

The Pennsylvania State University
The Graduate School

**OPTIMIZATION OF ACOUSTIC BLACK HOLES
FOR VIBRATION REDUCTION**

A Dissertation in
Acoustics
by
Cameron A. McCormick

© 2021 Cameron A. McCormick

Submitted in Partial Fulfillment
of the Requirements
for the Degree of

Doctor of Philosophy

December 2021

The dissertation of Cameron A. McCormick was reviewed and approved by the following:

Micah R. Shepherd
Associate Research Professor, Applied Research Laboratory
Dissertation Advisor
Chair of Committee

Stephen C. Conlon
Research Professor, Applied Research Laboratory

Mary I. Frecker
Reiss Chair of Engineering
Director, The Penn State Center for Biodevices
Professor of Mechanical Engineering

Julianna C. Simon
Assistant Professor of Acoustics and Biomedical Engineering

Victor W. Sparrow
Director, Graduate Program in Acoustics
United Technologies Corporation Professor of Acoustics

Abstract

A problem that frequently arises in noise control engineering is the need for lightweight structures that are also quiet. In many practical instances of transportation and aerospace design, stiff, lightweight panels are used to reduce overall weight. However, these stiff, lightweight panels are also highly receptive to vibration and can become efficient radiators of sound. Although traditional methods of passive damping can effectively combat this behavior, they usually lead to an increase in overall weight. By modifying the plate thickness according to a specific power law profile, so-called ‘acoustic black holes’ (ABHs) offer a compelling solution to the need for both broadband vibration reduction and reduced weight. ABH vibration absorbers have been shown to be effective at reducing structural vibration, often with a net reduction in overall weight. However, a comprehensive understanding of what constitutes optimal ABH design in finite structures has been historically lacking. This dissertation systematically investigates the many considerations of ABH design for vibration reduction. An optimization framework is developed that combines the modelling of structural dynamics with multi-objective evolutionary optimization, and this framework is used to determine the optimal designs for various ABH applications. By analyzing trends in the optimization results, the tradeoffs inherent in ABH design are illuminated in a more holistic manner than has been undertaken to date.

Table of Contents

List of Figures	vii
List of Tables	xv
Acknowledgments	xvii
Chapter 1	
Introduction	1
1.1 Historical background	1
1.2 Review of ABH literature related to optimal design	3
1.2.1 Review of parametric studies	3
1.2.2 Review of optimization studies	7
1.3 Theory	11
1.3.1 Analytical ABH theory	11
1.3.2 Reflection from an ABH termination	14
1.3.3 Multi-objective evolutionary optimization and the Borg algorithm	17
1.3.4 Optimization framework	20
1.4 Outline of the dissertation	22
Chapter 2	
A comparison of optimization algorithms: Gradient-based versus evolutionary	24
2.1 Introduction	24
2.2 Problem description	25
2.3 Theory	27
2.3.1 Transfer matrix method	27
2.3.2 Smoothing filters	30
2.4 Procedure	31
2.5 Results and discussion	33
2.5.1 $\omega \in \mathcal{W}_{\text{BB}}$	34
2.5.2 $\omega \in \mathcal{W}_{\text{LF}}$	37
2.5.3 $\omega \in \mathcal{W}_{\text{HF}}$	37
2.5.4 Comparing <code>fmincon</code> and Borg	40
2.6 Conclusions	45

Chapter 3	
Optimization of an ABH at the end of a cantilever beam	47
3.1 Introduction	47
3.2 Problem description	48
3.3 Theory	50
3.3.1 Damping model	50
3.3.2 Block Riccati transfer matrix method	53
3.4 Procedure	59
3.5 Results and discussion	61
3.6 Conclusions	69
Chapter 4	
Optimization and comparison of three styles of ABH	71
4.1 Introduction	71
4.2 Problem description	75
4.3 Procedure	76
4.4 Results and discussion	81
4.5 Additional comparison	87
4.5.1 Static compliance	87
4.5.2 Radiated sound power	91
4.5.3 Heterogeneous ABH parameters	92
4.6 Conclusions	93
Chapter 5	
A generalized measure of ABH performance	95
5.1 Introduction	95
5.2 Problem description	96
5.3 Theory	99
5.3.1 Elastodynamic theory	99
5.3.2 FEM formulation and preconditioning	101
5.3.3 Dissipated power formulation	106
5.4 Procedure	109
5.4.1 Overview of deal.II	110
5.4.2 Parallelization	112
5.4.3 Control of shear locking	114
5.5 Results and discussion	119
5.6 Conclusions	122
Chapter 6	
Conclusions	124
6.1 Summary of findings	124
6.2 Suggestions for future work	128

Appendix A
 Mesh convergence study **131**
 A.1 Introduction 131
 A.2 Test case description 133
 A.3 Results 135
 A.4 Conclusions 137

