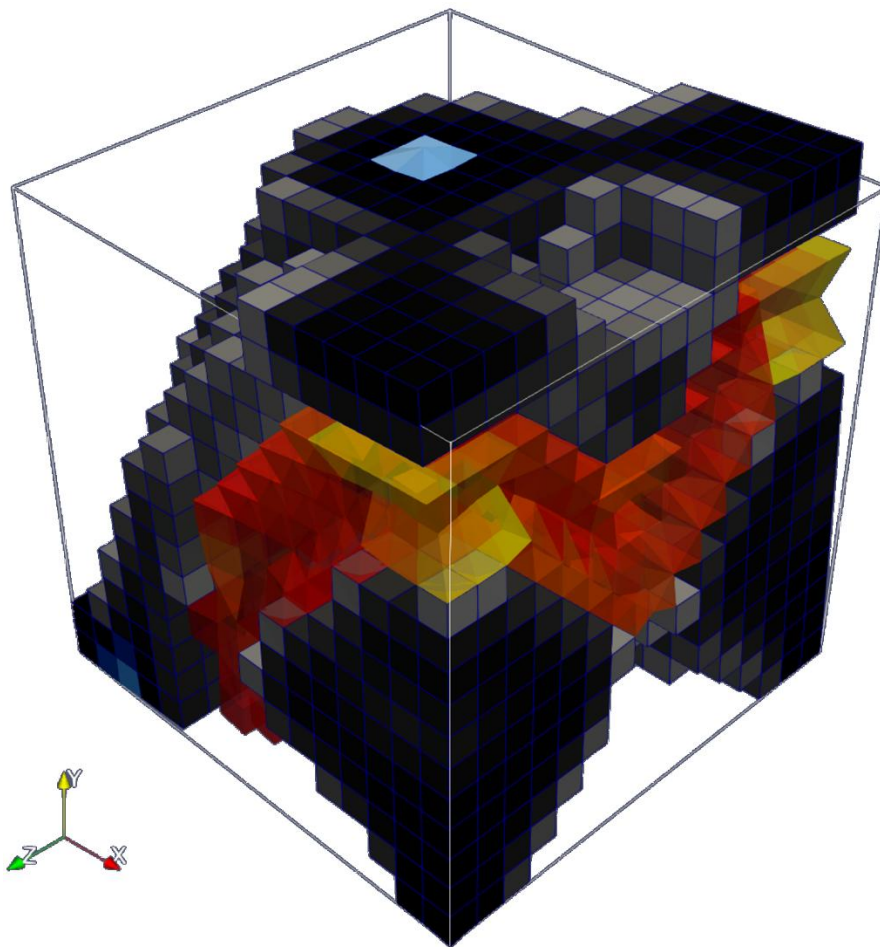# Department of Precision and Microsystems Engineering

Topology Optimisation for Passive Viscoelastic Vibration Isolators Using a Modal Design Approach

Justas Savinijus Petrulis

TUDelft Delft University of Technology

TOPOLOGY OPTIMISATION FOR PASSIVE VISCOELASTIC VIBRATION
ISOLATORS USING A MODAL DESIGN APPROACH

MSc Thesis
Faculty of Mechanical Engineering

Author: Justas Savinijus Petrulis (student ID - 6068189)
TNO Supervisor: Dr. ir. Arnoud Delissen
TU Delft Supervisor: Dr. Jieun Yang
Committee Member: Prof. dr. ir. Matthijs Langelaar

Delft, September 11, 2025

ABSTRACT

Vibration isolation is essential for spacecraft payloads, where both launch loads and on-orbit micro-vibrations can degrade performance. Passive isolators are favoured for their simplicity and reliability but face trade-offs between low-frequency performance, stiffness, and compactness. This thesis introduces a topology optimisation framework for designing passive viscoelastic vibration isolators that distributes aluminium and nearly incompressible rubber within the design domain, moving beyond conventional constrained-layer damping layouts used in literature.

The framework combines a modal design approach with a bi-material finite element formulation using selective reduced integration to handle incompressibility. Modal objectives are defined through rigid-body dominated eigenfrequencies of the combined isolator–payload system, with damping enforced via Q-factor constraints.

Each optimised layout consists of a main load-carrying structure that provides stiffness and smaller damping substructures that create localised regions of energy dissipation. They adapt their shape and location to the inertial properties of the setup and targeted vibration mode, strategically placing viscoelastic material where it contributes most to damping without compromising stiffness. In doing so, the framework enables novel damping mechanisms and tunable multi-directional performance, offering an insight for future passive isolation design in vibration-sensitive applications.

CONTENTS

1 Introduction

1.1 Problem statement

Vibration isolation is essential in modern spacecraft, directly influencing the performance and lifetime of sensitive payloads. Isolation systems protect instruments both during launch and in orbit. During launch, vibrations originate from multiple sources, including rocket motor ignition and shutdown, stage separation shocks, acoustic excitation during lift-off, and shock waves from pyrotechnic device cut-off (Shi et al., 2024). A key structural element in mitigating these loads is the Payload Attachment Fitting, which connects the satellite to the launch vehicle, providing mechanical support while attenuating harmful vibrations.

In orbit, micro-vibration isolation has become increasingly important due to rising requirements for payload stability and pointing precision, with some instruments demanding accuracies as high as 0.03 arcseconds (Shi et al., 2024). Micro-vibrations are low-level mechanical disturbances in a microgravity environment, typically spanning frequencies from below $1\,\mathrm{Hz}$ to $1\,\mathrm{kHz}$ (Zhang et al., 2011). The dominant range for spacecraft is 0.1-300 Hz, where disturbances above $30\,\mathrm{Hz}$ are classified as high frequency and those below as low frequency (Aglietti et al., 2004). In the near-vacuum of Earth's orbit, the absence of environmental damping allows such vibrations to persist, degrading the performance of precision instruments (Luo et al., 2023). Common sources include rotating components such as reaction wheels, control moment gyros, propulsion thrusters, solar array drive mechanisms, and movable mirrors (Jafari, 2018; Shi et al., 2024), as well as micro-thrusters and fluid flow within thermal control systems (Kwon et al., 2017). Since these subsystems are integral to spacecraft operation, their vibration output cannot be eliminated. Therefore, vibration isolation will remain a necessary component for ensuring the stability and accuracy of high-precision instruments in space.
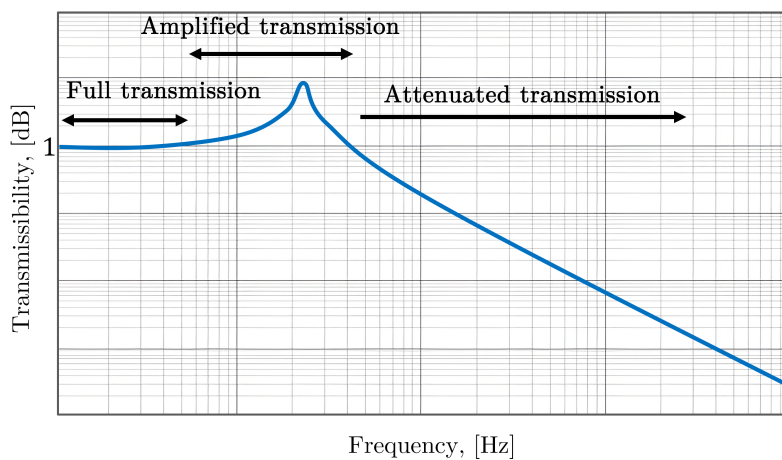
1.2 Background on vibration isolation

Figure 1: Typical transmissibility curve for a single-degree-of-freedom system.

A vibration isolator is a mechanical system designed to decouple a vibrating structure from the components mounted on it, reducing the transmission of dynamic loads to the payload by modifying the combined stiffness, damping, and mass properties of the system. These parameters influence the natural frequency and overall dynamic behaviour, with stiffness and mass defining the resonance point and damping controlling the magnitude of the response. The performance of a vibration isolator is often characterised by the transmissibility $T(\omega)$, defined as the ratio of the response amplitude of the payload to the base excitation amplitude at frequency ω. For an ideal single-degree-of-freedom system, isolation occurs when the excitation frequency exceeds the natural frequency by a sufficient margin. Below this point, transmissibility is dominated by static stiffness, and near resonance, the response is amplified. The curve in Figure 1 illustrates this typical behaviour, indicating regions of full, amplified, and attenuated transmission. Effective isolation begins once the unity transmissibility level is surpassed, with the slope of attenuation determined by the damping properties.
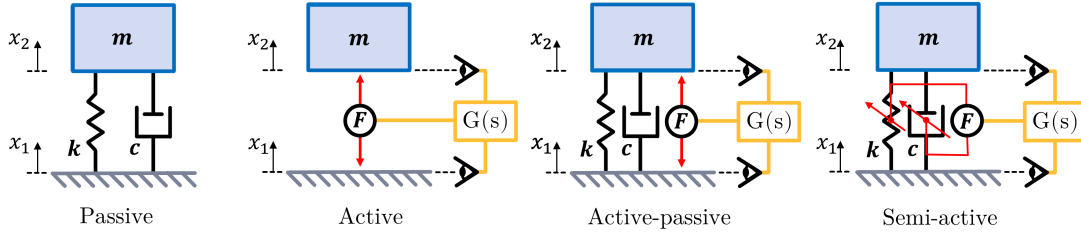
Figure 2: There are four main types of vibration isolators, all of which can be modelled as passive mass–spring–damper systems or with artificial stiffness and damping using a control loop.

Isolation strategies are generally classified as passive, active, semi-active, or active-passive, as illustrated in Figure 2. Active isolators rely on sensors, actuators, and control loops to counteract incoming vibrations, while active-passive systems combine such control with passive mechanical elements to improve robustness. These approaches can adapt to varying excitation conditions and provide effective attenuation at low frequencies, but they require continuous power and add considerable complexity, mass, and potential failure modes (Liu et al., 2015; Shi et al., 2024). Semi-active isolation represents a middle ground between active and passive techniques. By using materials or mechanisms that adjust their properties in response to external stimuli, semi-active systems offer tunable isolation performance with reduced energy consumption. Importantly, they can revert to a passive mode in the event of control failure, providing a fail-safe fallback (Jafari, 2018). Despite these advantages, they still face challenges in low-frequency performance and bring additional cost and system complexity (Liu et al., 2015).

In contrast, passive systems rely only on mechanical properties such as springs, viscoelastic elements, or pneumatic mounts (Shi et al., 2024; Liu et al., 2015). They are valued in space applications for their simplicity, reliability, and zero power consumption, though their performance is often limited at very low frequencies. Existing passive isolators face trade-offs between low-frequency isolation and static load-bearing capacity. They are also sensitive to temperature variations, which can affect performance in orbit. In spacecraft applications, three prominent groups

of passive isolators can be identified: those that leverage the viscous properties of materials, those that employ complex arrangements of mechanical springs, and those that utilise smart materials to dampen vibrations. Examples of these systems are shown in Figure 3, with further discussion of their strengths and drawbacks provided in Appendix A. For a broader overview of vibration isolation technologies in spacecraft, readers are referred to recent review articles by Shi et al. (2024) and Liu et al. (2015).
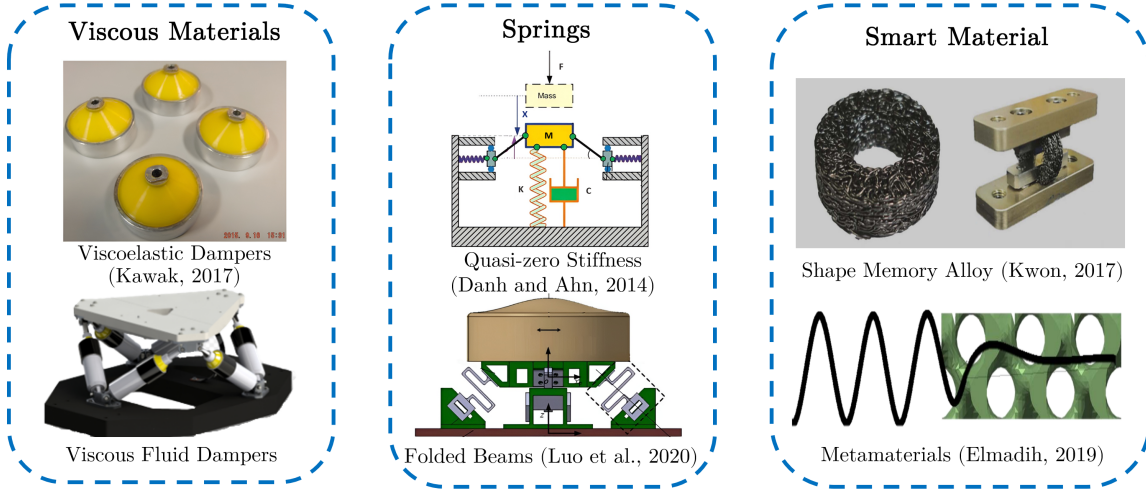
Figure 3: Three groups of passive vibration isolators researched for space applications: viscous-material based, mechanical spring arrangements, and smart-material systems (Kawak, 2017; Danh and Ahn, 2014; Luo et al., 2020; Kwon et al., 2017; Elmadih, 2019).

In spacecraft applications, passive isolation remains the most common approach due to its reliability and predictable behaviour in the space environment. Nevertheless, the demand for compact, multi-directional, and thermally stable isolators is driving interest in new structural concepts and material combinations. Within this context, this project focuses on the development of novel passive isolators. One promising research direction is the application of topology optimisation to vibration isolation. By tailoring both geometry and material distribution, topology-optimised designs can create lightweight, mission-specific isolators that combine the compliance of mechanical spring systems with the damping capability of viscoelastic materials. The following subsection 1.3 reviews the current state of the art in dynamic topology optimisation, the role of viscoelastic materials in vibration control, and prior work combining these approaches.

1.3 State of the art

Improving the vibration performance of a structure through changes to its geometry has long been of interest to engineers and scientists. Structural designers typically focus on meeting strength and stiffness requirements, while control engineers are more concerned with the dynamic characteristics of the structures they control (Ou and Kikuchi, 1996). Beyond controllability, dynamic response relates to the ability to withstand time-varying loads, prevent fatigue, and avoid resonance that can cause failure. In vibration-sensitive applications such as aerospace systems or pre-

cision manufacturing equipment, the ability to manage these characteristics is highly desirable. Designing structures that can effectively isolate vibrations is an important step toward improving performance and reliability in such systems. However, dynamic performance is more complex to characterise than static behaviour, and various indicators are used for its evaluation, including natural frequencies, frequency response, dynamic compliance, power flow response, and modal loss factor (Niu et al., 2018).

Topology optimisation provides a systematic approach for generating material layouts tailored to such performance objectives under defined constraints (Bendsøe and Sigmund, 2003; Sigmund and Maute, 2013). While early applications were primarily static, vibration-related topology optimisation has developed rapidly since the first study by Díaaz and Kikuchi (1992), which used the homogenisation method to maximise the fundamental eigenfrequency of a structure. Early work employed homogenisation and Evolutionary Structural Optimisation, with a key advance made by Pedersen (2000), who proposed a method to eliminate spurious eigenmodes, enabling broader use of the SIMP method for dynamic problems. Today, SIMP remains the most widely used approach, though alternative methods have been used and are well documented in the review by Zargham et al. (2016). These developments have created a mature set of methods for tailoring dynamic behaviour in purely elastic structures. However, extending such methods to systems that include viscoelastic materials introduces additional challenges, as their frequency- and temperature-dependent properties demand more advanced modelling and may alter the formulation of the optimisation problem. Previous research has addressed these challenges mainly through the optimisation of damping layer layouts, with fewer works considering viscoelastic materials as part of the structural domain. The following review first outlines the main developments in dynamic topology optimisation, then discusses modelling approaches for viscoelastic materials, and finally reviews work combining these areas.

1.3.1 Work done in dynamic topology optimisation

Dynamic topology optimisation spans a range of objectives, from shaping frequency response and minimising dynamic compliance to enhancing damping and tailoring eigenfrequencies (Zargham et al., 2016). This review concentrates on eigenfrequency based formulations because natural frequencies provide a compact description of dynamic behaviour, strongly influence low frequency response, and correlate with structural stiffness (Zhu, 2019; Bendsøe and Sigmund, 2003). Within this focus, the literature divides into eigenvalue maximisation and frequency gap formulations, alongside methods for addressing associated numerical issues.

Eigenvalue maximisation   Natural frequencies are a compact descriptor of a structure's dynamic behaviour, with the fundamental eigenfrequency strongly influencing low-frequency response (Zhu, 2019). Maximising this quantity is a well-established objective in topology optimisation (Bendsøe and Sigmund, 2003) and typically results in designs that are stiff both statically and dynamically. The first vibration-related SIMP study, by Pedersen (2000), targeted the fundamen-

tal eigenfrequency in a MEMS device. An alternative formulation, the mean-eigenvalue method introduced by Ma et al. (1995), optimises a weighted combination of several consecutive eigenfrequencies. This approach produces a more balanced improvement across the spectrum but yields smaller gains in the fundamental value (Du and Olhoff, 2007). Multi-objective formulations have also combined eigenvalue maximisation with static performance metrics, as in the missile structure design by Luo et al. (2006), where compliance and eigenfrequency were optimised simultaneously using a tolerance multilevel sequence approach. While maximising the fundamental eigenfrequency improves stiffness and controllability, it also broadens the high-transmissibility range below the first resonance, making it unsuitable for low-frequency vibration isolation. In such cases, shaping the distribution of eigenvalues can be more effective.

Frequency gap objectives   Another line of work focuses on increasing the separation between consecutive eigenfrequencies. This reduces the likelihood of resonance within specific frequency bands and has applications in vibration isolation (Lee and Youn, 2004), noise reduction (Zhang and Kang, 2013), energy harvesting (Noh and Yoon, 2012), and phononic crystal design (Yi and Youn, 2016). Jensen and Pedersen (2006) and Du and Olhoff (2005) developed formulations for maximising either the difference or the ratio between adjacent eigenfrequencies, with sensitivity analysis tailored for these objectives. Later, Du and Olhoff (2007) extended the approach to bi-material designs. More recently, Li et al. (2021) introduced differentiable frequency band constraints using a modified Heaviside function, enabling gradient-based optimisation while avoiding resonance in multiple specified bands. This formulation is objective-independent, allowing it to be applied alongside compliance, eigenfrequency, or volume constraints.

Numerical challenges   Dynamic topology optimisation is prone to two recurring numerical issues: localised eigenmodes and mode switching. Localised eigenmodes occur in low-density regions where stiffness-to-mass ratios become very small, producing spurious modes that distort the global response (Pedersen, 2000). Penalisation adjustments or replacing voids with weak material (Li et al., 2021) can mitigate this. Mode switching, where eigenvalues change order during optimisation, leads to discontinuities in sensitivities and instability in the solution process. The bound formulation (Olhoff, 1989) is widely used to address this, while weighted-sum eigenvalue objectives (Ma et al., 1995) provide an alternative. Both have been extended to multi-mode and frequency-gap problems (Jensen and Pedersen, 2006; Du and Olhoff, 2007).

1.3.2   Interlude: commonly used damping treatments

Practical damping applications of viscoelastic materials (VEMs) in structures are most commonly implemented as surface treatments: free-layer damping, where a VEM is bonded to the host surface and dissipates energy mainly through extensional deformation, and constrained-layer damping (CLD), where a stiff constraining layer forces shear in the VEM core and achieves high damping with thin layers (Nakra, 1998; Moreira, 2014). CLD generally offers superior material

efficiency but requires models that resolve the shear strain field in the viscoelastic layer (Moreira, 2014; Elmoghazy et al., 2024). These layouts, commonly referred to as *sandwich* layouts in the literature, have been widely used since the late 20[th] century and remain the best-established damping treatments (Johnson and Kienholz, 1982). Figure 4 shows the two classical layouts.
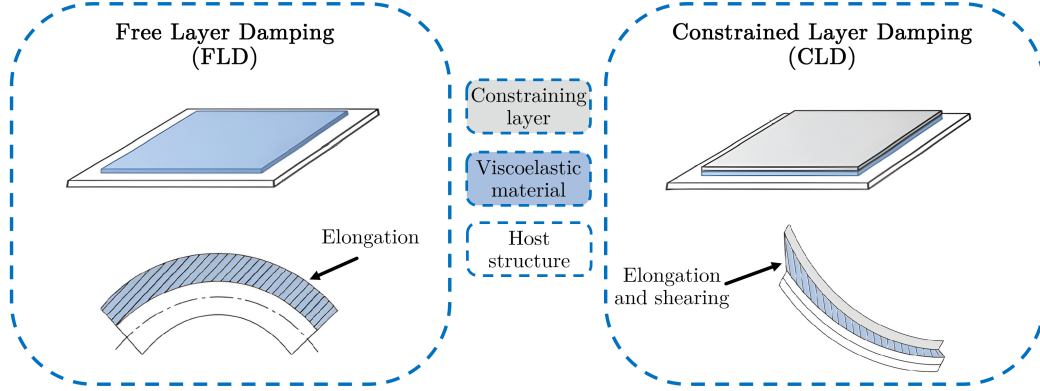
Figure 4: Basic configurations for viscoelastic damping treatments: Free Layer Damping (FLD) and Constrained Layer Damping (CLD). Adapted from Nakra (1998) and Moreira (2014).

