of empirical port loss corrections or extension to CFD would enable more accurate prediction of ported cabinet responses. Capturing viscothermal boundary layers and viscous damping would allow for the inclusion of transmission line cabinets. Asymmetric loudspeaker configurations require modelling of the full geometry, enabling analysis of asymmetric multi-driver configurations. Computational cost is expected to increase approximately $4\times$. Generality of the resonance frequency redistribution mechanism may be validated beyond single-driver sealed cabinets, across crossover-induced modal coupling, horn-loaded drivers with throat loading effects, and transmission lines with internal acoustic path interactions.

Optimisation methodology advances would refine design space exploration and accuracy of predictions. Weighted multi-objective optimisation formulations could be used to explore and quantify trade-offs between Q-factor reduction, directivity uniformity, and spatial response consistency. Preliminary single-objective results present minimal conflicts between the performance parameters, but explicit mapping would identify orthogonality limits. Exploration of alternative optimisation methods, such as genetic algorithms or multi-start strategies, could also provide reduction of local minima uncertainty inherent to gradient-based optimisation methods. The edge exclusion zones may be eliminated by implementing continuity conditions between panels on their shared edges, recovering part of the design space. A method for screening the optimisation metrics on initial iterations, such as evaluation of gradient values, should be employed to evaluate the potential effectiveness of the optimisation framework on a particular geometry. As such, low-potential optimisation would be identified without the computational cost of empirical analysis.

Industrial workflow integration would transition the framework from a stand-alone research method to practical design tool. This integration would involve exporting and post-processing of the optimised mesh to prepare it for manufacturing by decimating small features and refining mesh inaccuracies resulting from tessellation. Implementation of additional manufacturing constraints imposed by LFAM would improve manufacturability of the optimised geometry further, minimising manual intervention and verification. Such constraints include overhang angle limits, contour smoothness requirements, and ensuring that the structure can be printed to be airtight.



Conclusions

8.1. Summary of Key Findings

This thesis has investigated how Large Format Additive Manufacturing enables vibroacoustic optimisation of loudspeaker cabinets through geometries inaccessible to conventional manufacturing methods. A hybrid FEM/BEM simulation method was employed to develop a gradient-based shape optimisation framework, coupling structural dynamics with acoustic radiation to capture vibroacoustic behaviour. Prescribed deformation techniques preserved manufacturability constraints imposed by the LFAM process while exploring non-uniform panel curvatures. Experimental validation was performed using the Reflex Sub prototype, revealing resonance frequency predictions within 1-2% and SPL agreement within 3.16 dB. These results demonstrated sufficient comparative accuracy for trend-based optimisation, despite mesh-dependency of absolute Q-factor magnitudes.

Applying this framework, shape optimisation of a rectangular reference cabinet resulted in 41.9% Q^2 reduction over 180-200 Hz through redistribution of elastic strain energy. Structural modes shifted by +2.9%, +2.2%, and +>2%, relocating resonances away from evaluation frequencies. ESE spatial distributions transitioned from localised panel hotspots to diffuse energy sharing over multiple panels. Parametric baffle edge filleting reduced the squared directivity index by 27.3% over 20-200 Hz. Sensitivity analysis confirmed orthogonality between directivity and resonance suppression. This orthogonality motivated a sequential hybrid strategy, improving both performance metrics without compromise. Robustness of the optimisation to material properties was validated across $\pm 67\%$ loss factor variation, resulting in 15.3% performance variation, demonstrating that relative improvements remain consistent across LFAM process uncertainties. Re-running the optimisation using an alternative transverse material application confirmed that the modal detuning mechanism remains consistent across printing directions, while the magnitude of optimisation gain varies. This finding demonstrated both the generalisability of the framework and the influence of printing direction on optimisation potential.

The optimisation framework operates within a defined scope. Single-driver sealed cabinets optimised for low frequencies (20-200 Hz) are considered. Port modelling limitations from omitted viscothermal losses and non-linear acoustics restrict applicability to sealed configurations. Quarter-symmetric configurations are employed to reduce computational effort and the framework provides comparative rather than absolute performance assessment. Practical convergence was confirmed through plateau analysis (objective value CV 3.04%, gradient value CV 8.16%, geometric change CV 1.41%), rather than formal gradient tolerance of MMA. BEM-solid ill-conditioning prevented mathematical convergence verification, necessitating reliance on engineering sufficiency. Physical prototyping remains necessary to confirm transferability of simulated relative improvements to manufactured cabinets.