Appendix B
 FEM meshing calculations **140**

Bibliography **146**

List of Figures

1.1	Diagram of the ideal ABH wedge shape at the end of a uniform beam or plate. Note that Mironov’s original analysis was only concerned with the region $[0, x_i]$, agnostic to the rest of the structure.	1
1.2	Example Pareto front (solid curve) for a two-objective optimization problem [1]. The two objectives are f_1 and f_2 . The dashed curve represents the boundary of the realizable objective space, \mathbf{F}	18
1.3	Flowchart representation of the process for formulating an optimization problem. Note that this process is not prescriptive and is only a formalization of the author’s own process.	21
2.1	Graphical depiction of the current optimization problem. The beam geometry is represented by 100 beam segments, where each segment has constant thickness. The left-hand boundary condition is an imposed shear, while the right-hand boundary condition is zero displacement and zero rotation, i.e., fixed. The microphone region, (x_a, x_b) , has constant thickness h_{mic} throughout.	25
2.2	Flowchart representation of the design evaluation function called by both optimization algorithms.	32
2.3	Optimization results for the broadband case, $\omega \in \mathcal{W}_{\text{BB}}$, using <code>fmincon</code> . The top plot (a) shows the spectrum of $J_\omega(\omega, \mathbf{h})$ in decibels for the optimal beam design (green) and a reference beam with a uniform thickness of 5 mm (blue). The evaluation frequencies are shown as red hatch marks. The bottom left plot (b) shows the optimal thickness distribution (blue) compared to the reference beam (red). The bottom right plot (c) shows the displacement amplitude, $ U $, along the beam at $\omega = 2\pi \times 300$ Hz for the optimal beam design (green) and the reference beam (blue). The bounds of the microphone region are indicated by two red vertical lines.	34

- 2.4 Optimization results for the broadband case, $\omega \in \mathcal{W}_{\text{BB}}$, using Borg. The top plot (a) shows the spectrum of $J_\omega(\omega, \mathbf{h})$ in decibels for the optimal beam design (green) and a reference beam with a uniform thickness of 5 mm (blue). The evaluation frequencies are shown as red hatch marks. The bottom left plot (b) shows the optimal thickness distribution (blue) compared to the reference beam (red). The bottom right plot (c) shows the displacement amplitude, $|U|$, along the beam at $\omega = 2\pi \times 300$ Hz for the optimal beam design (green) and the reference beam (blue). The bounds of the microphone region are indicated by two red vertical lines. 35
- 2.5 Figure 3 from [2], corresponding to Figures 2.3 and 2.4. Note that Berggren et al. use the variables D and A in place of h and U , respectively. 36
- 2.6 Optimization results for the low-frequency case, $\omega \in \mathcal{W}_{\text{LF}}$, using `fmincon`. The top plot (a) shows the spectrum of $J_\omega(\omega, \mathbf{h})$ in decibels for the optimal beam design (green) and a reference beam with a uniform thickness of 5 mm (blue). The evaluation frequencies are shown as red hatch marks. The bottom left plot (b) shows the optimal thickness distribution (blue) compared to the reference beam (red). The bottom right plot (c) shows the displacement amplitude, $|U|$, along the beam at $\omega = 2\pi \times 300$ Hz for the optimal beam design (green) and the reference beam (blue). The bounds of the microphone region are indicated by two red vertical lines. 38
- 2.7 Optimization results for the low-frequency case, $\omega \in \mathcal{W}_{\text{LF}}$, using Borg. The top plot (a) shows the spectrum of $J_\omega(\omega, \mathbf{h})$ in decibels for the optimal beam design (green) and a reference beam with a uniform thickness of 5 mm (blue). The evaluation frequencies are shown as red hatch marks. The bottom left plot (b) shows the optimal thickness distribution (blue) compared to the reference beam (red). The bottom right plot (c) shows the displacement amplitude, $|U|$, along the beam at $\omega = 2\pi \times 300$ Hz for the optimal beam design (green) and the reference beam (blue). The bounds of the microphone region are indicated by two red vertical lines. 39
- 2.8 Figure 5 from [2], corresponding to Figures 2.6 and 2.7. Note that Berggren et al. use the variables D and A in place of h and U , respectively. 40

- 2.9 Optimization results for the high-frequency case, $\omega \in \mathcal{W}_{\text{HF}}$, using `fmincon`. The top plot (a) shows the spectrum of $J_\omega(\omega, \mathbf{h})$ in decibels for the optimal beam design (green) and a reference beam with a uniform thickness of 5 mm (blue). The evaluation frequencies are shown as red hatch marks. The bottom left plot (b) shows the optimal thickness distribution (blue) compared to the reference beam (red). The bottom right plot (c) shows the displacement amplitude, $|U|$, along the beam at $\omega = 2\pi \times 2300$ Hz for the optimal beam design (green) and the reference beam (blue). The bounds of the microphone region are indicated by two red vertical lines. 41
- 2.10 Optimization results for the high-frequency case, $\omega \in \mathcal{W}_{\text{HF}}$, using Borg. The top plot (a) shows the spectrum of $J_\omega(\omega, \mathbf{h})$ in decibels for the optimal beam design (green) and a reference beam with a uniform thickness of 5 mm (blue). The evaluation frequencies are shown as red hatch marks. The bottom left plot (b) shows the optimal thickness distribution (blue) compared to the reference beam (red). The bottom right plot (c) shows the displacement amplitude, $|U|$, along the beam at $\omega = 2\pi \times 2300$ Hz for the optimal beam design (green) and the reference beam (blue). The bounds of the microphone region are indicated by two red vertical lines. 42
- 2.11 Figure 6 from [2], corresponding to Figures 2.9 and 2.10. Note that Berggren et al. use the variables D and A in place of h and U , respectively. 43
- 2.12 Optimal thickness distribution (blue) and displacement amplitude at $\omega = 2\pi \times 2557.32$ Hz (red) for the high-frequency case, $\omega \in \mathcal{W}_{\text{HF}}$, using Borg. Clearly, at 2557.32 Hz the vibration energy is concentrated at the left end away from the microphone region, (x_a, x_b) 43
- 2.13 Optimal thickness distribution (blue) and displacement amplitude (red) for the high-frequency case, $\omega \in \mathcal{W}_{\text{HF}}$, using Borg. Analysis frequencies are $\omega = 2\pi \times 2191.23$ Hz (left) and $\omega = 2\pi \times 2902.02$ Hz (right). Whereas at 2557.32 Hz the vibration energy is concentrated at the drive point, here there is very little movement at the left end. 44
- 3.1 Graphical illustration of the optimization problem. The beam has overall length L_x and the base beam has thickness h_1 , while the ABH taper is determined by its length, L_{ABH} , its minimum thickness, h_0 , and its taper power, m . The composite structure is driven at a point, x_f , with a point force of magnitude F . The boundary at $x = 0$ is modelled as free and the boundary at $x = L_x$ is modelled as fixed. 48
- 3.2 Graphic of a two-layer beam. The neutral axis is denoted by a dashed line. 51