1.3.3 Topology optimisation of viscoelastic damping treatments

Early work on free- and constrained-layer damping used linear viscoelastic models to place damping material where it is most effective. For free-layer damping on shells under harmonic loading, Kang et al. (2012) reduced response amplitudes by optimising patch layouts in a reduced modal space. Kim et al. (2013) then maximised modal loss factors and validated the predictions experimentally. For constrained-layer damped plates, Fang and Zheng (2015) minimised peak resonant displacement and showed that updating the modal damping ratio during optimisation improves both damping and resonance suppression. Using sound radiation modes and modal strain energy, Zheng et al. (2016) further demonstrated reductions in radiated sound power with good experiment–simulation agreement. Together, these studies established that *sandwich*-type layouts are effective when the VEM is confined to surface treatments.

> **Note on viscoelastic material properties**
>
> Viscoelastic materials combine elastic energy storage with viscous dissipation, and their stiffness and damping vary strongly with excitation frequency and temperature. More details are given in Appendix B. In this thesis, however, a frequency- and temperature-independent viscoelastic model is adopted.

More recent work has increasingly allowed the viscoelastic properties to vary with frequency and/or temperature. Within CLD, Liu et al. (2018) co-optimised microstructure and layout, showing the well-known trade-off that greater damping typically reduces natural frequencies. In the eigenvalue-oriented stream, Li et al. (2022) formulated non-linear eigenvalue topology optimisation with frequency-dependent material behaviour, and Wu et al. (2024c) combined that with

frequency-band constraints (Li et al., 2021) and isolated-structure suppression (Wu et al., 2024b), yielding higher fundamental frequencies than those predicted by frequency-independent models. Multi-scale and temperature-aware variants report that natural frequencies are relatively stable across temperature, while modal loss factors are more sensitive to temperature (Wu et al., 2023; Zhang et al., 2022; Wu et al., 2024a).

Only a few studies have moved beyond the conventional strategy of restricting viscoelastic material to damping treatments. Among them, Inozume and Aihara (2021) applied multi-material topology optimisation to design lightweight, highly damped structures for the automotive sector. Zhang and Khandelwal (2019) targeted dissipative structures optimised for damping at finite strains, while Kolk et al. (2017) solved the complex eigenvalue problem to maximise the sum of modal loss factors in freely vibrating structures. These works demonstrate that viscoelastic material can be treated as a structural phase within the design domain, yet such approaches remain rare compared to the dominant *sandwich*-based formulations.

1.4   Research gap

Passive isolation remains attractive for spacecraft because of its reliability and zero power demand. However, existing designs still face persistent trade-offs between low-frequency attenuation, static bearing capacity, temperature sensitivity, and the need for compact, multi-directional performance (Liu et al., 2015; Shi et al., 2024). In parallel, topology optimisation for dynamic behaviour has matured considerably, from eigenvalue-based objectives and gap or band formulations to remedies for numerical artefacts. Yet nearly all damping-oriented studies continue to confine the viscoelastic material to sandwich treatments such as free- or constrained-layer damping (Kang et al., 2012; Kim et al., 2013; Fang and Zheng, 2015; Zheng et al., 2016; Li et al., 2022; Wu et al., 2024c; Wu et al., 2023; Zhang et al., 2022). Only a handful of works have explored free placement of viscoelastic material within the structural domain, and these have largely targeted applications outside vibration isolation. For spacecraft isolators, this remains almost completely unstudied. This motivates further investigation into coupling structural and damping phases directly, before addressing the added complexity of frequency- and temperature-dependent viscoelastic properties.

In short, very few vibration isolators for micro-vibration attenuation have been designed via topology optimisation, and existing studies overwhelmingly remain within sandwich layouts. Treating viscoelastic material as a distributed structural phase for compact, multi-axis isolation is largely unexplored.

1.5 Scope and assumptions

The following describes the scope and modelling assumptions adopted in this thesis for designing passive vibration isolators:

System idealisation.   The sensitive instrument's elastic modes are assumed to lie well above the rigid–body region of the combined isolator–instrument system. The instrument is therefore represented only by its lumped mass and inertia. No dynamic sub-structuring or frequency-response coupling is included.

Dynamics assumptions.   The analysis is restricted to small-amplitude, linear dynamics. Geometric nonlinearities, contact effects, and large-strain viscoelastic behaviour are not considered. Random excitation and broadband loads are also excluded and the focus is on modal response.

Configuration and design space.   A single vibration isolator is placed beneath the instrument bed, with load transfer through a fixed bolted interface. Bolt geometry is prescribed and not optimised. Topologies are studied in both 2D and 3D, using a square (2D) or cubic (3D) design domain with prescribed non-design mounting patches.

Material model.   A bi-material design is considered: aluminium as the structural phase and a damped, nearly incompressible rubber as the viscoelastic phase. Both materials are assumed isotropic. Aluminium is modelled as linear elastic and lightly damped, while the rubber is represented by a frequency-independent complex modulus at reference conditions. Near-incompressibility is handled by selective reduced integration in the finite element model.

Objectives and constraints.   The optimisation targets the rigid-body dominated modes of the combined isolator–payload system. The primary objective is to increase the associated eigenfrequencies, while ensuring sufficient damping by enforcing an upper limit on the modal Q-factor. A volume constraint is applied to control material usage. Mass minimisation is not considered, since only two fixed material phases are used (aluminium and rubber).

Out of scope.   The following effects are not considered in this work: multi-isolator layouts, bolt preload or creep, viscoelastic ageing, gravity unloading in orbit, launch-to-space loads, and manufacturing constraints.

1.6 Thesis structure

The thesis is organised as follows. section 2 introduces the behaviour of elastic and viscoelastic materials, their representation in terms of bulk and shear moduli, and the incorporation of

damping into finite element formulations, while section 3 addresses numerical instabilities from nearly incompressible materials and their mitigation, with a comparative study against COMSOL Multiphysics® presented in section 4. section 5 then outlines the topology optimisation algorithm, and section 6 formulates the modal design approach for rigid-body dominated eigenfrequencies and Q-factors. The optimisation results are reported in section 7, covering design strategies, parameter studies, numerical issues, and combined-objective examples in 2D and 3D. Finally, section 8 summarises the contributions, discusses limitations, and suggests directions for future research, while the appendices provide supplementary reviews, an additional viscoelastic material overview, and FEM validation details.

2 Material behaviour and finite element modelling

Accurate modelling of vibration isolation requires an understanding of both how materials respond under dynamic loading and how these responses are represented in numerical analysis. This chapter introduces the principles of elasticity and damping, with particular emphasis on nearly incompressible damped materials, and explains how these concepts are incorporated into finite element formulations.

2.1 Elasticity in isotropic materials

The elastic behaviour of isotropic materials can be expressed compactly using Hooke's law, which provides the link between stress and strain and underlies the conversion between different elastic constants. Since the mechanical response of such materials is fully described by two independent parameters, these are often chosen from Young's modulus E, Poisson's ratio ν, bulk modulus K, and shear modulus G. A general form of Hooke's law for isotropic materials is:

$$\varepsilon_{ij} = \frac{1+\nu}{E}\sigma_{ij} - \frac{\nu}{E}\sigma_{kk}\delta_{ij} \tag{1}$$

where ε_{ij} is the strain tensor, σ_{ij} the stress tensor, σ_{kk} its trace, and δ_{ij} the Kronecker delta. This expression separates the stress response into volumetric and deviatoric parts, directly relating to the bulk modulus K and shear modulus G.

Depending on the dimensionality and boundary conditions of the problem, different assumptions are applied. In two-dimensional formulations, two idealised cases are used:

- **Plane strain:** Assumes zero strain in the out-of-plane direction. Suitable for analysing thin cross-sections located far from free edges of long structures, where thickness deformation is constrained by surrounding material (e.g. dams, underground tunnels). Because plane strain is derived by constraining one strain component of the three-dimensional law, the conversion relations remain identical to the 3D case.

9

- **Plane stress:** Assumes zero stress in the out-of-plane direction. Suitable for thin structures where stresses through the thickness are negligible (e.g. thin plates or membranes). This assumption simplifies the material response by neglecting one stress component, which leads to modified conversion relations compared with the full 3D case.

These assumptions influence how the elastic constants E, ν and K, G are related. Plane stress produces modified conversion formulas, while plane strain retains the same relations as three-dimensional elasticity. The corresponding relations are summarised in Table 1.

Table 1: Relationships between bulk modulus K, shear modulus G, Young's modulus E, and Poisson's ratio ν for plane stress, plane strain, and full three-dimensional elasticity.

Case	$K(E,\nu)$	$G(E,\nu)$	$E(K,G)$	$\nu(K,G)$
Plane Stress	$\dfrac{E}{2(1-\nu)}$	$\dfrac{E}{2(1+\nu)}$	$\dfrac{4KG}{2K+G}$	$\dfrac{2K-G}{2K+G}$
Plane Strain / 3D	$\dfrac{E}{3(1-2\nu)}$	$\dfrac{E}{2(1+\nu)}$	$\dfrac{9KG}{3K+G}$	$\dfrac{3K-2G}{2(3K+G)}$

Health warning:

In this work, material parameters are initially specified in terms of E and ν and are consistently converted to K and G using the full 3D relations. Plane stress is only applied later in the constitutive matrix (subsection 2.5) when required by the problem, but its modified conversions are not used for material specification.

2.2   Material damping

When subjected to dynamic or cyclic loading, many materials exhibit internal energy dissipation, commonly referred to as material damping. Unlike purely elastic materials that store mechanical energy without loss, or purely viscous materials that dissipate energy as heat, most engineering materials display a combination of both behaviours. This dual character is evident in viscoelastic materials, which combine elastic and viscous behaviours and respond to loading with both immediate elastic deformation and delayed, time-dependent flow.

Material damping arises from various microstructural mechanisms. In metals, for example, energy is lost primarily through the movement of dislocations within the crystal lattice. Ceramics and glasses, on the other hand, are structurally rigid and dissipate very little energy, making them lightly damped. Polymers and rubbers are markedly different: they exhibit significant damping due to the internal friction caused by the movement and rearrangement of molecular chains. Thermoplastics allow for greater mobility of these chains and thus higher damping, while cross-linking in thermosetting plastics restricts motion and leads to slightly lower energy dissipation.

When a damped material is harmonically loaded, its response is typically also harmonic but occurs with a phase lag between the applied stress and resulting strain (Moreira et al., 2010). The stress-

strain curve under cyclic loading of VEMs displays a hysteresis loop, depicted in Figure 5, where the enclosed area indicates the energy dissipated as heat. Due to the viscous properties, the mechanical properties of viscoelastic materials are affected by the deformation rate. As a result, their deformation behaviour at varying rates is depicted by a series of curves instead of a single stress-strain curve (Elmoghazy et al., 2024). This behaviour allows damped rubbers to absorb and dissipate energy, making them particularly valuable in applications requiring vibration damping or dynamic load resistance (Moreira, 2014).
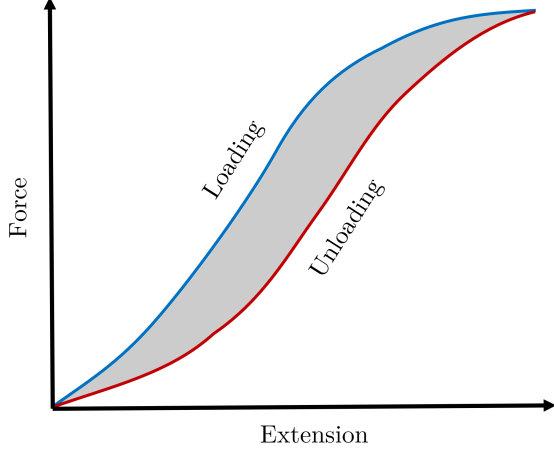

Figure 5: Typical hysteresis loop from cyclic loading of VEMs. (Elmoghazy et al., 2024)
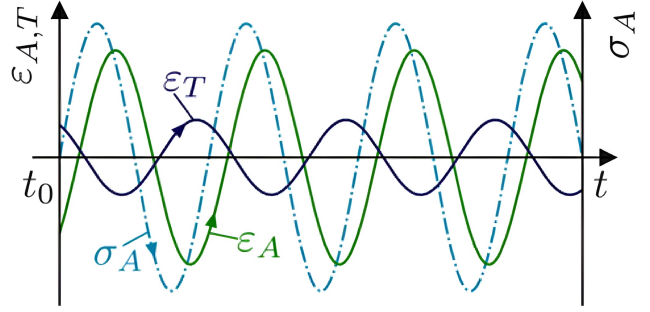
Figure 6: Harmonic uniaxial stress-controlled excitation $\sigma_A(\omega)$ with phase-shifted longitudinal $\epsilon_A(\omega)$ and transversal strain $\epsilon_T(\omega)$ of real materials. (Fauser et al., 2025)

For materials like rubber, which are nearly incompressible, volumetric changes under load are minimal. Damping occurs predominantly through shear deformation, with little to no energy loss through bulk compression. This makes the traditional representation using Young's modulus E and Poisson's ratio ν less suitable. Instead, a more practical and physically meaningful approach is to describe the material in terms of bulk modulus K and shear modulus G. These two parameters clearly distinguish between material's resistance to volumetric deformations, such as compression and expansion, and shearing behaviour. The relationship between these moduli is given by Equation 2.

$$\nu = \frac{3K - 2G}{2(3K + G)} = \frac{3 - 2\frac{G}{K}}{2(3 + \frac{G}{K})} \tag{2}$$

This equation reveals that the Poisson ratio is bounded within the interval $-1 < \nu < 0.5$. Materials with very high G/K ratios, such as designed lattice structures or certain polymer foams, can exhibit Poisson's ratios close to the lower bound. On the other end of the scale, as the ratio G/K tends to zero, the Poisson ratio approaches its theoretical upper limit of 0.5. Materials with very low shear stiffness relative to their bulk stiffness, therefore, approach the incompressibility limit. Rubber is a typical example of such a material. The transverse deformation behaviour under axial loading depends directly on the Poisson ratio. In dynamic applications, the response of incompressible damped rubbers is further complicated by time-dependent phase shifts in both axial and transverse

responses as shown in Figure 6. This phenomenon highlights the need for careful modelling when designing with damping materials (Fauser et al., 2025).

2.3 Complex modulus and damping representation in structural dynamics

The dynamic behaviour of damped systems is traditionally described by a second-order differential equation of motion:

$$\mathbf{M}\ddot{x} + \mathbf{C}\dot{x} + \mathbf{K}x = \mathbf{F}(t) \tag{3}$$

where $\mathbf{M}$, $\mathbf{C}$, and $\mathbf{K}$ are the mass, damping, and stiffness matrices respectively, and x is the displacement vector, and $\mathbf{F}(t)$ is the force vector. In this form, damping is introduced through the matrix $\mathbf{C}$, which is often assumed to represent viscous damping. However, for materials that exhibit internal energy dissipation, it is common to model damping as a material property rather than as a separate viscous force (Fauser et al., 2025). In such cases, where a single material is used, damping can be assumed to be proportional to stiffness, so that $\mathbf{C} = \eta\mathbf{K}$, where η is a dimensionless loss factor that quantifies the relative amount of energy dissipated per cycle. Substituting this into Equation 3 gives:

$$\mathbf{M}\ddot{x} + \mathbf{K}(1 + i\eta)x = \mathbf{0} \tag{4}$$

This formulation allows the damping effects to be condensed into a complex-valued stiffness matrix $\mathbf{K}^*$, which simplifies the analysis of frequency-dependent behaviour. This matrix can be constructed using the complex moduli of elasticity $\{E^*, K^*, G^*\}$, each of which satisfies linear viscoelastic theory for pure extension with the complex stress-strain relationship at the material level:

$$\sigma = E^*\varepsilon = \left(E' + iE''\right)\varepsilon = E'(1 + i\eta)\varepsilon \tag{5}$$

Here, E' is the storage modulus, representing the material's ability to store elastic (retrievable) energy, while E'' is the loss modulus, representing energy dissipation through internal friction. The ratio $\eta = \frac{E''}{E'}$ is again the loss factor. Many microstructural processes that contribute to material damping are not well modelled by viscous damping, and so the damping ratio varies with frequency, however it is assumed constant in the simplified material damping model used throughout this work. The complex modulus representation naturally captures the lag between stress and strain under harmonic excitation and remains consistent with the energy loss observed in the hysteresis loops of viscoelastic materials.

2.4 Generalised complex eigenfrequency problem and damping metrics

In the context of finite element simulations, describing dynamical systems as presented in Equation 4 using a complex stiffness matrix $\mathbf{K}^*$ enables modal or frequency-domain analysis with damping included naturally. In particular, when solving the complex generalised eigenvalue problem,

$$(\mathbf{K}^* - \lambda_n^2 \mathbf{M})\boldsymbol{\phi}_n = 0, \tag{6}$$

the eigenvectors $\boldsymbol{\phi}_n$ and eigenvalues λ_n of the n^{th} mode of vibration are complex. The eigenvectors $\boldsymbol{\phi}_n$ represent complex-valued mode shapes, which describe the spatial distribution and phase of displacements within the vibrating structure. In addition, the eigenvalues provide information about its resonance frequency and modal damping. Each eigenvalue can be expressed as

$$\lambda_n = \omega_n (1 + 2i\zeta_n) = \omega_{n,\text{real}} + i\omega_{n,\text{imag}}, \tag{7}$$

where ζ_n is the modal damping ratio, and $\omega_{n,\text{real}}$ and $\omega_{n,\text{imag}}$ are the real and imaginary parts of the damped natural frequency respectively. The damping ratio ζ_n quantifies the rate of modal decay and relates to energy lost per cycle. The damping ratio ζ_n, modal loss factor μ_n, and quality factor Q_n are interrelated as summarised in Table 2 and are commonly used to quantify the dynamic response of damped systems. The loss factor μ_n is particularly useful in frequency-domain applications, as it directly corresponds to the energy dissipated relative to stored energy over a radian of oscillation. Meanwhile, the quality factor Q_n gives an intuitive measure of how sharply a system resonates, with higher values indicating narrower resonance peaks and lower damping.

Table 2: Common damping metrics in vibration and modal analysis

Property	Symbol	Expression	Interpretation
Damping ratio	ζ_n	$\frac{\omega_{n,\text{imag}}}{\omega_n}$	Time-domain decay rate
Modal loss factor	μ_n	$\frac{2\omega_{n,\text{imag}}}{\omega_{n,\text{real}}}$	Energy loss per radian
Quality factor	Q_n	$\frac{1}{\mu_n}$	Sharpness of resonance

2.5 Elasticity matrices and constitutive relations in FEM

In finite element analysis, the relationship between stress and strain is written compactly in Voigt notation as

$$\boldsymbol{\sigma} = \mathbf{D}\,\boldsymbol{\varepsilon}, \tag{8}$$

where $\boldsymbol{\sigma}$ and $\boldsymbol{\varepsilon}$ are the stress and strain vectors, and $\mathbf{D}$ is the elasticity matrix derived from Hooke's law. Its explicit form depends on the kinematic assumption: plane strain, plane stress, or full three-dimensional elasticity.

For two-dimensional problems, the constitutive relations are shown in Equations 9 to 10. In each case, part (a) gives the complete stress–strain relation in terms of (E, ν), while part (b) shows the equivalent elasticity matrix expressed in terms of (K, G).