8.2. Contributions to Vibroacoustic Design

The present research advances vibroacoustic optimisation methodology through three principal contributions. First, resonance frequency redistribution was identified as the main mechanism enabling Q-factor reduction. Unlike traditional methods of resonance and vibration reduction, such as adding damping material or structural mass, this work employs spatially varying panel curvatures to detune vibration modes through localised stiffness variations. This approach reduces peak ESE magnitude

by 16.3% while redistributing the elastic strain energy spatially. This challenges conventional design heuristics favouring internal bracing or damping material, demonstrating that targeted geometric perturbations complement mass or damping treatments for modal control. The independence of this mechanism from material loss factor (η_{mat} remained constant as $\sum_{20}^{200} Q$ decreased by 12.6%) confirms that geometric optimisation is a strategy transferable beyond specific damping materials. This geometric approach to modal control is enabled by the shape freedom that LFAM offers. Conventional flat-panel construction could not achieve the required spatially varying curvatures, leaving internal bracing or damping treatments as the only viable options.

Second, prescribed deformation coupling between outer and inner surfaces of the cabinet enabled geometry perturbations of ± 20 mm while preserving consistent panel thickness. Additionally, Yeoh hyperelastic smoothing maintained mesh element quality under large deformations. The combination of these integrations unlocked geometries manufacturable within typical LFAM tolerances, which are inaccessible to typical flat-panel construction. The use of large-deformation mesh mechanics to accommodate shape perturbations exceeding panel thicknesses, without mesh failures and while preserving thickness, represents a novel integration within a vibroacoustic optimisation framework. Continuous surface optimisation that exploits LFAM's geometric freedom was achieved.

Third, the validated hybrid FEM/BEM framework presents comparative optimisation reliability despite mesh dependency of absolute Q-factors. Changes in discretisation shifted baseline Q^2 values by 5% while resonance frequency locations remained stable within 1-2%. This suggests that effective gradient-based shape optimisation may rely on predictive trends rather than absolute magnitude precision. If confirmed by experimental validation of the optimised geometry, comparative accuracy would suffice for geometry-driven improvements even when absolute performance prediction proves inaccessible.

8.3. Implications and Outlook

The presented framework's mechanism, spatial stiffness variations targeting specific resonance frequencies, extends beyond loudspeakers to other vibroacoustic applications where thin-walled structures are excited near resonances. Examples include soundboard vibrations in musical instruments, road noise in automotive panels, low-frequency rumble in HVAC ductwork, and electromechanical hum in electrical enclosures. LFAM enables design freedom beyond flat-panel construction, transforming structural optimisation from discrete element placement to continuous surface optimisation. This shift from bracing placement to panel shaping unlocks design space traditionally confined to empirical prototyping or expensive panel moulding processes. Furthermore, printing direction presents an unexplored design variable with demonstrated influence on optimisation potential. Future work should treat material orientation as an additional design variable, expanding the accessible design space further.


Realising this potential requires assessment of framework limitations prior to industrial adoption. Sealed-configuration scope restricts widespread applicability, necessitating empirical port corrections or CFD implementation to address Helmholtz resonance discrepancies. The assumption of quarter-symmetry limits applicability to symmetric driver and cabinet configurations. Extension to full-geometry requires increased computational effort but faces no methodological barriers. Most critically, experimental validation of optimised cabinets remains pending. Transferability of the simulated Q-factor reduction to measurable performance improvements requires prototype fabrication and anechoic testing. Manufacturing reliability should also be established through analysis of batch-to-batch variation of the Q-factor.

Despite these limitations, comparative studies validate the generality of the resonance frequency redistribution mechanism in the hybrid acoustic optimisation. Eigenfrequency-only optimisation presented in Section 6.4.7 achieves attenuation of structural resonance modes in the 20-200 Hz range at 0.52% computational cost compared to the hybrid acoustic method, but fails to capture full ASI behaviour. This motivates a two-stage workflow, with rapid exploration of the design space through eigenfrequency optimisation, followed by acoustic refinement for final improvements. This approach could improve accessibility to vibroacoustic optimisation, allowing rapid screening of candidates before committing to expensive computations. Geometry-driven resonance control, enabled by the emergence of LFAM, represents a strategy that is complementary to material damping and addition of mass. This thesis demonstrates LFAM's potential in the context of vibroacoustic problems and presents a validated pathway from simulated predictions to physical realisation.