3.3	Illustration of the instabilities that arise from the basic transfer matrix method. Pictured are drive-point mobilities for a simply-supported beam of length L	54
3.4	Flowchart representation of the design evaluation function used in the optimization.	60
3.5	Convergence study used to determine the number of segments per wavelength to use with the transfer matrix method. The relative error between the ‘true’ response of an example geometry and the response using the given number of segments per wavelength is shown as a solid line. The corresponding runtime of the solver in seconds is given as a dashed line.	62
3.6	Evolution of the Borg algorithm’s archive during the optimization, in terms of the design variables and objective value, $J(\mathbf{h})$, after each function evaluation.	63
3.7	Reference beam of uniform thickness (top) and optimal ABH thickness profile (bottom) with added damping layer (black). The left-hand edge of each beam is modelled as free and the right-hand edge is modelled as fixed in the calculation of U . Dimensions are not to scale.	63
3.8	Spatially-averaged squared velocity, $\langle \dot{U}^2 \rangle$, of a reference beam with uniform thickness (dashed) and of the optimal ABH thickness profile found by Borg (solid). Note that the scale of the x-axis is logarithmic.	64
3.9	Spatially-averaged squared velocity, $\langle \dot{U}^2 \rangle$, of a reference beam with uniform thickness (dashed) and of the ‘classically optimal’ ABH thickness profile that produces minimal reflection (solid). Note that the scale of the x-axis is logarithmic.	65
3.10	Average reflection coefficient magnitudes, $ R _{\text{avg}}$, for taper designs having length $L_{\text{ABH}} = 22.26$ cm and minimum thickness $h_0 = 0.635$ mm, but with varying taper power, m . The reflection coefficient was calculated in the same manner as [3]. Note that although the average values are close to one, the reflection coefficient magnitudes were shown to reach as low as 0.6 at particular frequencies.	66
3.11	Length of the ABH taper in which Equation (1.7) holds, relative to the total taper length, for the optimal design and the classically optimal design. Equation (1.7) is here considered satisfied when $\left \frac{1}{k^2} \frac{dk}{dx} \right \leq 0.3$	67

4.1	Three styles of one-dimensional acoustic black hole vibration absorbers (not to scale). The three styles are (a) Standard symmetric; (b) Standard nonsymmetric; and (c) Double-leaf. Added damping is indicated in black. Note that the colors used for each style in this figure will be used throughout the chapter to indicate each color's respective style.	72
4.2	Example mode shapes for nominal ABH designs (not to scale). Modal frequencies are (a1) 405.8 Hz and (a2) 551.5 Hz for the standard symmetric style; (b1) 405.6 Hz and (b2) 551.7 Hz for the standard nonsymmetric style; and (c1) 43.5 Hz, (c2) 232.4 Hz, and (c3) 373.9 Hz for the double-leaf style.	73
4.3	Modal densities for the nominal designs shown in Figure 4.2. (●) Standard symmetric. (■) Standard nonsymmetric. (▼) Double-leaf, axially symmetric modes. (▲) Double-leaf, axially anti-symmetric modes.	74
4.4	Example of how a given design might be divided for the purposes of meshing. The leftmost region will be the same for all designs, as its purpose is to guarantee nodes exist at $x = 0$ for the constraints and nodes exist at $x = x_f$ for the applied forces.	77
4.5	Flowchart representation of the recursive function used to mesh the ABH region. N is calculated according to Equation (4.5).	79
4.6	Flowchart representation of the design evaluation function used by Borg in the optimization process.	82
4.7	Approximate Pareto fronts for two objectives and three styles. (a) Standard symmetric. (b) Standard nonsymmetric. (c) Double-leaf. The objective values are also shown for an unmodified beam (x) and for an unmodified beam with the maximum damping allowed within the variable bounds (+).	83
4.8	Graphical depiction (not to scale) of the design of each style that minimizes J_2 ('Min Response' in Table 4.2). (a) Standard symmetric; (b) Standard nonsymmetric; and (c) Double-leaf. Also depicted is (d) An unmodified beam with the maximum damping allowed within the variable bounds (represented by a + symbol in Figure 4.7).	84
4.9	Trends for each design variable along the three Pareto fronts in Figure 4.7, from the design that minimizes J_1 ('min mass') to the design that minimizes J_2 ('min response').	85