Plane strain

$$\begin{bmatrix} \sigma_{11} \\ \sigma_{22} \\ \sigma_{12} \end{bmatrix} = \underbrace{\frac{E}{(1+\nu)(1-2\nu)} \begin{bmatrix} 1-\nu & \nu & 0 \\ \nu & 1-\nu & 0 \\ 0 & 0 & \frac{1-2\nu}{2} \end{bmatrix}}_{\mathbf{D}_{\mathrm{ps}}(E,\nu)} \begin{bmatrix} \epsilon_{11} \\ \epsilon_{22} \\ 2\epsilon_{12} \end{bmatrix} \quad (9a) \qquad \mathbf{D}_{\mathrm{ps}}(K,G) = \begin{bmatrix} K+\frac{4}{3}G & K-\frac{2}{3}G & 0 \\ K-\frac{2}{3}G & K+\frac{4}{3}G & 0 \\ 0 & 0 & G \end{bmatrix} \quad (9b)$$

Plane stress

$$\begin{bmatrix} \sigma_{11} \\ \sigma_{22} \\ \sigma_{12} \end{bmatrix} = \underbrace{\frac{E}{1-\nu^2} \begin{bmatrix} 1 & \nu & 0 \\ \nu & 1 & 0 \\ 0 & 0 & \frac{1-\nu}{2} \end{bmatrix}}_{\mathbf{D}_{\mathrm{pst}}(E,\nu)} \begin{bmatrix} \epsilon_{11} \\ \epsilon_{22} \\ 2\epsilon_{12} \end{bmatrix} \quad (10a) \qquad \mathbf{D}_{\mathrm{pst}}(K,G) = \begin{bmatrix} K+G & K-G & 0 \\ K-G & K+G & 0 \\ 0 & 0 & G \end{bmatrix} \quad (10b)$$

The full three-dimensional elasticity matrix follows the same constitutive relation (Equation 8), but extends to six stress–strain components in Voigt notation:

Full 3D elasticity

$$\begin{bmatrix} \sigma_{11} \\ \sigma_{22} \\ \sigma_{33} \\ \sigma_{12} \\ \sigma_{23} \\ \sigma_{31} \end{bmatrix} = \underbrace{\frac{E}{(1+\nu)(1-2\nu)} \begin{bmatrix} 1-\nu & \nu & \nu & 0 & 0 & 0 \\ \nu & 1-\nu & \nu & 0 & 0 & 0 \\ \nu & \nu & 1-\nu & 0 & 0 & 0 \\ 0 & 0 & 0 & \frac{1-2\nu}{2} & 0 & 0 \\ 0 & 0 & 0 & 0 & \frac{1-2\nu}{2} & 0 \\ 0 & 0 & 0 & 0 & 0 & \frac{1-2\nu}{2} \end{bmatrix}}_{\mathbf{D}_{3D}(E,\nu)} \begin{bmatrix} \epsilon_{11} \\ \epsilon_{22} \\ \epsilon_{33} \\ 2\epsilon_{12} \\ 2\epsilon_{23} \\ 2\epsilon_{31} \end{bmatrix} \tag{11}$$

with the equivalent (K, G) form

$$\mathbf{D}_{3\mathrm{D}}(K, G) = \begin{bmatrix} K + \frac{4}{3}G & K - \frac{2}{3}G & K - \frac{2}{3}G & 0 & 0 & 0 \\ K - \frac{2}{3}G & K + \frac{4}{3}G & K - \frac{2}{3}G & 0 & 0 & 0 \\ K - \frac{2}{3}G & K - \frac{2}{3}G & K + \frac{4}{3}G & 0 & 0 & 0 \\ 0 & 0 & 0 & G & 0 & 0 \\ 0 & 0 & 0 & 0 & G & 0 \\ 0 & 0 & 0 & 0 & 0 & G \end{bmatrix}. \tag{12}$$

2.6 Application of the elasticity matrix in FEM

In finite element formulations, the elasticity matrix $\mathbf{D}$ determines how each element resists deformation by linking local strains to stresses through Hooke's law. The stiffness contribution of a finite element is given by

$$\mathbf{K}_e = \int_{\Omega_e} \mathbf{B}^T \mathbf{D} \mathbf{B} \, dV \tag{13}$$

where $\mathbf{K}^{(e)}$ is the element stiffness matrix, $\mathbf{B}$ is the strain-displacement matrix derived from the element shape functions and geometry, and $\mathbf{D}$ is the constitutive matrix defined earlier in subsection 2.5. This expression follows from the principle of virtual work and is standard in linear elasticity-based FEM. For detailed derivations, see Zienkiewicz et al. (2010) or Bower (2010).

In practice, this volume integral is evaluated numerically using Gauss quadrature. For linear quadrilateral (2D) and brick (3D) elements, integration is performed over a reference coordinate domain $\{\xi_1, \xi_2, \xi_3\}$. In this domain, each element is mapped to a standard square or cube, bounded within $[-1, 1]^2$ or $[-1, 1]^3$, as illustrated in Figure 7. The general quadrature rule in 2D is:

$$\int_{-1}^{+1} \int_{-1}^{+1} f(\xi_1, \xi_2) \, d\xi_1 d\xi_2 \approx \sum_{i=1}^{N} \sum_{j=1}^{N} w_i w_j f\left(\xi_1^i, \xi_2^j\right) \tag{14}$$

where N is the number of Gauss points used in each coordinate direction. In practice, linear elements typically require only one or two points per direction, as higher-order integration (e.g. three points) quickly increases computational cost, especially in 3D. For 3D problems, the same rule extends by evaluating an additional sum over the third dimension. Applying this quadrature rule to the integral in Equation 13, the element stiffness matrix is approximated as:

$$\mathbf{K}_e \approx \sum_{q=1}^{n_{\mathrm{gp}}} w_q \mathbf{B}^T(\xi^q) \mathbf{D} \mathbf{B}(\xi^q) \det J(\xi^q) \tag{15}$$

where ξ^q denotes the Gauss integration points (in coordinates ξ_1, ξ_2, and ξ_3), w_q are the corresponding weights, and $\det J(\xi_q)$ is the determinant of the Jacobian matrix that transforms reference coordinates to the physical element geometry. This process results in a stiffness matrix $\mathbf{K}_e$

of size $n_{\text{dof}} \times n_{\text{dof}}$, where n_{dof} is the number of displacement degrees of freedom associated with the element's nodes.

(a) 2D Node Numbering

(b) 3D Node Numbering
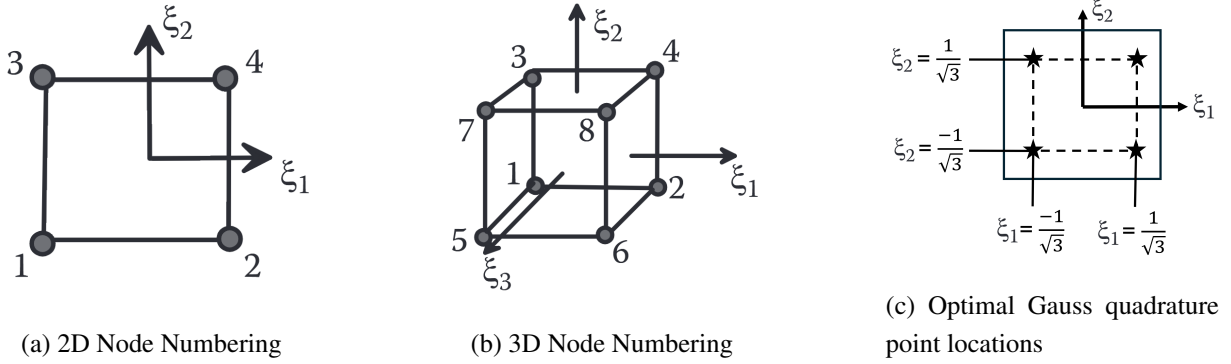
(c) Optimal Gauss quadrature point locations

Figure 7: Illustrations of node numbering in 2D and 3D elements, and integration points for Gaussian quadrature.

Once each $\mathbf{K}_e$ is computed, the global stiffness matrix is assembled using standard finite element procedures. Each element's contribution is inserted at positions determined by the node connectivity. Readers unfamiliar with the details of global assembly, shape function construction, or matrix indexing are referred to standard texts such as Zienkiewicz et al. (2010), Cook and Cook (2002), and Bower (2010).

3 Numerical instabilities in modelling incompressible materials

Modelling incompressible viscoelastic rubber using standard linear finite elements is likely to encounter numerical challenges. Such challenges are often referred to as locking phenomena. Finite elements are said to lock when they present an artificially stiff response to deformation, causing significant inaccuracies and poor convergence in numerical models (Bower, 2010). In general, there are two predominant types of numerical locking: shear locking and volumetric locking.

3.1 Shear locking and solutions

Shear locking occurs primarily in finite elements subjected to bending loads, especially noticeable in thin, elongated structures. Linear finite elements are limited by their straight edges, which prevent them from accurately approximating the curved deformation fields found in pure bending conditions. As a result, these elements exhibit artificially large shear strains rather than the expected bending deformation. This spurious shear strain significantly increases the element stiffness, causing the numerical model to underestimate deflections (Bower, 2010; ABAQUS, 2004).

This behaviour can be analysed in more detail by observing the deformation of a linear and second-order quadrilateral element subjected to pure bending. In the linear element seen in Figure 8, the line connecting the upper integration points elongates, indicating tensile longitudinal

stress (σ_{11}), while the lower line contracts, indicating compression in the same direction. Meanwhile, the vertical lines remain unchanged in length, confirming that the transverse normal stress (σ_{22}) remains zero, as expected for pure bending. However, the originally right-angled intersections between dashed horizontal and vertical lines crossing Gauss integration points become distorted, signalling the emergence of a spurious shear stress component (σ_{12}). The error originates from the inability of linear elements to curve, and the resulting strain energy is incorrectly stored as shear deformation rather than bending. Consequently, the overall stiffness of the element is overestimated, leading to reduced displacements. In contrast, quadratic quadrilateral elements, as shown in Figure 8, are capable of approximating the correct curvature under bending. Their edges can deform smoothly and capture intended bending behaviour without introducing spurious shear stresses. (ABAQUS, 2004)
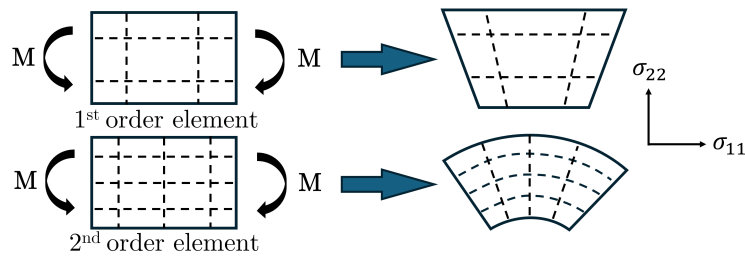
Figure 8: Linear quadrilateral element and quadratic quadrilateral element response under pure bending. Intersections of dashed lines represent Gauss integration points. Adapted from ABAQUS (2004).

Several effective strategies exist to mitigate shear locking:

Second-order elements: Second-order (quadratic) elements are effective as they support higher-order displacement interpolation and curved deformation modes. However, fully integrated, quadratic elements can also lock under complex states of stress (ABAQUS, 2004).

Mesh Refinement: Shear locking can be detected and solved by using mesh refinement and avoiding the use of elongated elements. However, it may significantly increase computational cost if numerous elements are required to achieve the necessary accuracy. This is particularly important for topology optimisation, where the iterative nature of the procedure amplifies the computational cost of using a fine mesh.

Incompatible mode elements: Incompatible mode elements provide additional strain modes specifically designed to capture bending deformation. These elements offer high accuracy at a lower computational cost compared to quadratic elements but require careful consideration as they are sensitive to element distortion. Interested readers are referred to Bower (2010) and ABAQUS (2004) which contain additional details on the method's formulation and example code implementing the method onto quadrilateral elements.

3.2 Volumetric locking and solutions

Volumetric locking arises when modelling nearly incompressible materials using standard displacement-based finite elements. This phenomenon becomes increasingly problematic as the Poisson's ratio ν approaches 0.5, and remedies should be considered at Poisson ratio values as low as 0.475 (ABAQUS, 2004). In this limit, the bulk modulus $K(E, \nu)$ tends towards infinity, as indicated by its expression in Table 1, meaning that even a small volumetric strain leads to disproportionately high hydrostatic stress. The hydrostatic stress (or -pressure) is defined as the negative average of the normal stresses:

$$p = -\frac{1}{3}(\sigma_{11} + \sigma_{22} + \sigma_{33}) = -\frac{1}{3}\operatorname{tr}(\boldsymbol{\sigma}). \tag{16}$$

In linear isotropic elasticity, it is also directly related to the volumetric strain by the constitutive relation:

$$p = -K\,\varepsilon_{\text{vol}}, \tag{17}$$

where $\varepsilon_{\text{vol}} = \varepsilon_{11} + \varepsilon_{22} + \varepsilon_{33} = \operatorname{tr}(\boldsymbol{\varepsilon})$. In a pure displacement formulation, volumetric strain is derived from the displacement field, which is interpolated using shape functions. However, standard interpolation schemes cannot ensure zero volumetric strain at all integration points, especially under full integration. As a result, even small inaccuracies in the displacement field may produce significant errors in volumetric strain. When scaled by a high bulk modulus as shown in Equation 17, these errors manifest as unrealistic pressure fields and overly stiff structural behaviour (ANSYS, 2017). This effect is independent of mesh size and cannot be resolved through refinement alone and specialised finite element formulations must be employed instead (Bower, 2010).

Several effective strategies exist to mitigate volumetric locking:

Reduced integration: Reduced integration is one of the simplest and most widely used techniques for alleviating volumetric locking. The idea is to integrate the element stiffness matrix using a lower-order quadrature rule than is normally required for exact integration. By reducing the number of Gauss integration points, the formulation avoids over-constraining the volumetric strain field, making it easier for the element to accommodate the near-incompressibility condition. This approach is highly effective in higher-order elements, such as quadratic quadrilaterals or bricks, where reduced integration can eliminate volumetric locking entirely and even improve the accuracy of the element (Bower, 2010).

In linear elements, such as four-node quadrilaterals or eight-node bricks, reduced integration leads to a rank-deficient stiffness matrix, enabling non-physical zero-energy deformation modes. This phenomenon is known as *hourglassing* due to the typical shape of the spurious mode. To stabilise the formulation, additional control measures are required. Common approaches include Selective Reduced Integration, B-bar Reduced Integration and Hourglass Control. The first two methods modify the virtual work formulation to separate the stress response into volumetric and deviatoric parts. The volumetric part, which is associated with the large bulk modulus K, is then averaged or integrated with a reduced scheme, while the deviatoric part, associated with the shear modulus

G, maintains full integration. This separation when deriving the finite element stiffness matrix $\mathbf{K}_{ai\,bk}^{(e)}$ can be written as

$$\mathbf{K}_{ai\,bk}^{(e)} = \underbrace{\int_{\Omega_e} \left(\mathcal{C}_{ijkl} \frac{\partial N^a}{\partial x_j} \frac{\partial N^b}{\partial x_l} - \frac{1}{3} \mathcal{C}_{ppkl} \frac{\partial N^a}{\partial x_i} \frac{\partial N^b}{\partial x_l} \right) d\Omega}_{\text{deviatoric part – full integration}} + \underbrace{\int_{\Omega_e} \frac{1}{3} \mathcal{C}_{ppkl} \frac{\partial N^a}{\partial x_i} \frac{\partial N^b}{\partial x_l} d\Omega}_{\text{volumetric part – reduced integration}}. \quad (18)$$

where $\mathcal{C}_{ijkl}$ is the fourth-order elasticity tensor, N^a and N^b are the element shape functions, and repeated indices imply summation. Using the full fourth-order tensor allows this formulation to extend to finite strain problems. Here, however, only small strain elasticity is considered, so the element stiffness reduces to its second-order tensor form $\mathbf{K}_{ab}^{(e)}$. This form was applied in the derivation of the finite element stiffness matrices presented in the thesis. For further theoretical details and implementation examples, the reader is referred to Bower (2010).

Mixed formulation: In standard displacement-based elements, pressure is indirectly determined from the volumetric strain, which in turn is derived from nodal displacements. However, for incompressible or nearly incompressible materials, volume changes are constrained or entirely prohibited, making it inappropriate to calculate pressure based solely on displacement data. As a result, pure displacement formulations become inadequate for such materials (ABAQUS, 2004). Mixed formulation attempts to remedy this by treating hydrostatic pressure of the solid as a primary variable and solve it explicitly using an additional degree of freedom. The displacement field is then used only to evaluate the deviatoric (shear) part of the strain, while the pressure field is interpolated separately and used to enforce the incompressibility constraint. This leads to a coupled system of equations at both element and system matrix level:

$$\begin{bmatrix} K_{uu} & K_{up} \\ K_{pu} & K_{pp} \end{bmatrix} \begin{Bmatrix} \Delta u \\ \Delta p \end{Bmatrix} = \begin{Bmatrix} \Delta F \\ 0 \end{Bmatrix} \quad (19)$$

In this system, the displacement increment Δu governs deviatoric deformation, while the pressure increment Δp applies a weak constraint on the volumetric strain. To ensure numerical stability, the interpolation functions for pressure are of one order lower than the interpolation function of strains or stresses (ANSYS, 2017). In linear elements, pressure is typically assumed to be constant and defined at the element centroid. This corresponds to constant interpolation for pressure and linear interpolation for displacements (Bower, 2010).

Second-order elements: Quadratic elements are generally less prone to volumetric locking because their higher-order interpolation allows them to represent complex deformation patterns, including nearly volume-preserving strains. While second-order elements do not inherently separate pressure and deviatoric responses like other advanced formulations do, their improved strain accuracy helps reduce locking severity by improving the accuracy of both strain and stress predictions throughout the element domain.

4 Finite element investigation on incompressibility

Rubber-like materials used in this thesis are nearly incompressible, which makes their finite element modelling prone to volumetric and shear locking. The topology optimisation framework is built on a custom implementation in PyMOTO (Delissen, 2023), which by default employs standard linear elements (quadrilaterals in 2D and bricks in 3D). While these elements are simple and efficient, they are also known to perform poorly for nearly incompressible materials, and therefore require verification before being applied to viscoelastic rubber components. Without such verification, optimisation outcomes could be dominated by numerical artefacts rather than the intended material response.

To address this, a comparative eigenfrequency analysis was carried out between PyMOTO and COMSOL Multiphysics®. The study examined benchmark models in two and three dimensions, spanning materials from aluminium to highly incompressible damped rubber. In COMSOL, three common locking mitigation techniques were systematically enabled or disabled:

- **reduced integration**

- **mixed formulation**

- **second-order elements**

The insights from this investigation guided the choice of volumetric locking mitigation scheme for subsequent topology optimisation. The following subsection 4.1 describes the simulation setup and the conventions used for comparing results.

4.1 Simulation setup and notation

Benchmark models: The verification was based on two simple benchmark geometries: a square domain in 2D and a cube in 3D. Each model was discretised using a 40×40 mesh in 2D and a $10 \times 10 \times 10$ mesh in 3D. These mesh sizes were chosen to match those employed during topology optimisation in the initial stages of this thesis. Fixed Dirichlet boundary conditions were applied along the bottom edge (2D) or surface (3D), while the remaining boundaries were left free. In PyMOTO, the domains were discretised using standard linear quadrilateral elements in 2D and linear brick elements in 3D.

Material information: Five materials were included in the study: one reference material (aluminium) and four rubber-like materials with Poisson ratios $\nu = \{0.48, 0.49, 0.4999, 0.499999\}$. Their stiffness parameters are listed in Table 3. Aluminium was treated as a compressible ($\nu = 0.3$) lightly damped elastic material with a loss factor ($\eta = 0.01$) applied to both bulk and shear moduli. The rubber materials were modelled using a viscoelastic formulation, with a real-valued bulk modulus and a complex-valued shear modulus following the methodology described earlier

Table 3: List of materials and their properties used for the incompressibility mitigation experiment

Material	E, [Pa]	η, [-]	ν, [-]	K, [Pa]	G, [Pa]
Aluminium	$68.900 \cdot 10^9$	0.01	0.30	$(5.741 + 6.315i) \cdot 10^{11}$	$(2.297 + 2.526i) \cdot 10^{10}$
Rubber$_{(\nu=0.48)}$	$2.462 \cdot 10^7$	1.1	0.48	$2.052 \cdot 10^8$	$(8.318 + 9.150i) \cdot 10^6$
Rubber$_{(\nu=0.49)}$	$2.462 \cdot 10^7$	1.1	0.49	$4.104 \cdot 10^8$	$(8.262 + 9.088i) \cdot 10^6$
Rubber$_{(\nu=0.4999)}$	$2.462 \cdot 10^7$	1.1	0.4999	$4.104 \cdot 10^{10}$	$(8.207 + 9.028i) \cdot 10^6$
Rubber$_{(\nu=0.499999)}$	$2.462 \cdot 10^7$	1.1	0.499999	$4.104 \cdot 10^{12}$	$(8.207 + 9.028i) \cdot 10^6$

in subsection 2.2. The rubber properties used in the simulations were derived from tensile tests conducted at TNO on an aerospace-grade elastomer SCVBR® supplied by ERIKS. These tests were performed at TNO across several temperatures, with the $20\,°C$ data point selected for this study.