References

- [1] Jun Fu, Minghui Xu, Zengfeng Zhang, Wenjie Kang, and Yong He. “Muffler structure improvement based on acoustic finite element analysis”. In: *Journal of Low Frequency Noise, Vibration and Active Control* 38.2 (2019), pp. 415–426.
- [2] Mansinh S Kumbhar, Atul Miskin, and Vishal Chaudhari. *Study of intake and exhaust system acoustic performance refinement with the help of vibro-acoustic analysis tool*. Tech. rep. SAE Technical Paper, 2010.
- [3] Peter Risby Andersen, N Aage, G Seongyeol, J Kook, et al. “Vibro-acoustic shape optimization for radiation problems: Application to compact loudspeakers”. In: *24th International Congress on Acoustics*. International Commission for Acoustics. 2022.
- [4] Ahmad H Bokhari, Martin Berggren, Daniel Noreland, and Eddie Wadbro. “Loudspeaker cabinet design by topology optimization”. In: *Scientific Reports* 13.1 (2023), p. 21248.
- [5] Peter Risby Andersen, Vicente Cutanda Henríquez, Niels Aage, and Junghwan Kook. “3D shape optimization of loudspeaker cabinets for uniform directivity”. In: *Structural and Multidisciplinary Optimization* 65.12 (2022), p. 343.
- [6] Rajitha Udawalpola and Martin Berggren. “Optimization of an acoustic horn with respect to efficiency and directivity”. In: *International journal for numerical methods in engineering* 73.11 (2008), pp. 1571–1606.
- [7] Peter Risby Andersen, Gyeong-Tae Lee, Daniel Gert Nielsen, Junghwan Kook, Vicente Cutanda Henríquez, Niels Aage, and Yong-Hwa Park. “Experimental characterization of a shape optimized acoustic lens: Application to compact speakerphone design”. In: *The Journal of the Acoustical Society of America* 153.4 (2023), pp. 2351–2351.
- [8] Efen Fernandez Grande. “Sound radiation from a loudspeaker cabinet using the boundary element method”. In: *DTU Orbit* (2008).
- [9] Peter Constantin and Ciprian Foiaş. *Navier-stokes equations*. University of Chicago press, 1988.
- [10] Qing-Huo Liu and Jianping Tao. “The perfectly matched layer for acoustic waves in absorptive media”. In: *The Journal of the Acoustical Society of America* 102.4 (1997), pp. 2072–2082.
- [11] Hermann von Helmholtz. “Über die Theorie der stationären oscillatorischen Bewegungen in Flüssigkeiten”. In: *Journal für die reine und angewandte Mathematik* 57 (1860), pp. 25–55.
- [12] Wai Cheng Lau and H Fratlin. “Hybrid FEM–BEM Techniques for Accurate Prediction of Sound Radiation in Complex Structures”. In: *Advanced Computational Acoustics Engineering* 2.1 (2024), pp. 20–25.
- [13] Xiang Xie and Yijun Liu. “Efficient multi-frequency solutions of FE–BE coupled structural–acoustic problems using Arnoldi-based dimension reduction approach”. In: *Computer Methods in Applied Mechanics and Engineering* 386 (2021), p. 114126.
- [14] COMSOL AB. *COMSOL Multiphysics 6.4: Acoustics Module User’s Guide - When to Use BEM*. Accessed: 2026-03-16. COMSOL AB. Burlington, MA, USA, 2024. URL: https://doc.comsol.com/6.4/doc/com.comsol.help.aco/aco_ug_pressure.05.149.html.
- [15] Dayton Audio. *SD270A-88 10” DVC Subwoofer 4 Ohm*. Accessed: 2026-03-16. URL: <https://www.daytonaudio.com/product/133/sd270a-88-10-dvc-subwoofer-4-ohm>.
- [16] Omni calculator. *Helmholtz Resonator Calculator*. Accessed: 2026-03-16. URL: <https://www.omnicalculator.com/physics/helmholtz-resonator>.
- [17] Alex Salvatti, Allan Devantier, and Doug J Button. “Maximizing performance from loudspeaker ports”. In: *Journal of the Audio Engineering Society* 50.1-2 (2002), pp. 19–45.
- [18] A. Neville Thiele. “Loudspeakers in Vented Boxes”. In: *Proceedings of the Institution of Engineers, Australia*. Vol. 22. 1971, pp. 189–196.
- [19] Wolfgang Klippel and Joachim Schlechter. “Measurement and visualization of loudspeaker cone vibration”. In: *Signal* 30 (2006), p. 20.
- [20] Wolfgang Klippel. *Cone Vibration and Radiation Diagnostics*. Tech. rep. AN-31. Accessed: 2026-03-16. Klippel GmbH, 2012. URL: https://www.klippel.de/fileadmin/_migrated/content_uploads/AN_31_Cone_Vibration_and_Radiation_Diagnostics.pdf.
- [21] Airtech Europe Sarl. *Dahltram® C-250CF Data Sheet*. Accessed: 2026-03-17. 2020. URL: <https://ceadgroup.com/wp-content/uploads/2022/05/Airtech-Dahltram%EF%BF%BD-C-250CF.pdf>.
- [22] Ahmed Arabi Hassen, Ralph B Dinwiddie, Seokpum Kim, Halil Levent Tekinalp, Vipin Kumar, John Lindahl, Pritesh Yeole, Chad Duty, Uday Vaidya, Hsin Wang, et al. “Anisotropic thermal behavior of extrusion-based large scale additively manufactured carbon-fiber reinforced thermoplastic structures”. In: *Polymer Composites* 43.6 (2022), pp. 3678–3690.
- [23] Bryce Doll. “Maximizing Bass Reflex System Performance Through Optimization of Port Geometry”. In: *Honors Undergraduate Theses* (2020).
- [24] Mithun Hari and Todd Michael. *Exploring Asymmetric Vent Designs in Loudspeaker Systems: A Combined COMSOL and Experimental Approach*. COMSOL Conference Paper. Accessed: 2026-03-17. 2024. URL: <https://www.comsol.com/paper/loudspeaker-vent-shape-optimization-a-study-using-a-lumped-speaker-model-136172>.