4.10	Dependence of J_2 on damping along the Pareto front. The dashed lines represent the length of damping, L_d , required to achieve 99% of the possible reduction in J_2 for each ABH style.	86
4.11	Spatially-averaged square velocity spectra for the design of each style that minimizes J_2 ('Min Response' in Table 4.2). (a) Standard symmetric. (b) Standard nonsymmetric. (c) Double-leaf. Also shown is the spectrum for an unmodified beam with the maximum damping allowed within the variable bounds (dashed).	88
4.12	Deformation (not to scale) corresponding to the values in Table 4.3. Shown is deformation due to axial compression for (a1) Standard symmetric, (b1) Standard nonsymmetric, and (c1) Double-leaf designs, as well as deformation due to torque for (a2) Standard symmetric, (b2) Standard nonsymmetric, and (c2) Double-leaf designs. Skeletons of the undeformed geometries are also shown in light grey.	90
4.13	Radiated sound power spectrum in one-third octaves for each of the designs that minimize J_2 ('Min Response' in Table 4.2). Also shown is the one-third octave radiated sound power for an unmodified beam with the maximum damping allowed within the variable bounds.	92
4.14	Approximate contour plot of J_2 as a function of the two taper powers for each half of the standard nonsymmetric ABH style. m_1 is the taper power of the left half of the ABH. m_2 is the taper power of the right half of the ABH. The remaining design variables are those under 'Min Response' in Table 4.2.	93
5.1	Graphical representation of the optimization problem under consideration. Depicted are (a) Top; (b) Bottom; and (c) Cross-sectional views of the geometry. The geometry is a square plate of side length a and thickness h_1 , with an embedded two-dimensional ABH feature and an attached free damping layer. The damping layer covers the same area of the ABH, which is a circle with radius R_{ABH} . The taper profile of the ABH is controlled by the minimum thickness, h_0 , and the power law exponent, m . Not pictured is the parameter γ , which alters the gradient of the ABH.	97
5.2	Flowchart representation of the multi-level parallelization framework used. The manager process interfaces with Borg to interpret the design variables and to calculate the objectives. Note that workers 1–4 share the same analysis frequencies but each own a different portion of the degrees of freedom. The same is true for workers 5–8. In this way, the construction of matrices and vectors is also parallelized.	113

5.3	An example of drive point mobility demonstrating artificial stiffness due to locking phenomena. The solid curve represents the mobility as calculated by NASTRAN with shear and volumetric locking control. The dashed line represents the result when no measures are taken to relieve locking. Clearly, the dashed line shows peaks that are shifted up in frequency, indicating increased stiffness.	114
5.4	A 2D reference cell. The points represent the quadrature points for 2×2 Gaussian quadrature. ξ and η (not to be confused with loss factor) are the coordinates in the reference domain.	116
5.5	A 3D reference cell. The points represent the quadrature points for $2 \times 2 \times 2$ Gaussian quadrature. ξ , η , and ζ are the coordinates in the reference domain. Although the axes are shown off to the corner for ease of readability, the origin is usually taken to be the center of the cell. . . .	117
5.6	Same mobility as in Figure 5.3, but now the dashed line includes selective substitution to relieve shear locking.	118
5.7	Approximate Pareto front for the current problem. Also indicated is the ‘knee’ of the front, at which the trade-off between the two objectives is approximately equal. Although the set of designs found by Borg does not include any with J_1 greater than 0.4, the scale of the abscissa is such that the figure shows all possible values for J_1 . Conversely, the values for J_2 are potentially infinite, and so it has been limited.	119
5.8	Variable trends for the Pareto front in Figure 5.7. The demarcation ‘max dissipation’ indicates the end at which J_2 is minimized, while the ‘min damping’ indicates the end at which J_1 is minimized. The location of the knee is also shown as a circular marker.	120
5.9	Ratio of power dissipated by the damping to power input to the system, as a function of ka . The spectrum for the ABH design that optimizes J_2 is shown as a solid line. The spectrum for a uniform plate with the equivalent added damping layer is shown as a dashed line. Note that the first modal frequency of the unmodified plate occurs at around $ka = 6$	121
A.1	Graphical depiction (not to scale) of the test geometry whose material parameters and dimensions are given in Table A.1. Green denotes the beam material, yellow denotes the Damping 1 material, and orange denotes the Damping 2 material.	133

A.2	Example mesh, showing the transition regions using pyramidal elements. This happens to be the mesh used for twelve elements per wavelength with linear elements.	134
A.3	Excitation displacement pulse (black), along with an example of the response at the start of the ABH taper (blue) and at the end of the taper (yellow).	135
A.4	Same time series as shown in Figure A.3, except that the signals have been filtered with a cosine filter of 0.5 ms width. Also shown are the envelopes calculated using the Hilbert transform (dotted lines).	136
A.5	Ratio of taper end amplitude to taper start amplitude for various combinations of mesh resolution and element type. Note that the theoretical ratio in the absence of any loss mechanisms would be approximately 7, according to WKB theory.	137
A.6	Relative L^2 error in the amplitude ratios shown in Figure A.5, using the results of 21 quadratic elements per wavelength as ‘true’. Note the logarithmic ordinate scale. Because the error between 21 quadratic elements per wavelength and itself is zero, it is not shown.	138

List of Tables

2.1	Parameters used in the optimization problem.	33
2.2	Summary of optimization results for the three frequency sets. Most optimal objective value, J , along with the mass, m , and compliance, C , of the corresponding design, normalized by their respective constraints. .	44
2.3	Performance comparison of the two algorithms. Because Borg will restart more frequently after the first restart, the objective value, J , and the number of function evaluations (NFEs) are given at the first restart. In contrast, the J and NFEs reported for <code>fmincon</code> are the averages for all restarts.	44
3.1	Parameters used in the optimization problem.	61
3.2	Optimal design variables for the problem described in Section 3.2, and the resultant objective function value.	62
3.3	Objective function values for the objective used in the optimization, $J(\mathbf{h})$, and two auxiliary objectives, $J_M(\mathbf{h})$ and $J_{KE}(\mathbf{h})$, measuring overall mass and average kinetic energy, respectively.	69
4.1	Variable and parameter bounds of the design problem and geometric and material properties of the system. *For the standard nonsymmetric style, $h_d = 3$ mm.	80
4.2	Designs from the three Pareto fronts in Figure 4.7 that minimize J_1 ('Min Mass') or minimize J_2 ('Min Response') for each of the three styles. . . .	84