Simulation labels: To enable systematic comparison between PyMOTO and COMSOL eigenfrequency results, several conventions and definitions were adopted throughout the investigation. The COMSOL simulations were configured to explore the impact of three commonly used locking mitigation strategies: reduced integration, mixed formulation, and second-order elements. Each of these settings was treated as a binary flag (ON/OFF), resulting in eight distinct combinations per material. For clarity, each configuration is identified using a three-part label. For example, `ON-OFF-OFF` denotes second-order elements enabled, reduced integration disabled, and mixed formulation disabled. The *gold standard* was defined as the case where all three techniques were active (`ON-ON-ON`). In total, approximately 80 simulation cases were completed, allowing for a broad comparison of performance across locking mitigation methods, material types, and dimensional assumptions.

Types of eigenmodes: In 2D, the first three eigenfrequencies were extracted, while six modes were evaluated in 3D, as illustrated in Figure 9. This selection was intentional, since higher-frequency modes typically involve local deformations that are less relevant for assessing vibration isolators with a payload attached at the top boundary. The extracted modes were classified according to their deformation characteristics and were found to maintain a consistent ordering by frequency across all materials and constitutive assumptions. Three fundamental vibration modes were common to both dimensions: *lateral bending*, *axial tension/compression*, and *planar rotation*. Extending the structure from 2D to 3D introduces an additional spatial axis and more axes of symmetry. This both duplicates some of the pre-existing modes and enables new ones to emerge. In 3D, both *lateral bending* and *planar rotation* split into two orthogonal variants due to the additional symmetry directions of the cube. More importantly, the third spatial axis gives rise to a new type of motion: *torsional rotation*, which cannot exist in 2D. This progression illustrates how increasing dimensionality expands the vibrational landscape while retaining equivalence with lower-dimensional analogues.
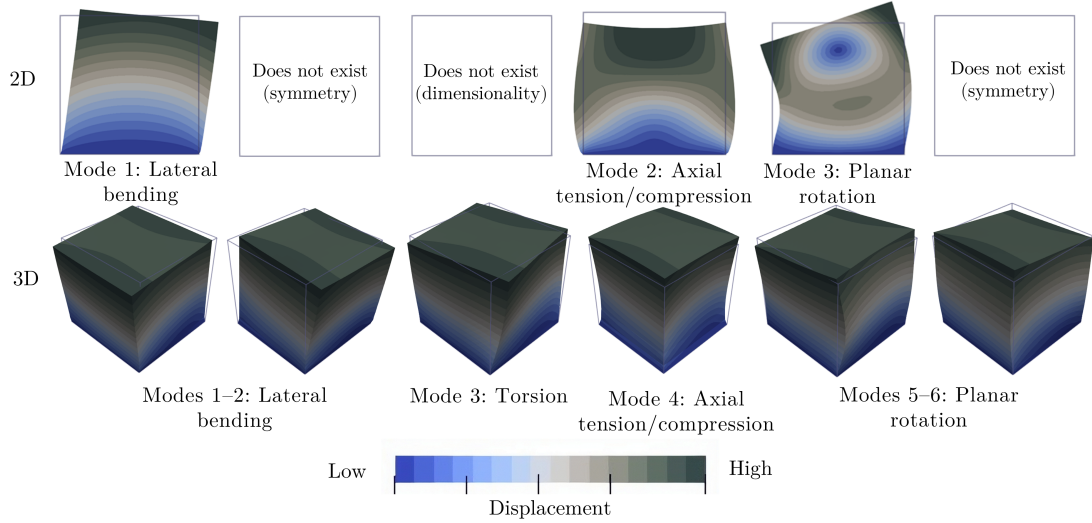
Figure 9: Lowest non-internal vibration modes of the bottom-fixed square (2D) and cube (3D). Three fundamental modes are shared between 2D and 3D (lateral bending, axial tension/compression, planar rotation), while the additional spatial axis in 3D introduces duplicate bending directions and a new torsional rotation mode.

4.2 Results and recommendation for FEM integration

Figure 10 summarises the average errors of six non-internal eigenfrequencies of a 3D cube obtained in PyMOTO and COMSOL across a range of Poisson ratios and locking mitigation settings. Each configuration is labelled using the `ON/OFF` convention introduced in subsection 4.1, and errors are reported relative to the COMSOL *gold standard* (`ON-ON-ON`) reference. Supplementary results for the plane strain assumption are provided in Appendix C. The findings show that all three methods alleviated volumetric locking, but with different levels of accuracy and computational cost:

Reduced integration:

[+] Had the best stand-alone performance of all methods. Achieved near-exact agreement with the COMSOL *gold standard* (`ON-ON-ON`).

[+] Reliable in both 2D and 3D, across all Poisson ratio values.

[+] Introduced no additional degrees of freedom or matrix coupling. As a result, it is the most computationally efficient method.

[+] Required minimal change to the FEM formulation, only element stiffness matrix integration needs to be adjusted.

[-] Different reduced integration implementations are more accurate than others. As seen in Figure 10, a small discrepancy can be seen between the `OFF-ON-OFF` results obtained with PyMOTO and COMSOL. PyMOTO used selective reduced integration, while COMSOL applied a more advanced reduced integration scheme with automatic hourglass control. The latter suppresses spurious zero-energy modes, leading to higher accuracy.

	1st order (OFF)			2nd order (ON)		
	Mixed (OFF)		Mixed (ON)	Mixed (OFF)	Mixed (ON)	
	Python	COMSOL				
Tag	OFF-OFF-OFF	OFF-OFF-OFF	OFF-OFF-ON	ON-OFF-OFF	ON-OFF-ON	
$\nu = 0.3$	0.90%	0.90%	0.62%	0.05%	0.02%	Reduced Integration (OFF)
$\nu = 0.48$	3.43%	3.43%	0.93%	0.50%	0.07%	
$\nu = 0.49$	5.40%	5.40%	0.95%	0.68%	0.07%	
$\nu = 0.4999$	37.55%	37.55%	0.98%	1.08%	0.07%	
$\nu = 0.499999$	42.16%	42.16%	0.98%	1.08%	0.07%	
Tag	OFF-ON-OFF	OFF-ON-OFF	OFF-ON-ON	ON-ON-OFF	ON-ON-ON	
$\nu = 0.3$	0.62%	0.18%	0.18%	0.02%	0.00%	Reduced Intergation (ON)
$\nu = 0.48$	0.93%	0.28%	0.28%	0.11%	0.00%	
$\nu = 0.49$	0.95%	0.29%	0.29%	0.29%	0.00%	
$\nu = 0.4999$	0.98%	0.29%	0.29%	0.12%	0.00%	
$\nu = 0.499999$	0.98%	0.29%	0.29%	0.12%	0.00%	

Average eigenfrequency relative error, [%]

Low   High

Figure 10: Results from 3D investigation on methods to improve volumetric locking.

Mixed formulation:

[+] Removed locking but at a slightly lower accuracy than reduced integration (`OFF-OFF-ON`).

[+] Reliable in both 2D and 3D, across all Poisson ratio values.

[-] Combining with reduced integration (`OFF-ON-ON`) provided no further accuracy gain over using only reduced integration (`OFF-ON-OFF`).

[-] Introduced an additional pressure variable per element, complicating global matrix assembly and increasing computational cost.

Second-order elements:

[+] Best accuracy for compressible material (aluminium). Second-order elements suffer less from discretisation error at small mesh sizes.

[+] Improved accuracy in rubber materials across the full range of Poisson ratios, however by less than other methods and had worse performance at very high Poisson ratios.

[+] In 3D, combining with reduced integration (`ON-ON-OFF`) or mixed formulation (`ON-OFF-ON`) had better agreement with the *gold standard* than any other stand-alone or two-method combination.

[-] Significantly increased the number of degrees of freedom and the global stiffness matrix size. Added computational cost makes them impractical for topology optimisation, where repeated assembly and solution are required.

Recommendation. Selective reduced integration consistently provided the best balance between numerical accuracy and computational efficiency. It performed reliably in both 2D and 3D, even for highly incompressible rubber materials with Poisson ratios $\nu \to 0.5$. Among all tested methods, it was the best stand-alone option in terms of its agreement with the COMSOL *gold standard* configuration (`ON-ON-ON`) while keeping the formulation simple, with no additional degrees of freedom and only a minor change in the element integration routine. Plane strain was adopted as the 2D modelling assumption, since it better represents the boundary conditions of rubber components than plane stress, while full 3D elasticity remains the most physically accurate option. For these reasons, reduced integration was implemented in the PyMOTO FEM framework and applied in all topology optimisation simulations presented in this thesis.

> **Note:**
>
> This study was not intended as a full characterisation of locking phenomena or FEM performance in every context. It was based on simplified benchmarks without mesh refinement or shear-locking analysis, using fixed boundary conditions representative of the optimisation use case. The aim was to identify a material assumption and locking mitigation strategy suitable for integration into the topology optimisation framework. Within this scope, the chosen approach was adequate for the intended purpose.

5 Algorithm overview

Topology optimisation algorithms generally follow a common iterative workflow built around a small set of components. The process begins with the definition of an objective function, constraints, material model, and governing physics. Although implementations differ in detail, most share the same basic structure: a design domain is initialised, densities are updated through optimisation, and equilibrium is enforced at each iteration. Convergence is declared once successive designs change only marginally or 100 iterations take place. The framework used in this thesis builds on this structure, combining filtering and projection, material interpolation, finite element analysis, sensitivity analysis, and iterative updates.
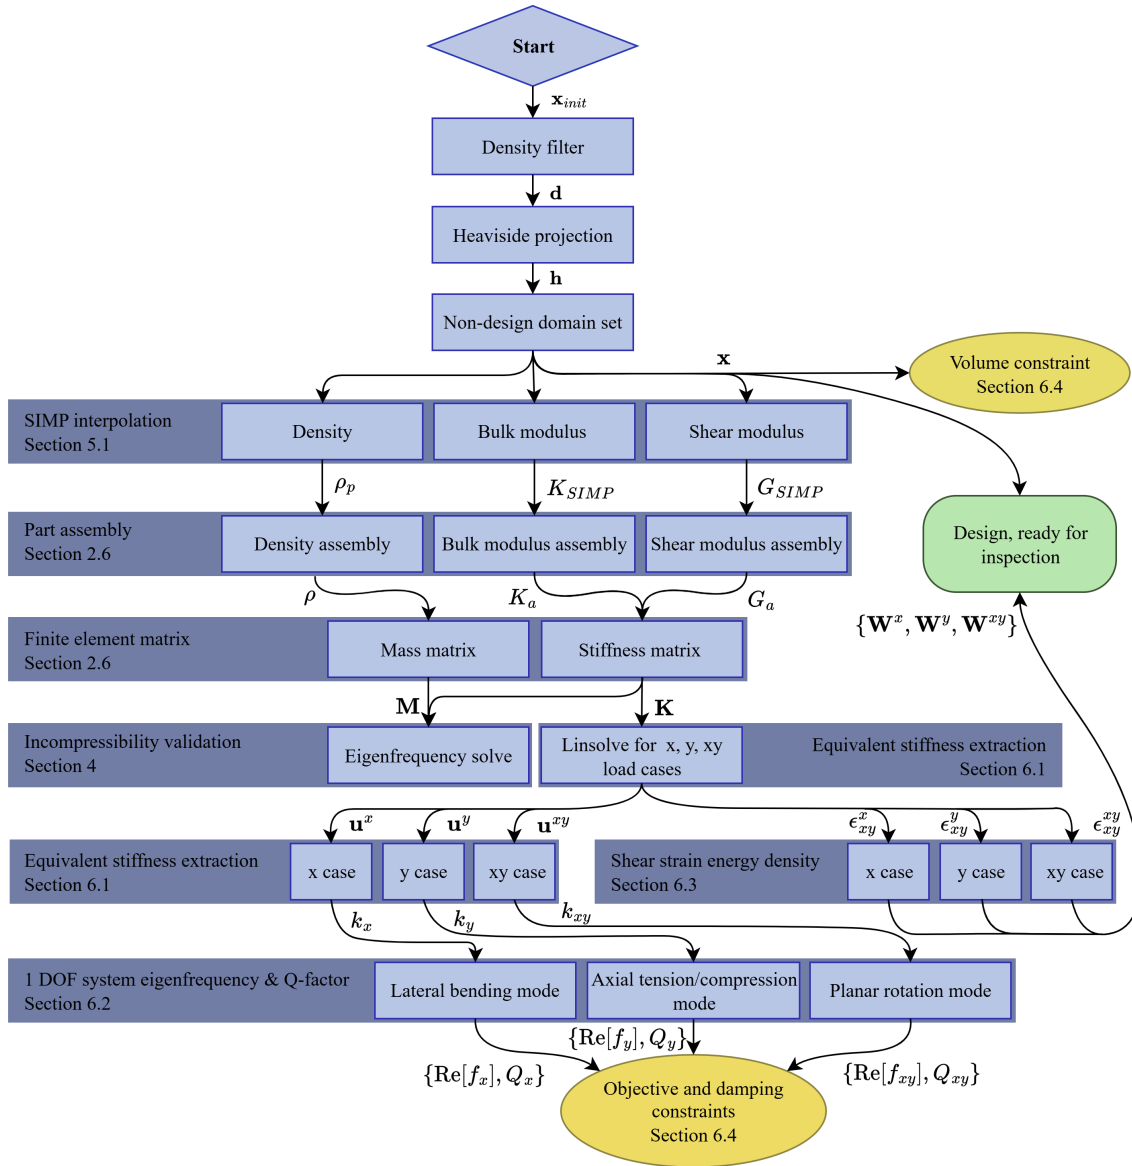

Figure 11: A graphical depiction of a single iteration within the optimisation loop. In each design cycle, the optimiser generates a new design vector $\mathbf{x}_{init}$ based on the evaluated objective and constraint values. The specific steps involved are described in detail in the relevant sections.

A density filter for a structured mesh was applied to eliminate non-physical artefacts such as checkerboard patterns and mesh-dependent designs (Bourdin, 2001). To further promote clear material layouts, the filtered densities were passed through a Heaviside projection scheme, which gradually sharpens intermediate values toward binary solutions (0 or 1) (Wang et al., 2011). The projection operator defined as

$$h_e(x_e) = \frac{\tanh(\beta \cdot t_c) + \tanh(\beta \cdot (x_e - 0.5))}{\tanh(\beta \cdot t_c) + \tanh(\beta \cdot (1 - t_c))}, \tag{20}$$

where x_e is the filtered density at an element and β controls the steepness of the projection, with the cut-off threshold fixed at $t_c = 0.5$. In the implementation used here, β was kept at 1 until iteration 30 and then progressively increased to $\beta = 10$ by iteration 100. This continuation strategy prevents premature convergence while ensuring that the final designs are nearly binary.

After projection, the design variable vector x is transformed into physical material properties using SIMP interpolation. The resulting layout is assembled and analysed using finite element analysis. Objective and constraint values are then evaluated from the structural response.

The optimiser used in this work is the Method of Moving Asymptotes (MMA) introduced by Svanberg (1987). Design updates rely on sensitivities of the objective and constraints with respect to the densities. These sensitivities are calculated automatically in the PyMOTO framework (Delissen, 2023) using backpropagation, allowing the order of modules to be rearranged without manual differentiation.

A schematic overview of a single iteration in the topology optimisation algorithm used in this thesis to generate 2D topology optimised designs with aluminium and incompressible damped rubber is shown in Figure 11. Equivalent 3D implementations were also realised during this work and follow the same procedure, with additional eigenmodes and components contributing to the strain energy density as discussed in Sections 7.5 and 6.3.

5.1 Implications of using bulk and shear modulus representation in SIMP

In traditional SIMP-based topology optimisation, material stiffness is interpolated using a single modulus such as Young's modulus E, while Poisson's ratio ν remains equal for both materials and fixed. This leads to the familiar penalised interpolation scheme (Sigmund, 2001):

$$E_{SIMP}(\rho_e) = E_0(1 - \rho_e^p) + E_1\rho_e^p, \tag{21}$$

where ρ_e is the design variable, p is the penalisation power, and E_0, E_1 are the Young's moduli corresponding to the void (rubber) and solid (aluminium) materials. This interpolation works well when stiffness is governed by a single scalar modulus such as E. In that case, the elasticity matrix $\mathbf{D}$ can be factored as $E \cdot \mathbf{D}'$, as shown in Equations 9, 10, and 11, where D' contains only geometric terms and a fixed Poisson's ratio. This allows the element stiffness matrix to be scaled directly by the interpolated $E_{SIMP}(\rho_e)$.

However, when the constitutive model is expressed in terms of bulk modulus K and shear modulus G, which contribute independently to the volumetric, shear and off-diagonal components of the elasticity matrix $\mathbf{D}$ (see Equations 9, 10, and 12). In other words, K and G appear in separate positions within $\mathbf{D}$ and cannot be factored out, so SIMP interpolation cannot be applied to the full matrix in one step as before. To accommodate this, unit element stiffness matrix $\mathbf{K}_e^u$ is decomposed into two components:

$$\mathbf{K}_e^u = \mathbf{K}_e^{(K)} + \mathbf{K}_e^{(G)}, \tag{22}$$

where $\mathbf{K}_e^{(K)}$ is computed assuming $K = 1, G = 0$ and $\mathbf{K}_e^{(G)}$ is computed with $K = 0, G = 1$. This separation and use of unit modulus values allows the contribution of bulk and shear to be scaled independently during global stiffness matrix assembly. Each contribution is interpolated using a SIMP-style expression:

$$K_{SIMP}(\rho_e) = K_0(1 - \rho_e^p) + K_1\rho_e^p, \tag{23}$$

$$G_{SIMP}(\rho_e) = G_0(1 - \rho_e^p) + G_1\rho_e^p. \tag{24}$$

These interpolated values are then used to scale the precomputed element stiffness matrices during global assembly:

$$\mathbf{K}_e(\rho_e) = K_{SIMP}(\rho_e) \cdot \mathbf{K}_e^{(K)} + G_{SIMP}(\rho_e) \cdot \mathbf{K}_e^{(G)}. \tag{25}$$

Note that this separation applies only to stiffness interpolation. Since density does not differentiate between volumetric and shear behaviour, no decomposition is required. Element material density ρ, which contributes to the mass matrix, is interpolated directly:

$$\rho_p(\rho_e) = \rho_0(1 - \rho_e) + \rho_1\rho_e \tag{26}$$

here, ρ_0 and ρ_1 represent the densities of rubber and aluminium. The use of the stiffness matrix, though, was used only for eigenfrequency analysis discussed in section 4.

6 Modal topology optimisation design approach

When designing a vibration isolator for sensitive equipment, it is important to model how the equipment and the isolator behave together as a single system. Considered separately, the equipment table and the isolator each have their own natural frequencies and mode shapes. Once the equipment is mounted on the isolator the dynamics change. The larger mass of the equipment shifts the resonance frequencies of the isolator to lower values. Analysing the combined system therefore provides a more realistic description of vibration transmission. The challenge is that the system may have many modes, and it is not practical to optimise all of them at once. Modal decomposition offers a solution by breaking the response into distinct modes that can be studied independently.

Through modal decomposition, even a complex system with multiple resonance peaks can be described as a sum of simple single-degree-of-freedom mass–spring–damper systems. Each mode then has only one resonance peak. The modes of the instrument–isolator system fall into two categories. *Rigid body* modes correspond to motion of the equipment table as a whole, with the isolator acting as a flexible support. These are usually modelled under the assumption that the isolator mass is negligible compared to the table. *Internal modes* involve deformation within the isolator and/or the table itself, without rigid body motion of the table. Rigid body modes appear at lower frequencies, while internal modes occur at higher ones, as shown in the combined frequency response in Figure 12. When the internal modes lie far outside the frequency range of interest, only the rigid body modes need to be considered in optimisation.