- [25] Yves Pene, Yoachim Horyn, and Christophe Combet. "Non-linear acoustic losses prediction in vented loudspeaker using computational fluid dynamic simulation". In: *Audio Engineering Society Convention 148*. Audio Engineering Society, 2020.
- [26] COMSOL AB. *COMSOL Multiphysics 6.4: Acoustics Module User's Guide*. Accessed: 2026-03-17. COMSOL AB, 2024. Chap. Pressure Acoustics, Frequency Domain: Mesh Settings. URL: https://doc.comsol.com/6.4/doc/com.comsol.help.aco/aco_ug_pressure.05.147.html.
- [27] COMSOL AB. *Performing a Mesh Refinement Study*. COMSOL Knowledge Base. Accessed: 2026-03-17. 2024. URL: <https://www.comsol.com/support/knowledgebase/1261>.
- [28] COMSOL AB. *COMSOL Multiphysics 6.4: Acoustics Module User's Guide - BEM-FEM Coupling*. Accessed: 2026-03-17. COMSOL AB, 2024. URL: https://doc.comsol.com/6.4/doc/com.comsol.help.aco/aco_ug_pressure.05.149.html.
- [29] COMSOL AB. *Loudspeaker Driver — Frequency-Domain Analysis*. COMSOL Model Library. Accessed: 2026-03-17. 2024. URL: https://www.comsol.com/model/download/1178751/models.aco.loudspeaker_driver.pdf.
- [30] Harry F Olson. "Direct radiator loudspeaker enclosures". In: *The Journal of the Acoustical Society of America* 23.5_Supplement (1951), pp. 623–623.
- [31] CEAD Group. *How do I choose the right nozzle size for large format 3D printing?* Accessed: 2026-03-17. CEAD Group, 2024. URL: <https://knowledge.ceadgroup.com/how-do-i-choose-the-right-nozzle-size-for-large-format-3d-printing>.
- [32] B&C Speakers. *18NW100 - LF Driver 18 inch, 4 Ω*. Accessed: 2026-03-17. B&C Speakers S.p.A., 2026. URL: <https://www.bcspeakers.com/en/products/lf-driver/18/4/18NW100>.
- [33] Mehmet Colakoglu. "Factors effecting internal damping in aluminium". In: *Journal of theoretical and applied mechanics* 42.1 (2004), pp. 95–105.
- [34] Harry Nyquist. "Certain topics in telegraph transmission theory". In: *Proceedings of the IEEE* 90.2 (2002), pp. 280–305.
- [35] Abderrazak Mejdji, Tiago Macarios, Bryce Gardner, Sébastien Chaigne, and Chad Musser. "A Proposed Coupled Vibro-Acoustic and CFD Modeling Solution for Complex-Shaped Mufflers Subject to Mean Flow". In: *DAGA 2015 Nürnberg* (2015).
- [36] COMSOL AB. *Modeling Speaker Drivers in COMSOL Multiphysics® – Lumped Methods*. COMSOL Learning Center. Accessed: 2026-03-17. 2024. URL: <https://www.comsol.com/support/learning-center/course/modeling-speaker-drivers-in-comsol-multiphysics-202>.
- [37] Oon H Yeoh. "Some forms of the strain energy function for rubber". In: *Rubber Chemistry and technology* 66.5 (1993), pp. 754–771.
- [38] COMSOL AB. *Bracket - Eigenfrequency Shape Optimization*. COMSOL Application Library. Accessed: 2026-04-01. 2026. URL: <https://www.comsol.com/model/bracket-eigenfrequency-shape-optimization-118721>.
- [39] Susana Rojas Labanda and Mathias Stolpe. "Benchmarking of optimization methods for topology optimization problems". In: *11th World Congress on Computational Mechanics, 5th European Conference on Computational Mechanics, 6th European Conference on Computational Fluid Dynamics*. 2014.
- [40] COMSOL AB. *COMSOL Multiphysics 6.4: Optimization Module User's Guide*. Accessed: 2026-03-17. COMSOL AB, Burlington, MA, USA, 2024. URL: <https://doc.comsol.com/6.4/doc/com.comsol.help.opt/OptimizationModuleUsersGuide.pdf>.
- [41] Compound Dynamics. *Large Format Additive Manufacturing (3D Printing) vs. Traditional Manufacturing: Choosing the Right Method for Your Product*. Accessed: 2026-03-20. Jan. 2025. URL: <https://www.compounddynamics.com/resources/large-format-additive-manufacturing-vs-traditional-manufacturing-choosing-the-right-method-for-you%7D>.
- [42] Jean Allard and Noureddine Atalla. *Propagation of sound in porous media: modelling sound absorbing materials*. John Wiley & Sons, 2009.
- [43] Lennart Moheit and Steffen Marburg. "Infinite elements and their influence on normal and radiation modes in exterior acoustics". In: *Journal of Computational Acoustics* 25.04 (2017), p. 1650020.
- [44] Karim Hamiche. *Efficient and Accurate Broadband FEM-based Vibro-acoustics (Part 2) – External Acoustics*. Accessed: 2026-03-18. Mar. 2018. URL: <https://blogs.sw.siemens.com/simcenter/efficient-and-accurate-broadband-fem-based-vibro-acoustics-part-2-external-acoustics/>.
- [45] Stephen Kirkup. "The boundary element method in acoustics: A survey". In: *Applied Sciences* 9.8 (2019), p. 1642.
- [46] Wang Peng, Jie Zhang, Meng Shi, Jiang Li, and Shaoyun Guo. "Low-frequency sound insulation optimisation design of membrane-type acoustic metamaterials based on Kriging surrogate model". In: *Materials & Design* 225 (2023), p. 111491.
- [47] Thomas Scelo. "Acoustic Modelling: A Brief Review". In: *New Zealand Acoustics* 16.2 (2003), pp. 14–18. URL: https://acoustics.org.nz/wp-content/uploads/Scelo_T_NZA2003.pdf.
- [48] COMSOL AB. *Pressure Acoustics Analysis via Perfectly Matched Layers*. Accessed: 2026-03-18. COMSOL AB, 2024. URL: <https://www.comsol.com/support/learning-center/course/modeling-pressure-acoustics-92/pressure-acoustics-analysis-via-perfectly-matched-layers-15081>.