4.3	Total absolute elemental strain values calculated by Equation (4.6) as a result of an axial compression force and a torque applied at the ends of the beam. Note that these are different from normal strain and shear strain. The designs are those that minimize J_2 ('Min Response' in Table 4.2), as well as an unmodified beam for reference.	89
5.1	Table of parameters used in the optimization.	109
A.1	Material parameters and dimensions of the test case. Uniform 1 is the uniform portion to the left of the ABH taper, and Uniform 2 is the uniform portion to the right of it. For the ABH taper profile, $\varepsilon = 3.93 \text{ m}^{-1}$ and $h_0 = 0.508 \text{ mm}$	133

Acknowledgments

Given the title of this section, it seems appropriate to first acknowledge that I am an able-bodied, cisgender, heterosexual, white man who was born into a wealthy, educated, Christian, English-speaking family in the United States. As such, there have been and continue to be many privileges and opportunities afforded me that are not afforded many, many others. It was never a question whether I would attend tertiary education, nor was it ever a question whether I could count on my family to support me both emotionally and financially. I never had to work a job unrelated to my studies in order to make ends meet. I have never gone a day without food, nor have I ever been forced to spend a night without a roof overhead. I have been incredibly privileged to be able to devote as much time as I wished to my studies. Whatever I have personally accomplished in completing this dissertation, it has been built on a foundation of extreme privilege.

I would also be remiss if I did not acknowledge the many people who have helped and supported me along my educational career. It is unlikely that I will remember everyone, but those who come to mind in this moment are my doctoral advisor, Dr. Micah Shepherd; Dr. Lynn Koplitz at Loyola University; Dr. Jordan Cheer at the ISVR; the other professors at Penn State, the ISVR, and Loyola; the staff at Penn State, the ISVR, and Loyola; my parents, bonus parents, grandparents, and brothers; my New Orleans family; my wife; and my many friends and colleagues around the world (especially the Biddies).

Thank you all.

Chapter 1 | Introduction

1.1 Historical background

Many applications of noise control engineering involve reducing the bending vibrations of beam- or plate-like structures. In addition to structural fatigue from the vibration itself, bending waves are the primary source of radiated sound and can result in high levels of noise [4]. A common means of achieving broadband damping of bending waves is to apply a layer of absorbing material (usually a viscoelastic polymer) to the entire surface of the structure [5]. However, this solution also means contributing additional weight to the structure, which can be disadvantageous in industries such as transport and aerospace. In this sense, so-called ‘acoustic black holes’ (ABHs) offer a compelling solution to the need for both broadband vibration reduction and reduced weight.

Mironov was the first to describe the theoretical possibility of zero reflection of bending waves from the tip of a wedge whose thickness, h , decreases to zero according to the relation $h(x) = \varepsilon x^m$ for constants $\varepsilon > 0$ and $m \geq 2$ [6]. A diagram of such a wedge is shown in Figure 1.1. For thin beams, the Euler-Bernoulli beam model is valid and the




Figure 1.1. Diagram of the ideal ABH wedge shape at the end of a uniform beam or plate. Note that Mironov’s original analysis was only concerned with the region $[0, x_i]$, agnostic to the rest of the structure.

wavenumber, k , phase speed, c_p , and group speed, c_g are*

$$k^4 = \frac{\omega^2 \rho A}{D}, \quad (1.1)$$

$$c_p = \frac{\omega}{k} = \left(\frac{\omega^2 D}{\rho A} \right)^{\frac{1}{4}}, \quad (1.2)$$

and

$$c_g = \frac{\partial \omega}{\partial k} = \frac{1}{2\omega} \frac{\partial}{\partial k} (\omega^2) = \frac{2 k^3 D}{\omega \rho A} = 2 \frac{\omega}{k} = 2c_p, \quad (1.3)$$

respectively, where ω is the radial frequency of interest, ρ is the density of the beam material, A is the cross-sectional area, and D is the bending stiffness, otherwise known as flexural rigidity. Using the above definition of group speed, the total transit time, τ , of a wave packet travelling from a point $x_i > 0$ to the origin is then

$$\tau = \int_{x_i}^0 -\frac{1}{c_g(x)} dx = \frac{1}{2\omega} \int_0^{x_i} k(x) dx. \quad (1.4)$$

In the case that the thickness of the wedge goes as $h(x) = \varepsilon x^m$, $k(x)$ is proportional to $(\varepsilon x^m)^{-\frac{1}{2}}$. From this, it is clear that the integral diverges for $m \geq 2$. This means that the transit time becomes infinite and the wave never reaches the tip of the wedge. Because it never reaches the tip, it can never reflect back, and therefore perfect absorption occurs.

Of course, limitations on real-world production techniques mean that such an ideal wedge is not possible, and there will always be a nonzero truncation at the tip of the wedge. In this case, the ABH effect is broken and can result in reflection of upwards of 70% of the incident energy [7]. However, Krylov theoretically [8] and experimentally [9] demonstrated that the shortcomings of real wedges can be drastically improved by adding thin layers of absorbing material near the tip. In such a case, the ABH vibration absorber works by reducing the bending wavelength and increasing the transverse amplitude, thereby focusing strain in a small region where damping material can more effectively absorb and dissipate energy. This is the principle behind the ‘acoustic black hole effect’ on which ABH vibration absorbers are based. The targeted application of damping material, combined with the removal of material to generate the ABH shape, means that ABH vibration absorbers can produce effective broadband vibration reduction with a net reduction in overall weight.