Figure 12 illustrates the approach. On the left, the rigid body motions of the combined system are shown: vertical and horizontal translations and rotation of the instrument table. In reality these motions are coupled, but after modal decomposition each is treated as an independent mode. The middle plots show how the modes contribute individually to the frequency response, and the right-hand sketches connect them to mechanical analogues with simple spring–damper diagrams.

Restricting optimisation to the rigid body modes makes the problem more manageable. In two dimensions there are three such modes: two translations and one rotation. In three dimensions there can be up to six. In ideal cases, where the instrument table and isolator layout are symmetrical, some of these modes are orthogonal and share identical eigenfrequencies. Each rigid body mode and its damping can then be optimised independently. Optimisation can target the modes most critical to performance, for example the horizontal translation mode if the equipment is especially sensitive to horizontal vibration. It is also possible to optimise several eigenfrequencies together using a mean-eigenvalue method similar to that of Ma et al. (1995), and to set damping constraints separately for each mode. In this way, isolators can be designed to perform well in the relevant frequency ranges without unnecessary complexity.
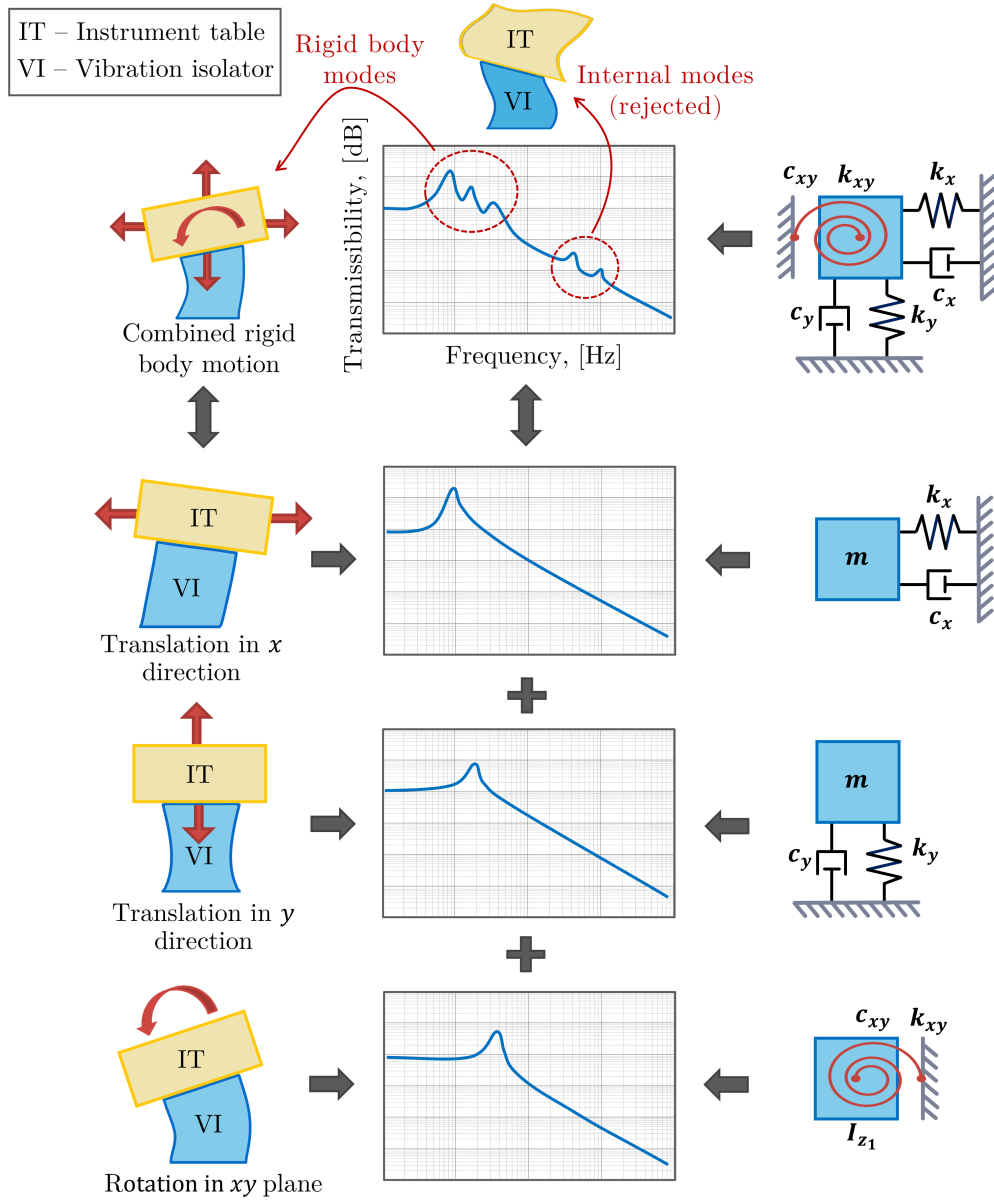
Figure 12: Modal Modelling Approach Outline. The combines isolator+instrument system is reduced into its rigid body modes, each of the mode is represented by a linear or torsional mass-spring-damper system. Internal modes are rejected from the analysis.

6.1 Equivalent stiffness extraction

In order to enable the modal topology optimisation, each isolator design has to be described as an equivalent complex spring stiffness. Once a new freeform design is updated by the optimiser it can by transformed into a complex scalar value to be further used as an objective or constraint. The equivalent stiffness is obtained by using the direct stiffness extraction method, where a force or torque is applied to a specific degree of freedom and the resulting displacement is measured.

To make the direct stiffness extraction more robust, a non-design region, referred to as the *fastening bolt*, is always included at the top of the design domain. This non-design region is set to house

aluminium material at every iteration. The purpose of the bolt is to provide a realistic interface for fastening the vibration isolator to the instrument table or to another system, as the isolator will be clamped in place by fastening a bolt in this region. The bolt non-design area serves several roles: it defines a stiff connection through which force and torque are transferred between components, sets the width for the torque couple in rotational stiffness extraction, and ensures that the force transfer happens over the entire bolt domain rather than a single node as shown in Figure 13. This region is made from aluminium, has a width and depth of $1\,\mathrm{cm}$ and a height of $1.5\,\mathrm{cm}$, and is centred within the top boundary of the $10\,\mathrm{cm} \times 10\,\mathrm{cm} \times 10\,\mathrm{cm}$ design domain.
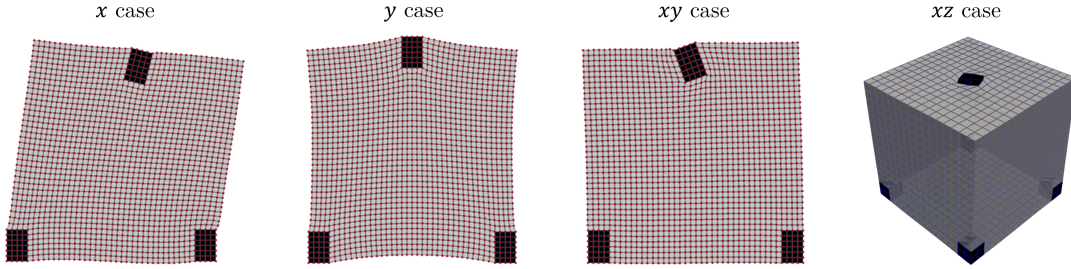

Figure 13: Deformation of finite element meshes during the extraction of equivalent stiffness for different loading cases. Shows deformations for determining stiffness under translational loads in x, y and rotation load applied in the planes xy and xz. The black material in the picture is aluminium and grey is rubber.

To estimate translational stiffness in the x, y, or z directions, a force of $1\,\mathrm{N}$ is applied at the centre of the top boundary. The bottom boundary is fully fixed or constrained by additional bolts. The resulting displacement of the loaded node u is then used to calculate the equivalent stiffness. For example, in the case of x loading, the equivalent stiffness in the x direction, k_x, is given by

$$k_x = \frac{F_x}{u_x}, \tag{27}$$

where F_x is the applied force in the x direction.

Rotational stiffness is determined using a torque couple. For planar rotation, equal and opposite forces are applied at two points on the top boundary, producing a torque of $1\,\mathrm{N\,m}$. The forces are placed symmetrically about the centre, and the *fastening bolt* width w_b defines the lever arm, as the forces act on the upper corners of the bolt. For torsional modes in 3D, two torque couples are applied at four points on the top surface, positioned in the xz plane. This ensures that the rotation of the top bolt due to the applied torque is symmetric.

The equivalent rotational stiffness is calculated from the induced rotation under the small-displacement assumption, where $\tan\theta \approx \theta$. For instance, the rotational stiffness in the xy plane is obtained from

$$k_{xy} = \frac{T_{xy}}{\theta_{xy}} = \frac{F_y w_b^2}{\Delta u_y}, \tag{28}$$

where T_{xy} is the applied torque, θ_{xy} is the resulting angular shift, w_b is the *fastening bolt* width, and Δu_y is the relative displacement in the y direction used to estimate θ_{xy}. In the case of torsional modes, the stiffness is evaluated using the average displacement difference from both torque couples.

6.2 Eigenfrequency of rigid body modes

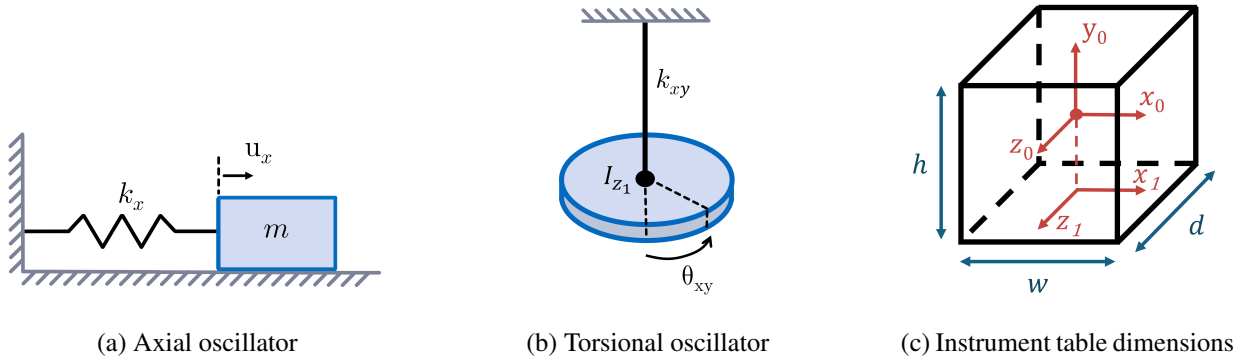(a) Axial oscillator   (b) Torsional oscillator   (c) Instrument table dimensions

Figure 14: Equivalent 1 dimensional mass-spring models and instrument table dimensions used in the analysis.

The dynamic response of the topology-optimised designs is evaluated using an equivalent 1 dimensional mass–spring oscillator model, with all parameters extracted from the direct stiffness method. For each mode, the equivalent stiffness is used together with the corresponding mass or moment of inertia of the instrument table to calculate the complex eigenfrequency.

For axial modes, the system is represented as a single-degree-of-freedom mass–spring oscillator, as shown in Figure 14a. For example, in the x case the complex eigenfrequency is given by

$$f_x = \frac{1}{2\pi} \sqrt{\frac{k_x}{m}}, \tag{29}$$

where k_x is the complex stiffness evaluated using the direct stiffness method for x loading and m is the mass of the instrument.

For rotational modes, the equivalent system consists of a torsional spring and a rigid body with a moment of inertia, as illustrated in Figure 14b. For example, once the equivalent rotational stiffness k_{xy} is determined as described in subsection 6.1, the eigenfrequency can is evaluated using

$$f_{xy} = \frac{1}{2\pi} \sqrt{\frac{k_{xy}}{I_{x_0}}}, \tag{30}$$

where I_{z_1} is the moment of inertia of the instrument table as it rotates around the axis of rotation upon torsional loading in the xy plane. In this work the table is modelled as a solid cube as shown in Figure 14c. Rotations in the xy and zy planes, as opposed to the xz plane, involve axes of rotation located at different positions relative to the centre of gravity. As a result, different expressions are required for the corresponding moments of inertia. The formulas used are given below:

$$I_{x_1} = m\left(\frac{h}{2}\right)^2 + \frac{1}{12}m(d^2 + 4h^2), \tag{31}$$

$$I_{y_0} = \frac{1}{12}m(w^2 + d^2), \tag{32}$$

$$I_{z_1} = m\left(\frac{h}{2}\right)^2 + \frac{1}{12}m(w^2 + 4h^2). \tag{33}$$

The numerical values used for the instrument table moments of inertia are listed in Table 4.

Table 4: Instrument table properties used for optimisation

Property	Mass	Height	Width	Depth	Moment of Inertia (Lateral Bending)		Moment of Inertia (Torsion)
Notation	m	h	w	d	I_{x_1}	I_{z_1}	I_{y_0}
Value	$5\,\mathrm{kg}$	$0.2\,\mathrm{m}$	$0.2\,\mathrm{m}$	$0.2\,\mathrm{m}$	$0.0833\,\mathrm{kgm}^2$	$0.0833\,\mathrm{kgm}^2$	$0.0333\,\mathrm{kgm}^2$

The Q-factor is evaluated from the ratio of the real and imaginary parts of the complex eigenfrequency:

$$Q_i = \frac{\mathrm{Re}[f_i]}{2\,\mathrm{Im}[f_i]}, \tag{34}$$

where the subscript i identifies the type of system the eigenfrequency is evaluated for, $i \in \{x, y, z\}$ is used for axial modes and $i \in \{xy, zy, xz\}$.

6.3 Shear strain energy density

In this thesis, strain energy density is reported only for the *shear* contribution. This choice follows from the fact that damping in the rubber phase is dominated by shear deformation, while bulk deformations contribute negligibly to dissipation. The energy loss is proportional to the strain energy stored, so regions with high strain energy density correspond to higher local dissipation. Focusing on the shear part highlights where deformation directly translates into material damping within the rubber material.

With complex-valued constitutive behaviour (to model damping), the elemental strain energy density is defined using the Hermitian form

$$W = \tfrac{1}{2}\,\varepsilon^\dagger \mathbf{D}_e\,\varepsilon, \tag{35}$$

where ε is the strain vector and $\mathbf{D}_e$ is the complex constitutive matrix for element e. Similarly to subsection 5.1, $\mathbf{D}_e$ is assembled from separately interpolated bulk and shear contributions,

$$\mathbf{D}_e = K_{\mathrm{SIMP}}\,\mathbf{D}_K + G_{\mathrm{SIMP}}\,\mathbf{D}_G, \tag{36}$$

with $\mathbf{D}_K$ and $\mathbf{D}_G$ being elasticity matrices computed using unit-bulk ($K = 1, G = 0$) and unit-shear ($K = 0, G = 1$), and $K_{\text{SIMP}}, G_{\text{SIMP}}$ the complex interpolated moduli.

In the implementation, shear strain energy density was evaluated in the optimised designs presented in section 7. The shear components are extracted in Voigt form and the conjugate is taken explicitly. In 2D this reduces to ε_{xy} only, while in 3D the vector $[\varepsilon_{yz}, \varepsilon_{xz}, \varepsilon_{xy}]^T$ is used together with the shear–shear block of $\mathbf{D}_e$. The real part of W_{shear} visualises stored elastic energy and the imaginary part visualises dissipated energy. Throughout this thesis, these fields are used solely for *visualisation* of where deformation and dissipation concentrate in the optimised layouts.

2D case. Using Voigt ordering $\boldsymbol{\varepsilon} = [\varepsilon_{xx}, \varepsilon_{yy}, \varepsilon_{xy}]^T$, the shear block reduces to the single component ε_{xy} which equates to a scalar shear modulus G_{SIMP} for isotropic materials. The shear strain energy density simplifies to

$$W_{\text{shear}}^{\text{2D}} = \tfrac{1}{2} G_{\text{SIMP}} \, \varepsilon_{xy} \, \varepsilon_{xy}^*, \tag{37}$$

where $(\cdot)^*$ denotes complex conjugation.

3D case. With Voigt ordering $\boldsymbol{\varepsilon} = [\varepsilon_{xx}, \varepsilon_{yy}, \varepsilon_{zz}, \varepsilon_{yz}, \varepsilon_{xz}, \varepsilon_{xy}]^T$, the shear vector is defined as $\boldsymbol{\varepsilon}_s = [\varepsilon_{yz}, \varepsilon_{xz}, \varepsilon_{xy}]^T$. For isotropic materials, the shear–shear submatrix of $\mathbf{D}_e$ reduces to

$$\mathbf{D}_s = G_{\text{SIMP}} \, \mathbf{I}, \tag{38}$$

where $\mathbf{I}$ is the 3×3 identity matrix. The shear strain energy density then reads as

$$W_{\text{shear}}^{\text{3D}} = \tfrac{1}{2} \, \boldsymbol{\varepsilon}_s^\dagger \mathbf{D}_s \, \boldsymbol{\varepsilon}_s. \tag{39}$$

6.4 Topology optimisation problem statement

The topology optimisation problem is formulated to maximise the vibration isolation performance of the design by targeting specific rigid body modes. The objective is to maximise the fundamental eigenfrequencies of one or more targeted modes, subject to constraints on modal damping and material usage. The optimisation problem described in Equation 40 reflects the approach used throughout this work:

$$\left.\begin{aligned} \min_{\mathbf{x}} \; &: \; \sum_{i \in S} \frac{1}{\text{Re}[f_i(\mathbf{x})]} \\ \text{subject to} \; &: \; Q_i(\mathbf{x}) \leq Q_{\max} \quad \forall i \in S \\ & \quad \frac{V(\mathbf{x})}{V_0} \leq v_f \\ & \quad 0 < x_{\min} \leq x_j \leq 1 \quad \forall j \end{aligned}\right\} \tag{40}$$

The design variables $\mathbf{x}$ represent the element-wise material densities, bounded below by $x_{\min}$ to prevent singularities in the global stiffness matrix $\mathbf{K}$. In this thesis, $x_{\min} = 10^{-9}$ was used.

The functions $f_i(\mathbf{x})$ denote the complex eigenfrequencies of the targeted modes within the set S. The set S contains all of the rigid body modes to be optimised, possible choices for axial modes include $\{x, y, z\}$ and for rotational modes $\{xy, zy, xz\}$.

The objective minimises the sum of the inverses of the real parts of eigenfrequencies to be optimised for, following the harmonic mean approach of Ma et al. (1995), which is effective for ensuring symmetry or balanced performance in multiple directions. This formulation emphasises maximising the lowest eigenfrequency while also promoting increases in higher-order modes.

For each mode i in the selected set S, the modal Q-factor must remain below the prescribed limit Q_{max} to ensure sufficient damping. Although Q-factor constraints are mainly applied to modes in the objective, they can also be imposed on other modes if required. A volume constraint limits the total material usage to a fraction v_f of the design domain. A value of $v_f = 1.0$ may be used for analysis when key material placement locations for damped rubber need to be identified.

This formulation is flexible and can be applied to either a single eigenfrequency or to multiple modes, depending on the requirements of the design problem. In this work, a single mode is targeted unless symmetry or multi-directional performance is required, in which case the set of included modes and corresponding constraints are adjusted. In this way, the optimiser can be tailored to the specific isolation requirements of the instrument or application.

7 Optimisation results

7.1 Design strategy and case definitions

The design space is large and non-convex, and adding many assumptions or objectives at once can trap the optimiser in local minima. This thesis adopts a staged strategy with as few early assumptions as possible. First, a stiffness-only run with a low volume fraction v_f (for example $v_f = 0.4$) and no imposed bolt locations reveals where the structure prefers to carry load and highlights candidate support regions. Based on that evidence and the intended loading, bottom bolt locations are then chosen (two corners in 2D, four corners in 3D by default). This method is not limited to fastening options at the bottom of the design domain. If integration into the larger system requires load paths to be transferred via side walls, the same procedure can be used to place side-mounted bolts. Once these non-design bolt regions are fixed, the design space is reduced and stricter goals are introduced: damping constraints, single or multimode objectives, and symmetry where required.