Alternative Methods for Modelling Vibroacoustic Exterior Radiation

This Appendix presents several alternative methods for modelling vibroacoustic radiation not explored in the present research. Each method is shortly described, followed by an evaluation against the FEM-PML or FEM/BEM method, depending on which method has closer resemblance.

Infinite elements (FEM variant)

Infinite elements extend a finite element mesh to infinity using special shape functions that decay like outgoing waves (e.g., Legendre modes on a spherical or cylindrical truncation surface). They enforce the Sommerfeld radiation condition asymptotically, allowing a smaller truncated fluid domain than PML while keeping sparse matrices, but require convex outer boundaries and can suffer from poor accuracy for evanescent modes or grazing angles [43].

Compared to PML, infinite elements has similar sparse solver efficiency but often requires fewer DOFs. The method is however less robust for complex geometries or low frequencies.

Absorbing boundary layers (AML) or ABCs

Absorbing layers (AML) apply a local impedance operator on a small convex truncation surface to mimic outgoing waves, often combined with high-order p-FEM for broadband accuracy. They approximate the radiation condition via a Robin-type boundary ($\frac{\delta u}{\delta n} iku = 0$), with higher-order versions capturing spherical spreading and dispersion [44].

This method is highly efficient for broadband exterior problems (2 orders fewer nodes than PML/infinite elements), ideal for convex domains, but is also less general than PML for non-convex or heterogeneous exteriors.