*Mironov’s original analysis used thin plate theory in one dimension instead of thin beam theory. However, the two wavenumbers differ only by a constant factor and so the end result is effectively the same.

Bowyer and Krylov confirmed experimentally that two-dimensional circular indentations whose thickness decreases according to the relation in Figure 1.1 can produce effective reduction in bending vibration [10] and sound radiation [11] in rectangular plates. Since then, much work has been done investigating the physical behavior and potential applications for such one- and two-dimensional ABH vibration absorbers. ABH literature includes investigations into the modal behavior of ABHs [12, 13], the wave energy localization properties of ABHs [14–16], how manufacturing imperfections and nonlinear behavior can impact ABH performance [17–20], as well as several unique analysis techniques applied to ABH behavior [21, 22]. Although research into ABHs has grown dramatically over the past two decades, research into the optimal design of ABH vibration absorbers has been lacking. As such, a comprehensive understanding of what constitutes optimal ABH design has not been demonstrated in the literature, especially for ABHs embedded within finite structures. The work in this dissertation seeks to fill this gap in understanding. By performing a series of ABH optimization studies and analyzing trends in the results, the tradeoffs inherent in ABH design can be illuminated in a more holistic manner than has been undertaken to date.

1.2 Review of ABH literature related to optimal design

As mentioned in Section 1.1, there has been a relative lack of research devoted to optimal ABH design. Moreover, some of the conclusions from what research does exist seem to contradict one another. This section will survey the ABH literature as it relates to ABH optimization through the year 2020. (Note that the work of this dissertation commenced in January of 2017.) Although they do not involve any formal or rigorous optimization, a number of publications are also included that could be broadly classified as ‘parametric studies’ insofar as they evaluate multiple ABH designs. By surveying the relevant ABH literature, the reader will be given context for the work presented in the following chapters of this dissertation.

1.2.1 Review of parametric studies

The first formal parametric study related to ABH design was published by Georgiev et al. in 2011 [3]. The goal was to identify the optimal amount of damping to apply to a one-dimensional ABH wedge at the end of a beam, in order to minimize reflection from the wedge. For a single ABH profile, the authors varied the thickness, length, loss factor,

and Young’s modulus of the damping layer, and then used an impedance matrix method to calculate the reflection coefficient at the end of the wedge. They concluded that the thickest damping layer did not produce the smallest reflection coefficient, and that the entire wedge did not need to be covered by damping material to produce minimum reflection. Unfortunately, they only considered a single ABH taper profile, and thus could not draw conclusions regarding relationships between the taper profile and the optimal amount of damping material. Moreover, their process was to first vary the thickness of the damping layer, then choose the thickness that gave minimal reflection, and use that same thickness for the rest of the parametric study. This means that other combinations of thickness and length – combinations which may be more effective – were not considered.

Unruh, Blech, and Monner published a parametric study of their own in 2015, in which they explicitly modeled the vibration reduction of a two-dimensional ABH in a rectangular plate [23]. Using the finite element method, they calculated the surface-averaged velocity response and radiated sound power from the plate due to a harmonic point force at one corner. They performed two parametric studies: for the first, they considered a single ABH and varied its radius and position on the plate; for the second, they considered two ABHs at fixed positions, and varied each one’s radius. The other taper parameters and the amount of damping material applied to each ABH was the same for all cases.

The authors concluded that the radius of the ABH had a more significant effect on the plate’s response and radiated sound than did the placement of the ABH on the plate. Moreover, they found that the benefits from the addition of a second ABH were significantly less than from the addition of the first ABH, suggesting that there might be diminishing returns from the addition of further ABHs. However, the only positions considered in the study were along the diagonal of the plate, in line with the drive point. As later studies would find [24], the optimal placement of the ABH with respect to a single drive point may be elsewhere, and may have a more significant effect than the radius of the ABH. Furthermore, the modeled damping layer was drastically thinner than Georgiev et al. found to be most effective. If the placement of the ABH had little effect on their results, it is conceivable that energy was not being effectively absorbed, which might have skewed their results. In short, it is difficult to draw conclusions from this study about the effects of ABH radius and position on performance.

In their 2015 paper studying the use of the normalized wavenumber variation (NWV) for ABH design [25], Feurtado and Conlon performed what amounts to a parametric

study, investigating how the NWV and reflection coefficient change for a one-dimensional ABH wedge with various lengths. Using a finite element model of a beam with an ABH wedge, they varied the length of the wedge, with and without damping and with and without an anechoic termination, and determined that the reflection coefficient tracks with NWV. Additionally, they found that reflection is minimized for NWV less than 0.3—a fact that would be used in future optimization studies [26]. Unfortunately, the authors only varied the wedge length, meaning that conclusions cannot be drawn regarding the optimal NWV as relates to other taper parameters like power-law exponent or minimum thickness.

Additionally, in support of the experimental investigation of their 2016 publication [27], Feurtado and Conlon performed a parametric study of the effect of damping thickness and percent coverage on the modal loss factor of a two-dimensional ABH unit cell finite element model in order to identify the optimal damping for experiments into the dynamic behavior of a plate with a 4-by-5 grid of such ABHs. Like authors before them, they confirmed that the optimal damping layer is neither the thickest nor the widest. In particular, Feurtado and Conlon concluded that the thickness of the free damping layer corresponds to six times the minimum thickness—in agreement with the damping theory of Ross, Kerwin, and Ungar [28]. Unfortunately, because the nature of the main investigation prescribed a given ABH geometry, there was no consideration of how the optimal damping might change by changing the ABH design parameters.