The results focus on optimising the rigid body motions of the combined isolator-instrument systems introduced in section 6, during which the isolator deforms with eigenmodes classified in subsection 4.1. Three cases are used throughout the 2D analysis. The *x case* targets the *lateral bending* mode in which the fastening bolt moves along the global x direction. The *y case* targets the *axial tension/compression* mode with dominant translation of the payload tip along y. The
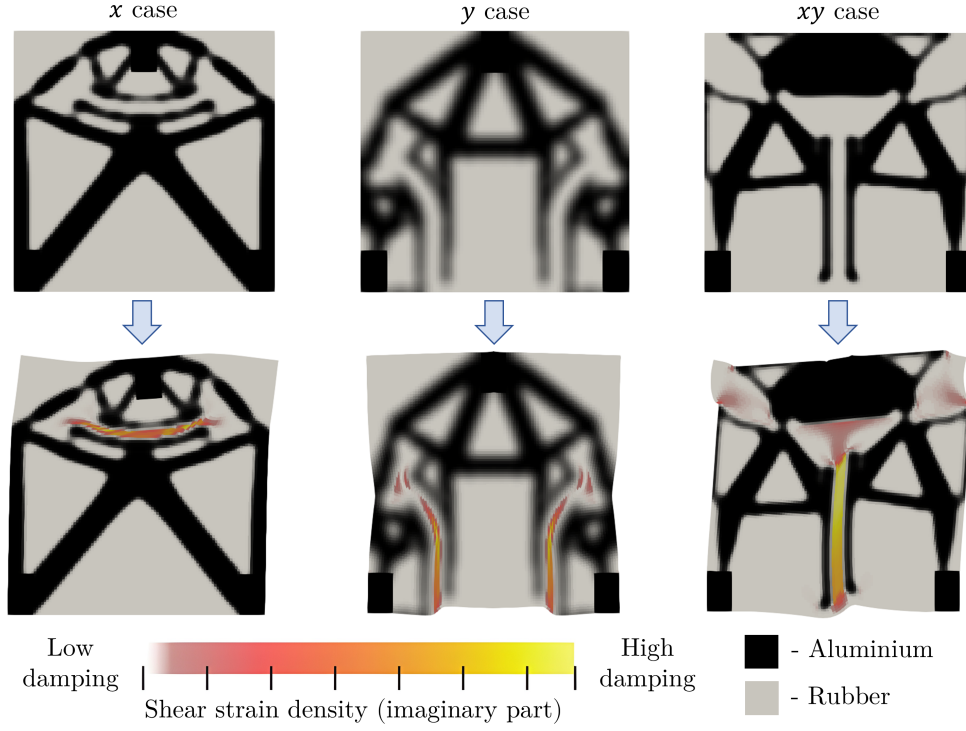
Figure 15: Optimised designs for x, y, and xy cases. Designs are optimised for eigenfrequency maximisation with a Q-factor constraint. The top row depicts undeformed designs and the bottom row shows the same designs in their displaced states under the targeted mode (note: grey colour is used for rubber to make their deformation more apparent). Colour map indicates high shear strain energy density regions, highlighting where heat dissipation occurs when the eigenmode is excited.

xy case targets the *planar rotation* mode, characterised by rotation in the xy plane about the z axis. For each case, the mode used in the objective is the one with the largest participation of the specified motion (x, y, xy) at the payload tip (or rotation of the top boundary) rather than by frequency ordering alone.

Once the damping constraint is combined with an eigenfrequency objective, more intricate designs emerge. While the overall layout still resembles stiffness-only solutions, new features appear that exploit mechanism-like behaviour under the targeted mode. Figure 15 shows representative examples for three core optimisation cases used in 2D (x, y, xy). Across all three cases, CLD-like damping features can be seen, with parallel aluminium plates and rubber layers between them. In many instances, the designs also incorporate hinge-like joints that act as compliant mechanisms, amplifying local motion in the damping regions. The optimisation parameters for these results are summarised in Table 5. The y-optimisation case is shown without Heaviside projection, for reasons discussed in subsection 7.3.

7.2 Effects of varying Q-factor and volume fraction

Because the topology optimisation framework works with structural and damped materials, volume fraction v_f and the damping limit $Q_{\max}$ offer complementary lenses on the designs. A low

Table 5: Optimisation parameters for the results in Figure 15.

Parameter	Volume Fraction	Q-factor	Resolution	Filter Radius	Heaviside	Projection Threshold
Notation	v_f	Q_{max}	–	$r_{\min}$	β	$\tilde{x}$
Value	0.4	5	100×100	4.5	$1 \to 10$ (iter. 30→100)	0.5

v_f exposes the critical aluminium load paths. Relaxing v_f and tightening $Q_{\max}$, in turn, shift the focus on damping and highlights where rubber is most effective for dissipation. To investigate the interplay between stiffness and damping in the optimised designs, a set of solutions was generated for the x optimisation case while varying both the volume fraction v_f and the maximum allowable $Q_{\max}$. The results are arranged as a grid in Figure 16, with v_f increasing from top to bottom and damping limit $Q_{\max}$ increasing from left to right, making the transition from damping-dominated to stiffness-dominated designs easy to follow.
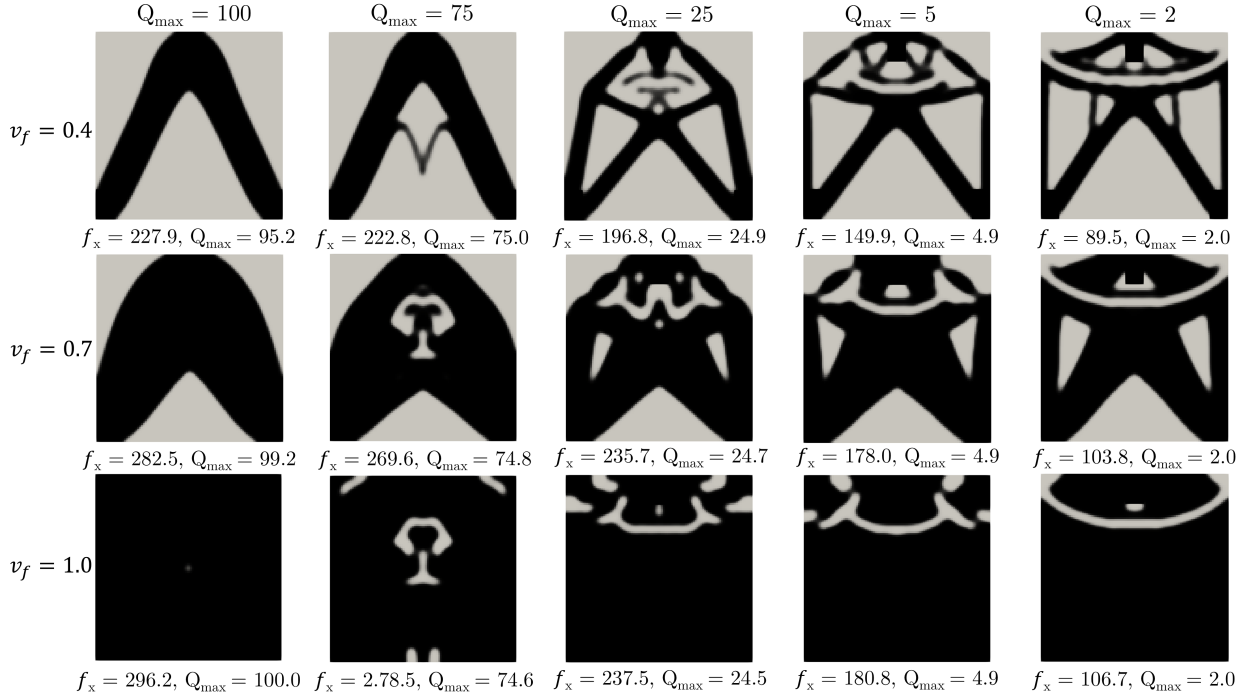
Figure 16: Grid of optimised designs to maximise the eigenfrequency of translational mode in the x direction. The grid highlights the effect of volume fraction v_f and damping limit $Q_{\max}$. The obtained eigenfrequency and Q-factor for each optimization configuration are shown below the corresponding design.

At low volume fraction ($v_f = 0.4$), most of the available aluminium is used to form a stiff load-carrying frame, leaving only a few key shear interfaces for damping. At moderate volume fraction $v_f = 0.7$, more of the domain is filled with stiff aluminium material, leaving smaller regions of rubber. At this intermediate volume fraction, some of the rubber regions act as damping sites while others serve mainly as weight-saving gaps. When volume fraction is unconstrained ($v_f = 1.0$), aluminium dominates the domain and rubber is placed only in key areas for damping, forming fewer

but enlarged or elongated regions of damping to satisfy the damping constraint. Varying $Q_{\max}$ shifts the balance between damping and stiffness. At high $Q_{\max}$ (e.g. $Q_{\max} = 100$), the damping constraint becomes negligible and the layouts revert to near-trivial compliance-minimising forms. As $Q_{\max}$ is reduced, minimal damping structures emerge at locations that are secondary for stiffness but optimal for damping. Lowering $Q_{\max}$ further encourages thinner hinges and more compliant motion that amplifies shear in the rubber. If pushed too far (e.g. $Q_{\max} = 2$), stiffness becomes secondary and extreme cases appear where the top bolt is essentially carried by rubber alone. This family of designs is referred to as a *detached head*, meaning there is no aluminium connectivity between the top fastening bolt and the base. Interestingly, this is when designs collapse toward the familiar CLD damping treatment from the literature. However, the curved shape and location of the CLD structure are still better adapted for the deformation caused by loading. A clear trade-off runs across the grid. Reducing $Q_{\max}$ improves damping but lowers achievable eigenfrequencies, whereas increasing $Q_{\max}$ boosts stiffness at the expense of damping. Finally, while higher v_f increases mass, since both material phases are dense materials ($\rho_0 = 1500\,\mathrm{kg/m^3}$, $\rho_1 = 2700\,\mathrm{kg/m^3}$), substantial weight reduction would require introducing a third, void-like material phase to the formulation.

7.3 Edge cases

7.3.1 Asymmetric designs

Some y-optimised runs converge to asymmetric layouts. This is most often observed when Heaviside projection is used, when the Poisson ratio is high ($\nu \geq 0.49$), or when moderate to high volume fractions are applied ($0.4 < v_f \leq 1.0$). This is noteworthy because most designs generated for x and y *cases* are symmetric and produce displacement only along the required axis of translation or rotation. However, in y *case* optimisation runs once the damping requirement is satisfied on one side, the remaining material is often allocated to a stiff connection elsewhere, which alters the tip trajectory and introduces an x component of motion. The behaviour does not align with the 1-DOF oscillator model in subsection 6.2. For asymmetric designs, the displacement under a y-directed load now has two components (x and y), which increases the eigenvalue. However, this is essentially an exploit of the absence of enforced symmetry constraints. The same effect is observed in both 2D and 3D y *case* optimisation results as shown in Figure 17.
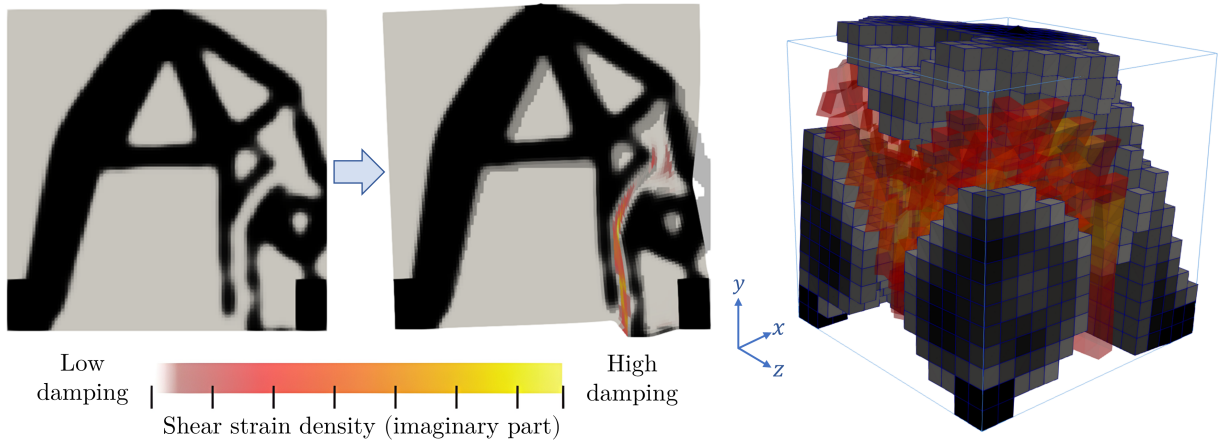
Figure 17: Asymmetric y-optimised designs in 2D and 3D. The 2D result is shown at a unperturbed (left) and perturbed (middle) state. The 3D results (right) is shown at a perturbed state with a blue wireframe showing the boundary of the unperturbed design.

7.3.2 Complications due to projection and incompressibility

This subsection examines the behaviour of the x-optimised case under near-incompressibility, where $\nu \to 0.5$ amplifies the bulk response of the rubber and can interact strongly with projection. The aim is to assess how projection influences the resulting layouts and whether path dependence arises when ν is varied. To this end, a sweep is performed with $Q_{\max} = 5$ and no volume constraint, comparing projected and non-projected cases across $\nu \in \{0.48, 0.49, 0.499, 0.49999, 0.499999\}$. The resulting designs are summarised in Figure 18, which also includes reinitialised runs using the $\nu = 0.48$ projected design as a starting point.

$\nu = 0.48$ $\nu = 0.49$ $\nu = 0.499$ $\nu = 0.4999$ $\nu = 0.499999$

Heaviside OFF

$f_\mathrm{x} = 147.4$, $Q_\mathrm{max} = 5.0$ $f_\mathrm{x} = 149.8$, $Q_\mathrm{max} = 5.0$ $f_\mathrm{x} = 185.0$, $Q_\mathrm{max} = 5.0$ $f_\mathrm{x} = 189.5$, $Q_\mathrm{max} = 5.0$ $f_\mathrm{x} = 204.3$, $Q_\mathrm{max} = 5.0$

Heaviside ON
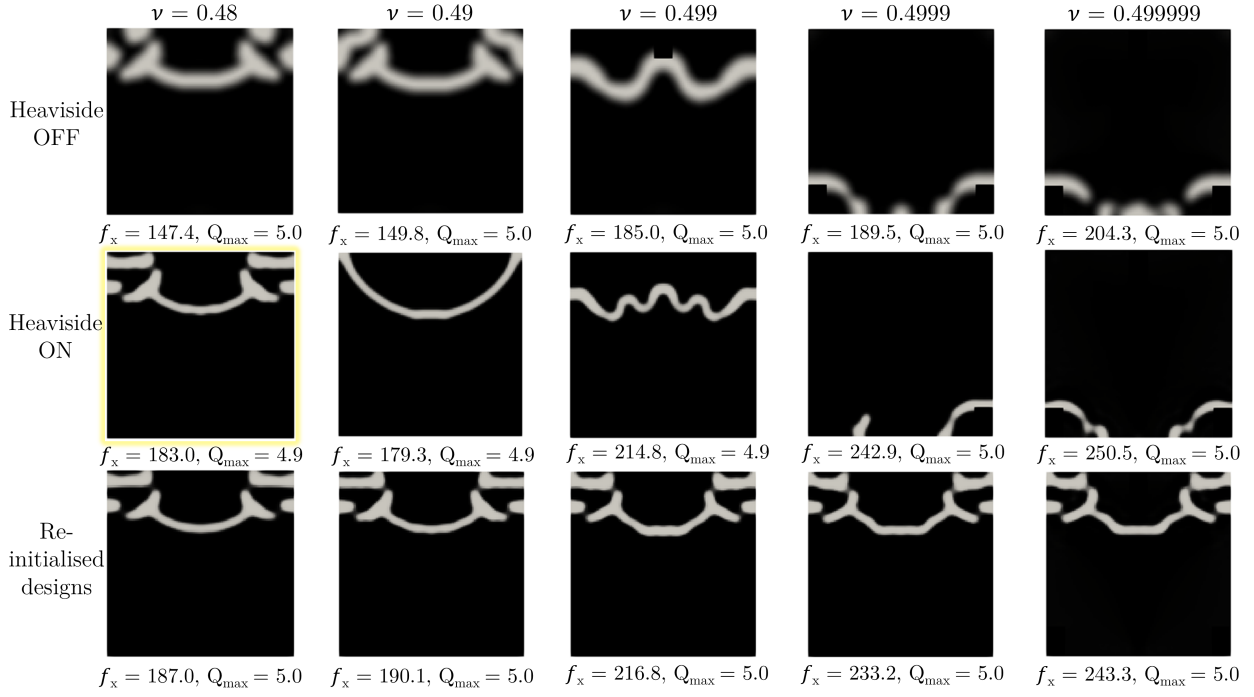
$f_\mathrm{x} = 183.0$, $Q_\mathrm{max} = 4.9$ $f_\mathrm{x} = 179.3$, $Q_\mathrm{max} = 4.9$ $f_\mathrm{x} = 214.8$, $Q_\mathrm{max} = 4.9$ $f_\mathrm{x} = 242.9$, $Q_\mathrm{max} = 5.0$ $f_\mathrm{x} = 250.5$, $Q_\mathrm{max} = 5.0$

Re-initialised designs

$f_\mathrm{x} = 187.0$, $Q_\mathrm{max} = 5.0$ $f_\mathrm{x} = 190.1$, $Q_\mathrm{max} = 5.0$ $f_\mathrm{x} = 216.8$, $Q_\mathrm{max} = 5.0$ $f_\mathrm{x} = 233.2$, $Q_\mathrm{max} = 5.0$ $f_\mathrm{x} = 243.3$, $Q_\mathrm{max} = 5.0$

Figure 18: Grid of x-optimised designs for increasing Poisson ratio, $Q_\mathrm{max} = 5$, no volume constraint. Top row: direct optimisation without projection. Middle row: direct optimisation with projection. Bottom row: projected runs reinitialised from the $\nu = 0.48$ design (highlighted in yellow).

Projection. Projection results in same designs with sharper boundaries and higher objective values. Designs also appear more trustworthy because the non-physical mixed aluminium-rubber phase is reduced. However, projection can push otherwise stable x-case results into asymmetry, as seen by comparing the $\nu = 0.4999$ results with and without projection in Figure 18. Projection also adds difficulty for the optimiser: the β continuation introduces small discontinuities that perturb sensitivities and may drive solutions to local optima. An example of this effect is seen in the $\nu = 0.49$ projected run, as it converges to a *detached head* layout, which has no aluminium connectivity between the top fastening bolt and the bottom supports. Even though the $\nu = 0.49$ rubber is stiffer, the *detached head* solution yields a lower objective than the $\nu = 0.48$ variant. The $\nu = 0.499$ projected run attains a higher objective, but this coincides with the rapid increase of the bulk modulus of the rubber as $\nu \to 0.5$, so comparisons across ν must be interpreted with care.

Poisson ratio. Across the sweep, designs at $\nu = 0.48$ are well-behaved, whereas for $\nu \geq 0.49$ many results converge to detached head layouts. At $\nu \in \{0.4999, 0.499999\}$ an *amplified shear* substructure appears, in which motion at the top of the isolator is amplified below and imposes strong shear on rubber above the bottom bolts.

Path dependence. The third row of Figure 18 reuses the $\nu = 0.48$ projected design as the initial condition. At $\nu \in \{0.49, 0.499\}$ these reinitialised runs achieve higher objectives than the directly

optimised detached head cases, indicating that the latter are likely local optima. For $\nu \geq 0.4999$, however, the reinitialised runs yield lower objectives than direct optimisation, suggesting that the detached head layouts represent the global optimum within this model.