Rayleigh integrals (analytical BEM special case)

Rayleigh integrals compute the exterior field directly from the normal velocity on a planar radiator via the free-space Green's function: $p(x) = \int_S v_n(y) \frac{i\omega\rho e^{ik|x-y|}}{4\pi|x-y|} dS_y$. They exactly satisfy the radiation condition for half-space problems with flat sources, enabling fast modal radiation impedance calculations for plates with arbitrary supports [45].

Compared to BEM, this method is limited to simple (planar) radiators and low to moderate frequencies. It is extremely efficient for vibroacoustics of baffles/panels but not applicable to general 3D scatterers.

Lumped/acoustic network models

Lumped/acoustic network models are low-frequency approximations representing the structure-fluid system via equivalent circuits (impedances, mobilities) or transfer matrices, treating radiation as a piston-like load (analogous to the piston assumption used in the present research), they capture global radiated power and impedance without meshing [46].

This method is ultra-fast for modal analysis or optimisation below coincidence frequency, yet proves inaccurate for detailed fields or high frequencies, unlike FEM/BEM.

High-frequency methods: Ray tracing and radiation modes

Ray-tracing models energy propagation geometrically using a high-frequency approximation, tracing rays from vibrating surfaces with reflection and diffraction for statistical averaging. Acoustic radiation modes (ARMs) are eigenmodes of the radiation operator on the source surface, capturing dominant radiating contributions (where few modes suffice at low to moderate k) for model order reduction or far-field prediction [47].

Ray tracing excels above wavelength \sim radiator size and typically complements FEM/SEA. ARMs reduce DOFs in FEM/BEM hybrid methods for efficient far-field vibroacoustics, but both lose near-field detail.

B

Simulation Results using the FEM-PML Method

This Appendix describes the implementation and preliminary results of the FEM-PML method, applied to the problem considered in the present research of solving the ASI problem for a rectangular reference cabinet. The implementation of the FEM-PML method is discussed, followed by the key findings and observed limitations. Background information on the FEM-PML method is provided in Section 2.1.1.

Implementation and Model Setup

Implementation of the FEM-PML method starts with defining the physics nodes and model setup required in COMSOL. Simultaneously, the parameters for the volumetric mesh and actual PML have to be determined, along with more global parameters and modelling considerations. The model description is kept superficial for brevity. Chapter 3 presents the FEM/BEM method with a similar modelling approach in more depth.

The FEM-PML method is likewise set up with a vibroacoustic model as described in Chapter 2 and shown in Figure 2.1. The cabinet is rectangular to match the shape that is most often seen in loudspeakers. The same driver is chosen as for the FEM/BEM method, the SD270A-88 by Dayton Audio [15].

Physics

The model setup for an ASI problem using FEM-PML requires 2 physics nodes, namely **Pressure Acoustics** and **Structural Mechanics**. Pressure acoustics is applied to the interior and exterior domain representing air. Structural mechanics is applied to the cabinet structure. The discretization of both nodes is set to be **quadratic**, as this is the recommended default for properly discretizing wave pressure [48]. The discretization also dictates the required mesh resolution, with a common recommendation being to use 5-6 elements per wavelength, as introduced in Section 3.6.

Geometry and Domain Definitions

Both cabinet and driver are modelled using SOLIDWORKS. The cabinet features a cut-out on the front panel, or *baffle*, to house the driver. The panels are 10 mm thick and the Helmholtz resonator, when included, is modelled as a tube tuned to a frequency of 36 Hz. Note that preliminary FEM-PML simulations omitted the resonator to simplify the model. The driver is modelled according to drawings provided by the manufacturer, but with several features omitted to lower element-count (see Figure 3.1). The diaphragm itself is modelled as a single surface with no thickness. This is done to simplify the application of excitation to this surface in the pressure acoustics node. Complex deformation modes of the diaphragm are not of interest and modelling this as a solid is not necessary. The cabinet and driver are *not* joined in a single body, their geometries are kept separate so that the physics nodes can be applied accordingly. Note that these preliminary simulations did not utilise symmetry and considered the full geometry.

The assembly of cabinet and driver is imported into COMSOL Multiphysics®. Within the geometry GUI, a sphere is modelled around the assembly to represent the exterior acoustic domain. The size of this sphere should follow the convention presented in Section 2.1.1 and have a radius of 8.08-17.15 m. This nets a domain setup that is prohibitively large for computation on consumer-grade hardware and diminishes any performance gained by not including BEM coupling. To explore whether a smaller sphere suffices for the present research, a radius of 4 m is adopted.