In consideration of the potential applications for ABH vibration absorbers in the aerospace industry, Dorn, Blech, and Langer published a parametric study in 2017 of a two-dimensional ABH in a curved honeycomb panel with fiberglass face sheets [29]. The thickness of the honeycomb core itself was tapered according to a power-law profile, while the two face sheets were held at a uniform thickness. The diameter and position of the ABH were varied within the curved panel, each having three possible values, for a total of nine combinations. For each case, 25% of the ABH area included an additional damping layer. Harmonic point forces were applied to each of the four corners of the panel, and a spatially-averaged square velocity response was calculated and compared to a reference panel without an ABH.

The authors concluded that the presence of the ABH in the panel improved the velocity response very weakly when compared to a uniform panel with the equivalent damping material distributed at the corners. As the authors implied, this could very well be due to the low frequency range considered, which was only up to 1 kHz. It may be that the ‘cut-on’ frequency of the ABH is above this frequency and that the ABH

is not providing any benefit. Although the ABH did not seem to improve performance significantly, the authors did find that what improvement it made was more strongly dependent on the position of the ABH than on the size, which is in opposition to the conclusions of Unruh, Blech, and Monner [23].

Wang et al. performed an investigation in 2019 of a one-dimensional ABH at the edge of a circular disk [30]. The authors proposed the use of an ABH as a means to dramatically reduce the reverberant field produced when performing acoustic emission testing—a common problem in the study of frictional mechanisms [31]. To demonstrate the efficacy of this technique, the authors performed a parametric study of the ABH taper length relative to the plate size for three different sensor locations, using transient analysis of a finite element model and then verifying their results experimentally. In order to reduce the complexity of the parametric study, the authors used orthogonal array testing (OAT) to capture the results of all 27 possible combinations in only nine test cases.

It is worth noting, however, that OAT assumes any effect is the result of the interaction between parameters, which is to say that it is assumed no significant effect is produced by the change of a single parameter alone. This, of course, is not true of ABH design, as evidenced by the literature already discussed. Wang et al. nevertheless concluded that the most critical parameter for the proposed technique is the ratio of taper length to plate diameter, which should be relatively large. The authors therefore confirmed what had been established by past research—i.e., that a longer taper is more effective for vibration reduction.

Also in 2019, Hook, Cheer, and Daley presented a parametric study of a one-dimensional ABH wedge at the end of a long beam [32]. Using a finite element model of the beam and ABH, the authors estimated the reflection coefficient from the end of the wedge by analyzing the dynamic response of the model at two different locations along the beam and assuming only forward- and backwards-propagating waves. By varying the length, minimum thickness, and power-law exponent of the taper profile, the authors identified frequency regions in which the reflection coefficient is minimized, as well as indicating that the variation in wedge length affected reflection significantly more than did variation in minimum thickness. Moreover, the authors found that regions of minimum reflection were especially distinct for different values of power-law exponent.

Bearing this behavior in mind, the authors performed modal analysis of just the ABH wedge for various power-law exponent values, and found that the bands of minimum reflection directly corresponded to the modal frequencies of the wedge. This finding

corroborates the trend of higher structural losses associated with ABH modes [12, 33]. Moreover, although the authors spend much of their discussion considering the optimal combination of parameters to minimize reflection in the broadband, the real strength of this study lies in the trends identified, which can help to shed light on the tradeoff between optimal performance in the broadband and optimal performance in the narrowband.

In their 2019 work, Ouisse et al. sought the optimal damping configuration for a given beam with an ABH termination, using a process that is more formally a parametric study [34]. Given the base structure of an aluminum beam with an attached quadratic ABH wedge, the authors modelled the effects of varying the length and thickness of an attached shape memory polymer (SMP). What was unique about this paper is that before varying the length and thickness of the SMP layer, the authors first modelled the effect of controlling the temperature within the SMP. The reason for this is that the SMP material properties are highly temperature-dependent. The loss factor of the material can be tuned by adjusting the temperature, although the Young’s modulus will also be affected. The authors chose a nominal damping design, and then modelled the reflection coefficient for the ABH beam with SMP damping at five temperatures ranging from 30°C to 70°C.

Deciding that the SMP material performed best near 70°C, the authors then performed a parametric study on how the length and thickness of the SMP material would affect the average reflection coefficient of the composite ABH beam. They calculated the average reflection coefficient for two frequency ranges (0–200 Hz and 0–1000 Hz) in order to compare the ‘optimal’ configurations for each range. The authors found that the two measures produced conflicting results, with the two configurations differing by at least 20%. Nevertheless, the authors were successful in showing that the control of material properties afforded by the use of an SMP damping layer shows great potential for the design of effective ABH vibration absorbers.

1.2.2 Review of optimization studies

The first formal optimization study regarding ABH design for vibration reduction was published by Rothe et al. in 2016 [35]. In this optimization study, the authors sought the optimal position of a two-dimensional ABH in a rectangular plate with respect to the surface-averaged velocity response due to a harmonic point force at the corner. The authors found that a single ABH at the optimal position could reduce the response by 51.9% relative to a uniform plate, while the same design at a suboptimal position could worsen the response by 36.4%.

It is worth noting that although the ABH design, damping, and drive point location in this study were similar to those investigated by Unruh, Blech, and Monner [23], the optimal position identified in this study is one very different from any of the positions considered in that parametric study. In particular, the optimal position was away from the diagonal. This is an example of where a formal optimization study can provide a fuller picture than a parametric study, which is inherently limited in scope. Unfortunately, no attempt was made to explain the ABH’s performance as a function of position. Furthermore, this optimization study only considered a single ABH geometry, meaning that no information can be gleaned regarding interactions between the shape parameters and the optimal position of the ABH.