Amplified shear substructure. Direct optimisation results for $\nu \geq 0.4999$ reveal a behaviour that seems to exploit the computational model. Figure 19 highlights the damping mechanism that emerges: a very thin rubber channel (two elements across) subjected to amplified motion as the fastening bolt deflects. The local kinematics resemble incompressible fluid-like expulsion or suction of rubber as the channel volume varies. This *amplified shear* mechanism combines high stiffness (higher eigenvalues than *attached head* layouts) with high dissipation (large shear strain energy density). Its extreme thinness suggests the optimiser may be exploiting numerical or modelling limitations, potentially linked to element-scale resolution, shear locking, or the strong bulk response at $\nu \approx 0.5$. In practice, such a configuration would be vulnerable to damage or loss of adhesion, so the substructure is considered an unvalidated artefact rather than a viable design feature.
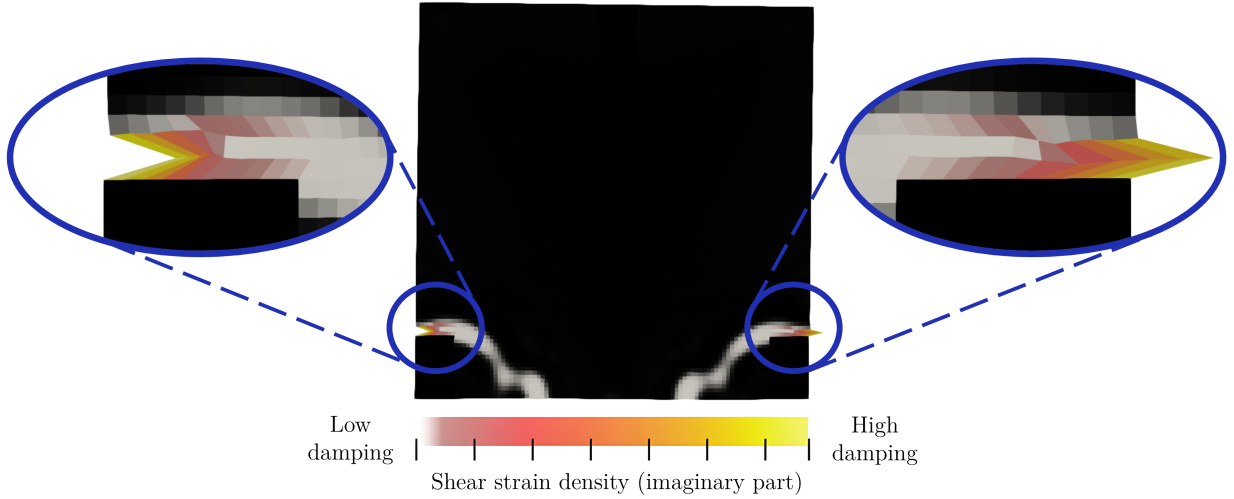

Figure 19: Zoom-in of the *amplified shear* substructure at $\nu = 0.499999$ (projection on, $Q_{\max} = 5$, no volume constraint).

7.4 Damping substructures

Two main groups of damping substructures are identified: *parallel plate* and *pinch*. The parallel plate type is the most common and generally the most effective. It is characterised by relative sliding motion during deformation, with the plate shape adapting to the dominant vibration mode. Plates may be straight, curved along an arc, or elongated into channels. Their length typically scales with the damping requirement and available space in the region of largest motion. Detached head variants also occur, in which parallel plates carry the top fastening bolt without direct aluminium connectivity to the supports.
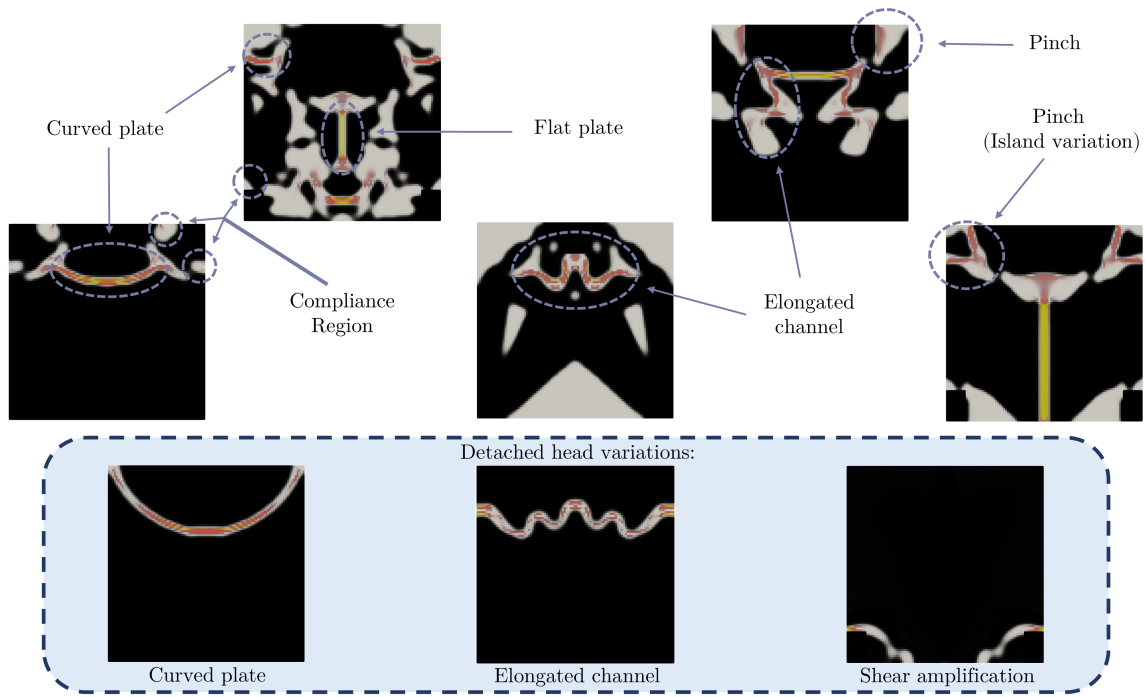
Figure 20: Catalogue of identified damping substructures with representative examples.

The second group comprises *pinch* substructures, formed by two oblique aluminium boundaries that compress or stretch a rubber layer between them as they come closer or move apart. This produces a pinching action and shear in the rubber, promoting dissipation. Pinch substructures can be reinforced by an intermediate island that further enhances shear. In some cases, regions of rubber function less as damping sites and more as compliant zones that allow larger motion elsewhere. A *detached head* variation of the pinch type was also noted, termed *shear amplification*, and was introduced earlier in subsection 7.3.

Optimised designs often contain multiple damping substructures that engage simultaneously. Occasionally, both parallel plate and pinch mechanisms are combined, distributing shear and dissipation across the domain. This behaviour, however, is not universal and depends on optimisation settings. The main damping substructures and their variants are summarised in Table 6 and illustrated in Figure 20.

Table 6: Classification of identified damping substructures

Variation	Parallel plate	Pinch
Connected	Flat plate; Curved plate; Elongated channel	Pinch; Pinch (island variation); Compliance region
Detached head	Curved plate; Elongated channel	Shear amplification

41

7.5 Combined Objectives

So far, results have been limited to optimising a single eigenfrequency. However, as outlined in subsection 6.4, the framework can accommodate multiple eigenfrequencies and damping constraints within a single run. The selection of which eigenfrequencies to maximise or constrain can be tailored to the application or used to enforce symmetry, particularly in 3D designs.

In 3D, six rigid-body-dominated motions are relevant, but two pairs of them are equivalent. The *x case* and *z case* correspond to lateral bending of the isolator in orthogonal directions, while the *xy case* and *zy case* correspond to planar rotation about orthogonal axes. To ensure better multi-directional performance, it is natural to optimise the eigenfrequencies of the *x* and *z* cases together, or of the *xy* and *zy* cases, while also applying damping constraints to the targeted modes.

In 2D, one may choose to maximise two eigenfrequencies simultaneously, for example the *x case* and *xy case*, while constraining damping in the remaining mode. Many such combinations are possible depending on the intended application and the requirements of the design. Two representative examples of multi-objective optimisation runs are shown in Figure 21. These results are included as a brief demonstration of the broader capabilities of the framework, which are left open for further exploration once more concrete case studies are established.
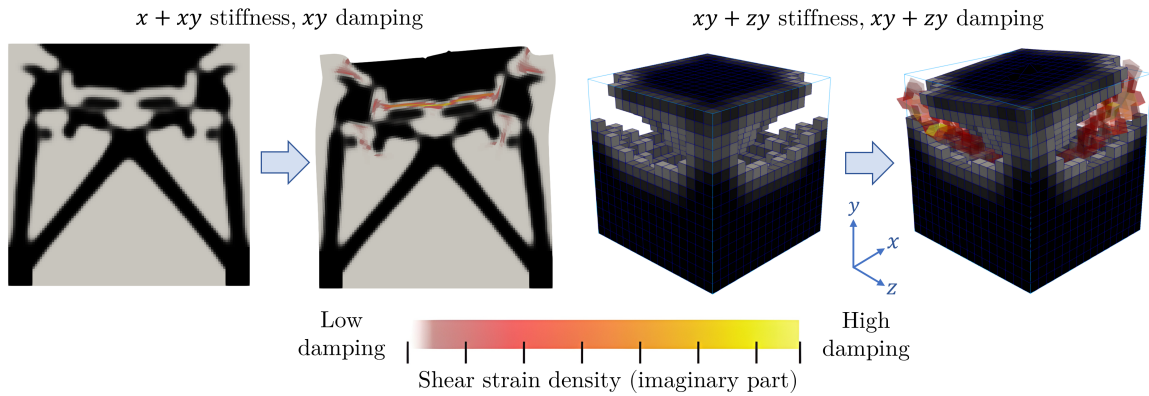

Figure 21: Examples of optimisation runs involving multiple eigenfrequencies and damping constraints.

8 Conclusion and discussion

8.1 Main contributions

1. **Framework development and implementation:**

   (a) Developed a topology optimisation framework that freely distributes damped and structural material within a structural domain, moving beyond conventional CLD *sandwich* configurations that dominate literature.

   (b) Enabled adaptability to boundary and loading conditions, including varied mounting point locations.

   (c) Integrated incompressible viscoelastic material damping properties into the finite element formulation using a real bulk modulus and a complex shear modulus.

   (d) Investigated remedies for numerical locking in FEM-based topology optimisation and identified reduced integration as most effective.

   (e) Applied incompressibility handling by implementing selective reduced integration into the PyMOTO FEM modules.

   (f) Developed a consistent method for visualising shear strain energy density in optimised designs to identify high damping regions.

   (g) Established a versatile design procedure, providing a basis for future research that combines objectives and constraints for specific applications.

2. **Findings and results:**

   (a) Framework generated unique, physically meaningful designs that have not been reported in the literature before. The design region is not confined within the CLD layer, however substructures similar to CLD now emerge within the isolator layout at strategically placed locations.

   (b) The choice of rigid body modes to optimise for within the topology optimisation routine can tune the isolator's performance according to the inertial properties of the sensitive instrument unit.

   (c) Discovered and categorised damping substructures, including *parallel plate*, *pinching*, and mixed mechanism configurations. Showed that these damping mechanisms are case-specific and often non-intuitive, producing results unlikely to be conceived by a human designer.

8.2 Limitations:

The framework and the assumptions on which it relies introduce a number of limitations and open questions. These relate to the accuracy of the finite element modelling, the absence of manufacturing and practical constraints, the simplified treatment of viscoelastic material behaviour, and

the idealised nature of modal topology optimisation approach and boundary conditions. Each of these areas is discussed in turn below.

Accuracy   The accuracy of the finite element modelling remains an important open question. While the framework reproduces reference results for simple cubic or prismatic rubber blocks, as discussed in section 4, it is less certain that this accuracy holds for the irregular and intricate geometries produced by topology optimisation. Incompressibility is currently modelled using a reduced integration approach, which mitigates but does not entirely eliminate volumetric locking, and localised locking effects may still occur in complex subregions. Shear locking has not been systematically investigated, either in simple test blocks or in the optimised designs, leaving another potential source of inaccuracy unquantified. The damping model itself, which treats the bulk modulus as real and the shear modulus as complex, has a sound theoretical basis and has been used in previous studies, but its validity in a topology-optimised setting remains to be confirmed. Furthermore, the assumption of perfect adhesion between aluminium and rubber in the model neglects practical issues such as limited interfacial bonding and differential thermal expansion, which in real structures could lead to partial or complete de-bonding. Experimental testing is therefore essential to confirm that the simulated dynamic behaviour is representative of what would occur in practice.

Manufacturability   The current optimisation framework does not include any manufacturability constraints, meaning that some features in the resulting designs may be impractical or impossible to produce. No stress constraints are imposed, so highly localised stress concentrations could arise, leading to premature failure in real applications. Static or launch loading of the isolator has also not been considered, even though sustained loads could influence both the structural integrity of the aluminium regions and the long-term behaviour of the rubber. From a manufacturing perspective, the most feasible approach for these designs would likely involve casting or additive manufacturing of the aluminium framework, followed by moulding rubber into the resulting cavities. However, this process would be incompatible with certain configurations produced by the optimiser, such as isolated aluminium islands or fully enclosed voids, which cannot be accessed for rubber casting. Incorporating geometric constraints, penalisation schemes, or load-based restrictions to prevent such unmanufacturable or mechanically unstable features would therefore be a necessary step towards transitioning the method from a purely academic tool to a design process suitable for industrial use.

Viscoelasticity   The framework also simplifies the behaviour of viscoelastic rubber, omitting several effects that could significantly influence real-world performance. Frequency-dependent properties are not included, even though the stiffness and damping characteristics of viscoelastic materials can vary with excitation frequency. Capturing this behaviour would require solving a non-linear eigenfrequency problem in which both the elastic moduli and the loss factor vary with frequency. Temperature dependence has likewise been neglected, despite its relevance not only

in satellite applications, where thermal gradients and fluctuations are common, but also due to self-heating in the rubber during operation. As the damping substructures dissipate energy, their temperature could rise, altering their mechanical properties in ways not captured by the current model. Long-term effects such as creep, which can cause gradual shifts in geometry and material properties under sustained loads, are also not represented. In some applications, pre-tensioning is used to mitigate creep-induced changes, but this has not been considered here.

Modal decomposition   The modal topology optimisation design approach described in section 6 rests on several simplifying assumptions. The connection between the vibration isolator and the instrument table is modelled as acting solely through the mounting bolts, whereas in reality the entire top surface of the isolator would be in contact with the instrument. This broader contact area could introduce additional coupling effects and alter the dynamic response. The internal modes of both the isolator and the instrument table are also disregarded, under the assumption that they are far from the target modes of interest. However, in some cases these internal modes could lie close enough in frequency to influence overall performance. A further simplification is the reduction of the full geometry to a single complex stiffness value, following the one-degree-of-freedom complex stiffness–mass oscillator model outlined in subsection 6.2. In this reduction, the isolator mass is neglected, since it is replaced by an equivalent massless spring. While this mapping is convenient, its accuracy for highly non-uniform, topology-optimised layouts remains uncertain. In symmetric and homogeneously damped structures the approximation is often adequate, but in designs where stiffness and damping are distributed unevenly, the reduction may not fully capture the true dynamic behaviour. Additionally, some optimised designs, such as the asymmetric ones, exhibit multi-directional motion when excited, deviating from the pure single-axis behaviour assumed in the reduced model.

Lack of data   Finally, the study was conducted in a largely academic context, with limited use of real-world application data. The rubber properties used in the simulations were derived from tensile tests on the aerospace-grade elastomer SCVBR® supplied by ERIKS. These tests were performed at TNO across several temperatures, with the $20\,°C$ data point selected for this study. The mass and inertia values of the instrument table were likewise approximated, rather than obtained from a specific hardware configuration.Representative micro-vibration spectra, which would ideally serve as design targets, were not available for this project, as publicly accessible datasets remain scarce and specifications are rarely shared in detail. Similarly, no constraints on size, preferred mounting arrangement, or mass budget were imposed either. Consequently, the designs explored here were shaped entirely by the optimisation objectives and constraints defined within the study. This ensured a clear focus on the methodological aspects, while leaving scope for future work to incorporate application-specific requirements and more detailed input data.

References

ABAQUS, Inc. (2004). *Getting Started with ABAQUS v6.5*. User Manual and Technical Documentation, Chapter 4.1: Element formulation and integration.

Aglietti, G S et al. (2004). "Model building and verification for active control of microvibrations with probabilistic assessment of the effects of uncertainties". en. In: *Proceedings of the Institution of Mechanical Engineers, Part C: Journal of Mechanical Engineering Science* 218.4, pp. 389–399. DOI: `10.1177/095440620421800404`.

ANSYS, Inc. (2017). *ANSYS Mechanical APDL Theory Reference*. Tech. rep., Section 3.6: General Element Formulations.

Bendsøe, Martin P. and O. Sigmund (2003). *Topology optimization: theory, methods, and applications*. en. Berlin ; New York: Springer. DOI: `10.1007/978-3-662-05086-6`.

Bourdin, Blaise (2001). "Filters in topology optimization". en. In: *International Journal for Numerical Methods in Engineering* 50.9, pp. 2143–2158. DOI: `10.1002/nme.116`.

Bower, Allan F. (2010). *Applied mechanics of solids*. en. Boca Raton, Fla.: CRC Press.

Cook, Robert D. and Robert Davis Cook, eds. (2002). *Concepts and applications of finite element analysis*. eng. 4. ed. New York, NY: Wiley.

Danh, Le Thanh and Kyoung Kwan Ahn (2014). "Active pneumatic vibration isolation system using negative stiffness structures for a vehicle seat". In: *Journal of Sound and Vibration* 333.5, pp. 1245–1268. DOI: `10.1016/j.jsv.2013.10.027`.

Dealy, J. and Donald Plazek (2009). "Time-temperature superposition-a users guide". In: *Rheol. Bull.* 78, pp. 16–31.

Delissen, Arnoud (2023). *pyMOTO: Modular framework for topology optimization with semi-automatic derivatives*. DOI: `10.5281/zenodo.7708738`.

Díaaz, Alejandro R. and Noboru Kikuchi (1992). "Solutions to shape and topology eigenvalue optimization problems using a homogenization method". en. In: *International Journal for Numerical Methods in Engineering* 35.7, pp. 1487–1502. DOI: `10.1002/nme.1620350707`.

Du, Jianbin and Niels Olhoff (2005). "Topological optimization of continuum structures with respect to . . . eigenfrequencies". In.

— (2007). "Topological design of freely vibrating continuum structures for maximum values of simple and multiple eigenfrequencies and frequency gaps". en. In: *Structural and Multidisciplinary Optimization* 34.2, pp. 91–110. DOI: `10.1007/s00158-007-0101-y`.

Elmadih, Waiel (2019). "Additively manufactured lattice structures for vibration attenuation". en. In.

Elmoghazy, Yasser Hamed et al. (2024). "Computational Modelling and Analysis of Effect of Viscoelastic Materials on Damping and Vibrational Behaviors of Composite Structures—An Extensive Review".

en. In: *Archives of Computational Methods in Engineering* 31.5, pp. 2611–2662. DOI: `10.1007/s11831-023-10057-4`.

Fang, Zhanpeng and Ling Zheng (2015). "Topology Optimization for Minimizing the Resonant Response of Plates with Constrained Layer Damping Treatment". en. In: *Shock and Vibration* 2015.1, p. 376854. DOI: `10.1155/2015/376854`.

Fauser, Dominik et al. (2025). "Complex Poisson's ratio for viscoelastic materials: direct and indirect measurement methods and their correlation". In: *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences* 481.2309, p. 20240543. DOI: `10.1098/rspa.2024.0543`.

Inozume, Seita and Tatsuhito Aihara (2021). "Maximising Mode Damping Ratios in Vehicle Frames Under Load Conditions via Multi-Material Topology Optimisation". In: *Transactions of the Japan Society of Mechanical Engineers* 87.897, pp. 21–00001. DOI: `10.1299/transjsme.21-00001`.

Jafari, Behzad (2018). "Whole Spacecraft Vibration Isolation System: A Comparison of Passive vs. Semi-Active Vibration Isolation Designs". en. In.

Jensen, Jakob S. and Niels L. Pedersen (2006). "On maximal eigenfrequency separation in two-material structures: the 1D and 2D scalar cases". In: *Journal of Sound and Vibration* 289.4, pp. 967–986. DOI: `10.1016/j.jsv.2005.03.028`.

Johnson, Conor D. and David A. Kienholz (1982). "Finite Element Prediction of Damping in Structures with Constrained Viscoelastic Layers". en. In: *AIAA Journal* 20.9, pp. 1284–1290. DOI: `10.2514/3.51190`.

Kang, Zhan et al. (2012). "On topology optimization of damping layer in shell structures under harmonic excitations". en. In: *Structural and Multidisciplinary Optimization* 46.1, pp. 51–67. DOI: `10.1007/s00158-011-0746-4`.

Kawak, B. J. (2017). "Development of a low-cost, low micro-vibration CMG for small agile satellite applications". In: *Acta Astronautica* 131, pp. 113–122. DOI: `10.1016/j.actaastro.2016.10.021`.

Kim, Sun Yong, Chris K. Mechefske, and Il Yong Kim (2013). "Optimal damping layout in a shell structure using topology optimization". In: *Journal of Sound and Vibration* 332.12, pp. 2873–2883. DOI: `10.1016/j.jsv.2013.01.029`.

Kolk, Max van der et al. (2017). "Multi-material topology optimization of viscoelastically damped structures using a parametric level set method". en. In: *Journal of Vibration and Control* 23.15, pp. 2430–2443. DOI: `10.1177/1077546315617333`.

Kwon, Seong-Cheol, Mun-Shin Jo, and Hyun-Ung Oh (2017). "Experimental Validation of Fly-Wheel Passive Launch and On-Orbit Vibration Isolation System by Using a Superelastic SMA Mesh Washer Isolator". en. In: *International Journal of Aerospace Engineering* 2017.1, p. 5496053. DOI: `10.1155/2017/5496053`.