The outer layer of this sphere is thickened to form the PML region. The outermost surface of the sphere is defined as `Exterior Field Calculation` boundary for extracting radiation data. The sphere and cabinet interior form the total acoustic domain. The cabinet itself makes up the solid domain. To this solid domain, a PET-G from the built-in material library is applied as a homogeneous material. A small loss factor is included as a means of modelling damping. The exact material used in the final framework that includes an approximation of the anisotropic nature of 3D printed structures was described in Section 3.4.

Acoustic and Mechanical Excitation

The acoustic and solid domain can be excited in several ways in modelling terms. Here, a simplified approach can be taken. The TS parameters introduced in Section 3.3 are likewise used to construct an a lumped mechanical system. Like in the final implementation, the SPL and impedance data of the driver are used to reconstruct an equivalent normal velocity of the modelled diaphragm to replicate the known response. Mechanical excitation is also applied to the driver mounting surfaces.

Meshing and Study Configuration

The geometry is meshed according using 6 free tetrahedral elements per wavelength. For the solid domain, use of the `normal` preset presented convergence in Section 3.6. The same resolution is used in the FEM-PML method. The mesh generated by the structural definition is of higher resolution than acoustically required and dominates the element-count. The PML sphere is meshed according to the same conventions. The actual PML is meshed using a swept strategy with 8 elements along the thickness, which is in accordance to the recommendations provided by COMSOL [48]. PML performs coordinate stretching and the mesh has to match the stretching direction, which is the motivation for using a swept mesh.

A frequency domain study is carried out for a range of **20-200 Hz**, with a sampling interval of 2 Hz. Due to the relatively high element-count of the sphere, an iterative solver is utilised to forego memory limitations. The solver configuration is akin to those described in the present research, leveraging GMRES iterative solvers with PARADISO direct preconditioning.

B.1. Key Findings and Observed Limitations

The key metrics of on-axis SPL, directivity, and the Q-factor are evaluated for the FEM-PML simulation. Initial testing of individual components showed adequate levels of success in terms of efficiency and numerical accuracy, justifying further exploration of the method. For evaluating the applicability of the final method, fully coupled simulations are considered here. These results include the Helmholtz resonator, since the initial scope of the present research did include it. Also, the FEM/BEM coupling responsible for the Helmholtz response, a known and acknowledged limitation, is not applicable here. Variation of the resonator aided in revealing the numerical inaccuracies that led to the ultimate decision to abandon the FEM-PML method.

While initial testing proved successful, several small discrepancies were observed. These were attributed to numerical noise or modelling artifacts. The difficulty of recognizing these defects was heightened by not having real-world measurement data to validate the simulation results to. The discrepancies manifested most clearly as asymmetric responses from a symmetric system, such as in the directivity or Q-factor.

The three key metrics (on-axis SPL, directivity, and Q-factor) for the fully coupled system are visualised in Figure B.1a, b, and c, respectively. The SPL of the full system is compared to the FEM/BEM data and shows smooth fluctuations, deviating from the FEM/BEM result. The directivity, when evaluated using COMSOLs built-in polar plot, reveals asymmetry at certain frequencies, such as in Figure B.1b. The

Q-factor in itself does not show odd values in this case, but proved incoherent when performing mesh convergence studies. The convergence criterion presented in Section 3.7.1 was never met, resonance frequency locations did not stabilise. Furthermore, the locations of the resonance frequencies do not line up with the FEM/BEM result.




Figure B.1: Responses obtained for the rectangular cabinet using a FEM-PML framework, showing the SPL (a), the directivity using a polar plot (b), and the Q-factor (c)

Inclusion and variation of the Helmholtz resonator confirmed the suspicion of numerical artifacts. The resonator is considered outside of scope for this research, but tests were performed with slight geometrical variations to assess numerical stability. The resonator is designed to lie on the width-wise symmetry plane. Offsetting the resonator to the left and the right of the symmetry plane is expected to yield identical yet mirrored results in all metrics, but this was not the case.

This result further confirmed that the FEM-PML simulation method is inadequate for providing accurate near-field, far-field, and structural responses in the context of this research. Theory regarding the FEM-PML strategy indeed contradicts the setup described here, validating the inaccuracies. Correct setup with a minimally sized PML sphere would increase element count by $8\times$. In a FEM-only setup, the computational cost scales mostly linearly with element count, implying an $8\times$ increase in computation time required, increasing solver time from 20 minutes to 160 minutes for the full 20-200 Hz range, compared to an average time of 45 minutes for the FEM/BEM method.