Lee, W.-S. and S.-K. Youn (2004). "Topology optimization of rubber isolators considering static and dynamic behaviours". en. In: *Structural and Multidisciplinary Optimization* 27.4, pp. 284–294. DOI: `10.1007/s00158-004-0376-1`.

Li, Quhao et al. (2021). "Topology optimization of vibrating structures with frequency band constraints". en. In: *Structural and Multidisciplinary Optimization* 63.3, pp. 1203–1218. DOI: 10.1007/s00158-020-02753-7.

Li, Quhao et al. (2022). "Nonlinear eigenvalue topology optimization for structures with frequency-dependent material properties". In: *Mechanical Systems and Signal Processing* 170, p. 108835. DOI: 10.1016/j.ymssp.2022.108835.

Liu, Chunchuan et al. (2015). "Recent advances in micro-vibration isolation". In: *Mechanical Systems and Signal Processing* 56-57, pp. 55–80. DOI: 10.1016/j.ymssp.2014.10.007.

Liu, Jikai et al. (2018). "Current and future trends in topology optimization for additive manufacturing". en. In: *Structural and Multidisciplinary Optimization* 57.6, pp. 2457–2483. DOI: 10.1007/s00158-018-1994-3.

Luo, Haitao et al. (2023). "Design and experiment of micro-vibration isolation system for optical satellite". In: *European Journal of Mechanics - A/Solids* 97, p. 104833. DOI: 10.1016/j.euromechsol.2022.104833.

Luo, Qing, Jun Wu, and Caizhi Fan (2020). "Design and Modeling of Micro-vibration Isolation for Spacecraft Flywheel". In: *2020 3rd International Conference on Mechatronics, Robotics and Automation (ICMRA)*, pp. 12–17. DOI: 10.1109/ICMRA51221.2020.9398352.

Luo, Zhen, Jingzhou Yang, and Liping Chen (2006). "A new procedure for aerodynamic missile designs using topological optimization approach of continuum structures". In: *Aerospace Science and Technology* 10.5, pp. 364–373. DOI: 10.1016/j.ast.2005.12.006.

Ma, Zheng-Dong et al. (1995). "Topological Optimization Technique for Free Vibration Problems". en. In: *Journal of Applied Mechanics* 62.1, pp. 200–207. DOI: 10.1115/1.2895903.

Moreira, R. A. S. (2014). "Structural Dynamics and Viscoelastic Passive Damping Treatments". en. In: *Modern Mechanical Engineering: Research, Development and Education*. Ed. by J. Paulo Davim. Berlin, Heidelberg: Springer, pp. 89–107. DOI: 10.1007/978-3-642-45176-8_5.

Moreira, R.A.S., J.D. Corte-Real, and J. Dias Rodrigues (2010). "A Gneralized Frequency-Temperature Viscoelastic Model". en. In: *Shock and Vibration* 17.4-5, pp. 407–418. DOI: 10.1155/2010/463963.

Nakra, B. C. (1998). "VIBRATION CONTROL IN MACHINES AND STRUCTURES USING VISCOELASTIC DAMPING". In: *Journal of Sound and Vibration* 211.3, pp. 449–466. DOI: 10.1006/jsvi.1997.1317.

Niu, Bin et al. (2018). "On objective functions of minimizing the vibration response of continuum structures subjected to external harmonic excitation". en. In: *Structural and Multidisciplinary Optimization* 57.6, pp. 2291–2307. DOI: 10.1007/s00158-017-1859-1.

Noh, Jin Yee and Gil Ho Yoon (2012). "Topology optimization of piezoelectric energy harvesting devices considering static and harmonic dynamic loads". In: *Advances in Engineering Software* 53, pp. 45–60. DOI: 10.1016/j.advengsoft.2012.07.008.

Olhoff, N. (1989). "Multicriterion structural optimization via bound formulation and mathematical programming". en. In: *Structural optimization* 1.1, pp. 11–17. DOI: `10.1007/BF01743805`.

Ou, J. S. and N. Kikuchi (1996). "Integrated optimal structural and vibration control design". en. In: *Structural optimization* 12.4, pp. 209–216. DOI: `10.1007/BF01197358`.

Pedersen, N.L. (2000). "Maximization of eigenvalues using topology optimization". en. In: *Structural and Multidisciplinary Optimization* 20.1, pp. 2–11. DOI: `10.1007/s001580050130`.

Rade, D. A. et al. (2019). "Passive Vibration Control Using Viscoelastic Materials". en. In: *Nonlinear Structural Dynamics and Damping*. Ed. by Juan Carlos Jauregui. Cham: Springer International Publishing, pp. 119–168. DOI: `10.1007/978-3-030-13317-7\_5`.

Shi, H. T et al. (2024). "Vibration isolation methods in spacecraft: A review of current techniques". In: *Advances in Space Research* 73.8, pp. 3993–4023. DOI: `10.1016/j.asr.2024.01.020`.

Sigmund, O. (2001). "A 99 line topology optimization code written in Matlab". en. In: *Structural and Multidisciplinary Optimization* 21.2, pp. 120–127. DOI: `10.1007/s001580050176`.

Sigmund, Ole and Kurt Maute (2013). "Topology optimization approaches". en. In: *Structural and Multidisciplinary Optimization* 48.6, pp. 1031–1055. DOI: `10.1007/s00158-013-0978-6`.

Svanberg, Krister (1987). "The method of moving asymptotes—a new method for structural optimization". en. In: *International Journal for Numerical Methods in Engineering* 24.2, pp. 359–373. DOI: `10.1002/nme.1620240207`.

Wang, Fengwen, Boyan Stefanov Lazarov, and Ole Sigmund (2011). "On projection methods, convergence and robust formulations in topology optimization". en. In: *Structural and Multidisciplinary Optimization* 43.6, pp. 767–784. DOI: `10.1007/s00158-010-0602-y`.

Wu, Fan, Pu Xue, and Xin Zhang (2024a). "Dynamic topology optimization of two-phase viscoelastic composite structures considering frequency- and temperature-dependent properties under random excitation". In: *Engineering Optimization* 0.0, pp. 1–28. DOI: `10.1080/0305215X.2024.2346988`.

Wu, Fan et al. (2023). "Multi-scale concurrent topology optimization of frequency- and temperature-dependent viscoelastic structures for enhanced damping performance". en. In: *Structural and Multidisciplinary Optimization* 66.11, p. 234. DOI: `10.1007/s00158-023-03681-y`.

Wu, Qiangbo, Quhao Li, and Shutian Liu (2024b). "A method for eliminating local modes caused by isolated structures in dynamic topology optimization". In: *Computer Methods in Applied Mechanics and Engineering* 418, p. 116557. DOI: `10.1016/j.cma.2023.116557`.

— (2024c). "Dynamic topology optimization for structures exhibiting frequency-dependent material properties with prescribed frequency forbidden band". In: *Computer Methods in Applied Mechanics and Engineering* 432, p. 117439. DOI: `10.1016/j.cma.2024.117439`.

Yi, Guilian and Byeng D. Youn (2016). "A comprehensive survey on topology optimization of phononic crystals". en. In: *Structural and Multidisciplinary Optimization* 54.5, pp. 1315–1344. DOI: `10.1007/s00158-016-1520-4`.

Yun, Kyeong-Soo and Sung-Kie Youn (2017). "Multi-material topology optimization of viscoelastically damped structures under time-dependent loading". In: *Finite Elements in Analysis and Design* 123, pp. 9–18. DOI: 10.1016/j.finel.2016.09.006.

Zargham, Sajjad et al. (2016). "Topology optimization: a review for structural designs under vibration problems". en. In: *Structural and Multidisciplinary Optimization* 53.6, pp. 1157–1177. DOI: 10.1007/s00158-015-1370-5.

Zhang, Dongdong et al. (2022). "Topology optimization of constrained layer damping plates with frequency- and temperature-dependent viscoelastic core via parametric level set method". In: *Mechanics of Advanced Materials and Structures* 29.1, pp. 154–170. DOI: 10.1080/15376494.2021.1938302.

Zhang, Guodong and Kapil Khandelwal (2019). "Design of dissipative multimaterial viscoelastic-hyperelastic systems at finite strains via topology optimization". en. In: *International Journal for Numerical Methods in Engineering* 119.11, pp. 1037–1068. DOI: 10.1002/nme.6083.

Zhang, Xiaopeng and Zhan Kang (2013). "Topology optimization of damping layers for minimizing sound radiation of shell structures". In: *Journal of Sound and Vibration* 332.10, pp. 2500–2519. DOI: 10.1016/j.jsv.2012.12.022.

Zhang, Z., G.S. Aglietti, and W. Zhou (2011). "Microvibrations induced by a cantilevered wheel assembly with a soft-suspension system". In: *AIAA Journal* 49.5, pp. 1067–1079. DOI: 10.2514/1.J050791.

Zheng, Weiguang et al. (2016). "Topology optimization of PCLD on plates for minimizing sound radiation at low frequency resonance". en. In: *Structural and Multidisciplinary Optimization* 53.6, pp. 1231–1242. DOI: 10.1007/s00158-015-1371-4.

Zhu, Li (2019). *Dynamic topology optimization of plate for vibration suppression*.

Zienkiewicz, Olgierd C., Robert L. Taylor, and Jianzhong Zhu (2010). *The finite element method: its basis and fundamentals*. en. 6. ed., reprint., transferred to digital print. Amsterdam Heidelberg: Elsevier.

A Strengths and drawbacks of passive vibration isolation methods

Table 7: Passive vibration isolation methods with their strengths and drawbacks. Adapted from Shi et al. (2024).

Isolation Method	Strengths	Drawbacks
Viscoelastic materials	Low weight, size and cost, adaptable to small satellites	Low performance at low frequency, sensitive to temperature
Viscous fluid damper	High isolation performance, can be integrated into Stewart hexapod or Octo-platforms	Viscous fluid dependence on temperature, effect of vacuum environment on fluid properties
Mechanical Springs	Light weight, cost and simple to integrate	Trade-off between bearing capacity and isolation performance
Quasi-zero stiffness mechanism	Good isolation performance at low frequencies, no effect of temperature variation on isolation performance, adaptability to small satellite platforms	Effective only with small external excitation
Shape memory alloy	Delamination in case of SMA blades	In case of SMA washer, the tightening torque may affect the characteristics of the washer
Folded beams	Good suppression effect at high frequency, light weight	Low static stiffness to support load, poor performance at low frequency
Metamaterials	Light weight, small volume and easy to integrate, effective over wide frequency range	Limited bearing capacity

B Temperature- and/or frequency-dependent viscoelastic material properties

Viscoelastic materials (VEMs) uniquely combine elastic and viscous behaviours, exhibiting both instantaneous and time-dependent responses to external loads. Purely elastic materials store mechanical energy without dissipation, whereas purely viscous materials dissipate energy as heat. In purely viscous materials, stress is proportional to the strain rate, whereas in elastic materials, stress is proportional to strain (Yun and Youn, 2017). When harmonically loaded at a predefined frequency, the material exhibits a harmonic response at the same frequency, but with a phase lag between the applied load and the resulting strain. This phase lag is directly related to the material's damping characteristics (Moreira et al., 2010).

The stress-strain curve under cyclic loading of VEMs displays a hysteresis loop, depicted in Figure 22, where the enclosed area indicates the energy dissipated as heat. Due to the viscous properties, the mechanical properties of viscoelastic materials are affected by the deformation

rate. As a result, their deformation behaviour at varying rates is depicted by a series of curves instead of a single stress-strain curve (Elmoghazy et al., 2024). This behaviour allows VEMs to absorb and dissipate energy, making them particularly valuable in applications requiring vibration damping or dynamic load resistance (Moreira, 2014).
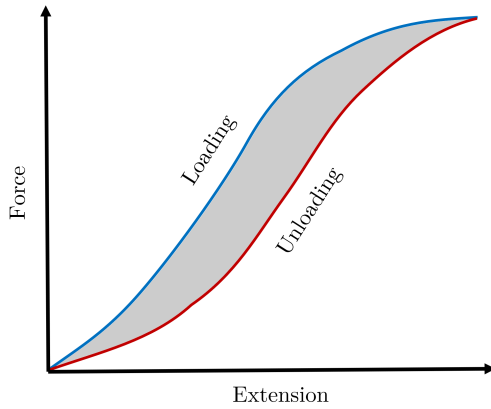

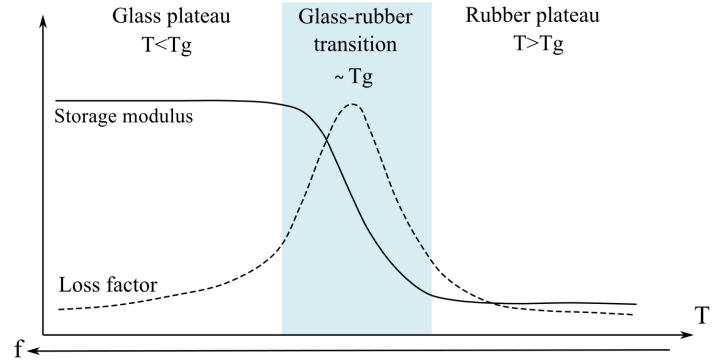Figure 22: Typical hysteresis loop from cyclic loading of VEMs. (Elmoghazy et al., 2024)

Figure 23: The influence of frequency and temperature on the storage modulus and loss factor. (Rade et al., 2019)

The material characteristics of viscoelastic materials have a clear non-linear dependency on frequency and temperature (Moreira et al., 2010). The behaviour of VEMs as a function of temperature can be divided into three distinct regions shown in Figure 23: the glassy state, the glass-to-rubber transition zone, and the rubbery state. In the glassy state, the material exhibits a high storage modulus and a low loss factor, indicative of its rigidity and minimal energy dissipation. The glass-to-rubber transition zone is marked by a significant decrease in the storage modulus and a peak in the loss factor, reflecting substantial energy dissipation and material softening. Finally, in the rubbery state, both the storage modulus and the loss factor become low and nearly constant. These changes are associated with different levels of mobility of the internal polymer chains. At temperatures bellow glass transition temperature T_g and high frequency excitations the polymer chains are rigid and result in high stiffness. When the frequency of applied load is reduced or the material's temperature is increased, polymer chains loosen resulting in a lower material stiffness (Moreira et al., 2010). For engineering applications, viscoelastic materials are typically chosen to operate within the transition zone or the rubber-like phase, as these regions offer the best damping performance (Rade et al., 2019). In applications, where temperature- or frequency- dependent properties are to be considered, various viscoelastic material models can be used such as the generalised Maxwell model (Wu et al., 2023; Zhang et al., 2022) or the Golla-Hughes-McTavish (Wu et al., 2024a; Li et al., 2022; Wu et al., 2024c). However, the characterisation of these models require experimental evaluation of parameters to fit the behaviour of a real-life VEMs to be used.

Time-temperature superposition principle

Interestingly, the dependencies on temperature and frequency exhibit a similar behaviour. In particular, amorphous polymers, one of the most common classes of viscoelastic materials in en-

gineering, demonstrate an equivalence between time or frequency effects and temperature effects above their glass transition temperature T_g (Rade et al., 2019). This equivalence is formally described by the time-temperature superposition principle, which provides a framework to predict material behaviour under varying temperature and frequency conditions. According to this principle, the material's response at a specific time t (or frequency ω) and temperature T can be related to its behaviour at a "reduced" time t_r (or frequency ω_r) and a corresponding "reduced" temperature T_r (Dealy and Plazek, 2009). In the frequency domain, the time-temperature superposition principle can be expressed as shown in Equation 41:

$$
\begin{cases}
\omega_r = a_T\left(T, T_r\right)\omega \\
E'\left(\omega_r, T_r\right) = a_T\left(T, T_r\right) E'\left(\omega, T\right) \\
E''\left(\omega_r, T_r\right) = a_T\left(T, T_r\right) E''\left(\omega, T\right)
\end{cases}
\tag{41}
$$

where $a_T(T, T_r)$ are horizontal shift coefficients applied to the isotherms of storage and loss moduli at temperature T to estimate material properties at another temperature T_r. Consequently, by shifting along the frequency axes, the functions describing the storage modulus and loss factor at various temperatures can be combined into master curves, as shown in Figure 24. Time-temperature superpositions allows to estimate viscoelastic material behaviour at any temperature using test data gathered at other temperatures, which is useful for measuring material properties beyond the capabilities of the test setup (e.g. high temperatures) by testing at a decreased frequency or vice versa.
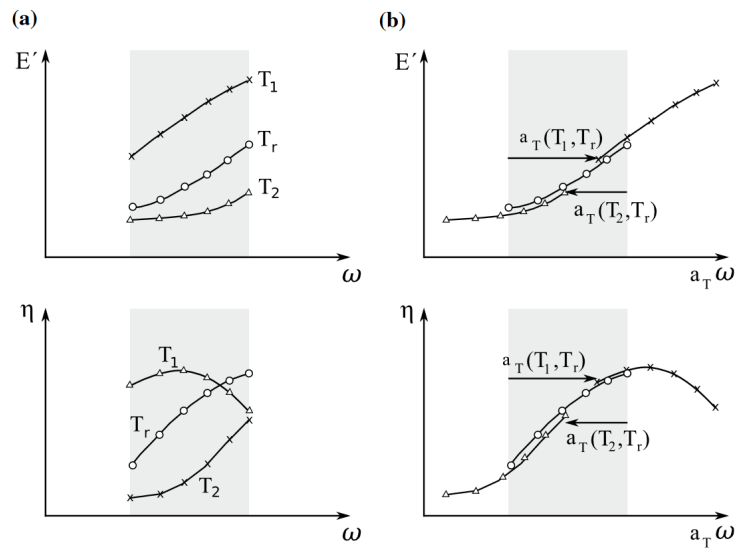

Figure 24: Time-temperature superposition principle: (a) storage modulus and loss factor isotherms at temperatures T_1, T_2, and T_r ($T_1 > T_r > T_2$), (b) temperature shifted storage modulus and loss factor at the reference temperature of T_r. (Rade et al., 2019)

C   Additional results from FEM validation

The plane strain assumption was partially investigated to compare eigenfrequency errors across Poisson ratios and mitigation methods. Due to technical limitations, complete results could not be obtained, and a full set of plane stress results was not tabulated. The available data nonetheless illustrate the tendency of plane strain models to exhibit locking at high Poisson ratios and the effect of some locking mitigation techniques in 2D.

	O^1 (OFF)			O^2 (ON)		
	u-p (OFF)		u-p (ON)	u-p (OFF)	u-p (ON)	
	Python	COMSOL				
Tag	OFF-OFF-OFF	OFF-OFF-OFF	OFF-OFF-ON	ON-OFF-OFF	ON-OFF-ON	Reduced Integration (OFF)
ν = 0.3	0.070%	0.070%	0.044%		0.000%	
ν = 0.48						
ν = 0.49	0.600%	0.778%	0.098%		0.000%	
ν = 0.4999	19.995%	19.995%	0.105%		0.000%	
ν = 0.499999	73.181%	73.185%	0.105%		0.000%	
Tag	OFF-ON-OFF	OFF-ON-OFF	OFF-ON-ON	ON-ON-OFF	ON-ON-ON	Reduced Intergation (ON)
ν = 0.3	0.037%	0.008%	0.008%			
ν = 0.48						
ν = 0.49	0.087%	0.003%	0.003%			
ν = 0.4999	-0.014%	0.003%	0.003%			
ν = 0.499999	-0.014%	0.002%	0.002%			

Average eigen | frequency relat | ive error

Figure 25: Supplementary results for the plane strain assumption. Some data points are missing due to incomplete simulations, but the overall trend illustrates increasing errors at higher Poisson ratios without locking mitigation.

Comment on material assumptions.   For plane stress, eigenfrequencies were consistently lower and largely insensitive to Poisson's ratio, which agrees with the fact that the bulk modulus $K(E, \nu)$ does not diverge as $\nu \to 0.5$. However, plane stress neglects out-of-plane constraints and is therefore not representative of rubber components. The partial plane strain results shown above confirm the expected locking behaviour as $\nu \to 0.5$, with large errors in the `OFF-OFF-OFF` configuration. For completeness, the 3D case is included in the main text (Figure 10), where similar trends appear together with additional discretisation error due to the coarser mesh. In practice, plane strain was adopted as the 2D assumption, while full 3D elasticity remains the most physically accurate representation.