C

Exploratory Simulation Method using an Acoustic Buffer Domain

In order to lower computation times, some effort was spent on reducing the BEM-solid coupling contribution to the overall ASI problem. This coupling dominates assembly and solver times and is governed by the number of boundary elements that couple the cabinet solid domain to the exterior air domain. Literature and testing shows that both acoustic FEM-BEM coupling and *acoustic* FEM-solid coupling do not suffer the same computational penalty. A method is explored to forego the expensive BEM-solid by implementing an acoustic transitional region, or *buffer domain*.

The core idea is to decouple the BEM-solid interaction by layering different domains that couple more efficiently. To do this, an acoustic FEM domain is added around the cabinet domain, creating two couplings instead of the original BEM-solid coupling:

- Cabinet Solid (FEM) - Buffer Domain (FEM)
- Buffer Domain (FEM) - Exterior Air Domain (BEM)

A visualisation of this buffer domain within context of the full ASI model can be seen in Figure C.1. Note the circular cutout made for the driver diaphragm. This cutout is made to maintain the BEM definition of the diaphragm, being the application of an equivalent normal velocity as acoustic excitation to the domain. Making the diaphragm boundaries internal to the FEM buffer domain would not allow for this normal velocity to be applied since it can only act upon transitional boundaries to any acoustic domain.




Figure C.1: Schematic rendering of the 3D model used for vibroacoustic simulation in COMSOL, including domain definitions and the acoustic buffer domain.

C.1. Key Findings and Observed Limitations

Omission of the expensive BEM-solid coupling indeed lowered computation times initially by approximately 60%. The total element count is increased by the buffer domain as its volume needs to be discretized according to the 6-element-per-wavelength guideline, but the overall computational effort is still lowered despite the increase in DOFs solved for. Preliminary simulations quickly revealed discrepancies between the results obtained using this buffer domain and the standard FEM/BEM method as described in Chapter 3. Figure C.2a shows the SPL response for both methods and Figure C.2 the Q-factor response. Validation of the buffer domain method would be implied through agreement with the known-good FEM/BEM method, yet both responses show deviations. The SPL shows fluctuations from the baseline, predominantly at 84, 118, and 142 Hz, resulting in a mean absolute error of **8.79**. The Q-factor shows similar resonance locations but large fluctuations in value, with the biggest outlier at 186 Hz showing a Q-factor of **243.7**, a **14.8** \times increase over the FEM/BEM result. These results, primarily the SPL response, invalidate the buffer domain method.

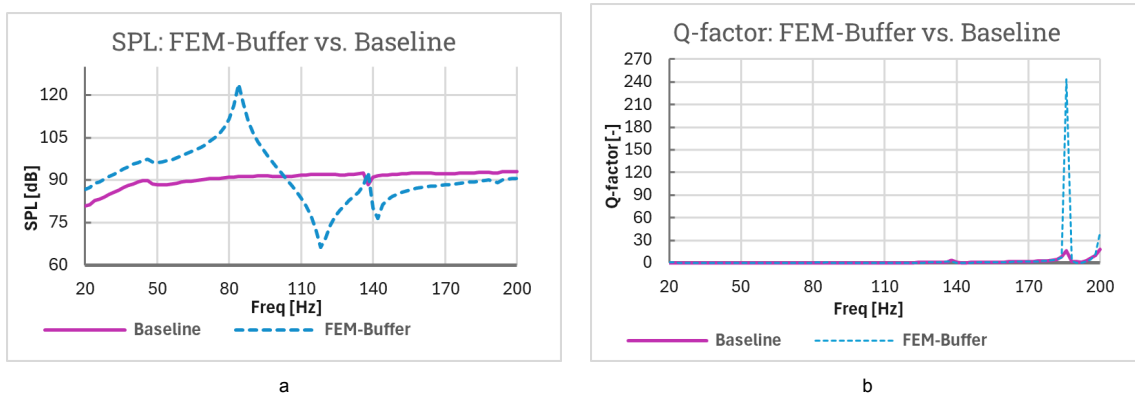


Figure C.2: Comparison in response between the buffer domain method and FEM/BEM method, showing the on-axis SPL (a) and the Q-factor (b)

Upon consideration, it is hypothesised that the buffer domain is analogous to the PML sphere in a FEM-PML method, except that the outer boundaries of the domain are not defined using Exterior Field Calculation but instead through the acoustic FEM-BEM coupling. The same shortcoming then also applies, in that the buffer domains needs adequate size to properly resolve and propagate waves of up to 17.15 m in length. The observed discrepancies may be interpreted as analogous to PML-modes.