Another formal optimization study was published in 2016 by Ih et al., in which the authors sought the optimal combination of the taper profile and damping for a one-dimensional ABH at the edge of a plate [24]. Using the finite element method, the authors calculated the spatially-averaged square velocity response due to a harmonic point force at the other end of the plate. For the first part of the optimization, the power-law exponent of the ABH taper and the length of the minimum-thickness portion were allowed to vary with no added damping layer. 55 random combinations of the two variables were evaluated, and then kriging (i.e., Gaussian process regression) was used to interpolate the rest of the search space. In this way, the authors visually identified regions of optimality within the variable bounds and concluded that an optimal combination produces a normalized wavenumber variation (NWV) between 0.45 and 0.55. Note that this is in opposition to the findings of Feurtado and Conlon, who concluded that optimal performance is obtained for an ABH with NWV less than 0.3 [25].

In the second part of the optimization, Ih et al. chose the best performing of the 55 random ABH designs, and added a free damping layer to the bottom. The thickness and length of this damping layer were allowed to vary, while maintaining a bending stiffness less than that of the plate. Again, a finite number of random thickness/length combinations were evaluated and kriging was used to interpolate the remaining search space. Ultimately, the authors concluded that damping dimensions that produced between 10 and 20 grams per meter width were optimal. They selected one of the randomly-generated combinations to be used as the ‘optimal’ ABH vibration absorber design, and found that this ABH minimized the objective function by 1.4 dB compared to a heavily damped uniform plate.

Of course, there are limitations to the conclusions that can be drawn from this optimization study. For one, although the authors generated a model of the search space

through kriging, they chose not to use this model for any quantitative evaluation. Instead, they considered only the finite number of designs used to generate the model, despite the fact that their model predicts designs that perform better. Furthermore, their process first addressed the ABH geometry, and then separately addressed the damping, meaning that any interaction between the two was not considered.

One of the most rigorous and comprehensive ABH optimization studies came from Shepherd, Feurtado, and Conlon in 2016 [26]. In this study, the authors applied formal multi-objective optimization using a multi-objective evolutionary algorithm (MOEA) to identify the tradeoff between minimizing reflection from a one-dimensional ABH and satisfying the theoretical condition on NWV. Building off of the findings of Feurtado and Conlon [25], Shepherd, Feurtado, and Conlon set as their two objectives 1) the frequency at which the reflection coefficient is less than 0.4, and 2) the frequency at which the NWV is less than 0.4. The entire geometry of the ABH taper was allowed to vary, including the minimum and maximum thicknesses.

By running the MOEA for 30,000 evaluations, a set of optimal geometries was found that illustrates the tradeoff between the two objectives. The authors confirmed that the longest possible taper is optimal and that the smallest possible maximum and minimum thicknesses is optimal. This means that the tradeoff between the two objectives is almost entirely dependent upon the power-law exponent, a parameter that was not considered in previous investigations [25]. Additionally, the authors found that significant reduction in reflection can occur despite violating the condition on NWV, and that these reductions can be realized at frequencies lower than previous analytical predictions [12]. This optimization study pointed the way towards future studies into the tradeoffs inherent in ABH design for vibration reduction, and demonstrated that MOEAs are a powerful tool for this kind of analysis.

Although only a single ABH profile was considered, in 2019 Ma and Cheng performed a thorough topology optimization of the damping attached to a two-dimensional ABH feature [36]. Utilizing their earlier work on a Rayleigh-Ritz solver with Daubechies wavelet basis functions to calculate the plate dynamics, and the Rayleigh integral to calculate radiated sound power, the authors used the method of solid isotropic material with penalization (SIMP) together with the optimality criteria method. In short, the plate geometry was discretized into 900 square elements and each element was allowed to have applied damping or no applied damping, with a constraint on the total volume of damping material. The optimization process was first carried out for individual frequencies ranging from 100 to 4000 Hz.

The authors found that, compared to a reference damping configuration of a circular patch at the center of the ABH, the optimized damping distribution could reduce the radiated sound power by up to 4 dB at a given frequency. They did note that the mean square velocity was increased at certain frequencies compared to the reference damping configuration. As a result, they performed a similar optimization, using the mean square velocity as the objective, rather than radiated sound power, and found that the latter objective precipitated a distribution that had significantly more damping at the center of the ABH for all but the highest frequencies. This is to be expected, since the bending wavespeed in the ABH will be significantly less than that in the base plate, making the ABH a naturally poor radiator for all but the first mode or ‘piston’ mode. As such, there is not much need to add damping material to the center of the ABH in order to reduce noise, since it radiates poorly anyway. However, the velocity in the ABH will be significantly greater than that in the base plate, making it the most effective region to introduce damping if one wishes to minimize velocity. Ultimately, the most important result of these optimization studies was to underpin the effect that careful design of added damping material can have on ABH performance, and that the optimal design varies with frequency.

In 2020, Ma and Cheng explored ABH profiles of a more general polynomial shape [37]. The profile was defined using Lagrange interpolation of five points to define a 4th-order polynomial. Four of the x -coordinates were fixed (including the origin and the length of the taper) and two of the y -coordinates were fixed (the minimum and maximum thicknesses); the remaining four values were used as design variables, including the thickness at the origin. This 4th-order polynomial was then rotated radially to form a circular ABH embedded in a plate. The objective function chosen was the reciprocal of the sum of the first 100 modal loss factors. Interestingly, an additional constraint placed on the problem was that the total accumulated phase from the origin to the length of the taper is greater than that of some reference ABH profile. The authors claim that this constraint ensures the optimized taper preserves the features of an ABH, although this argument is questionable.

The authors used the Nondominated Sorting Genetic Algorithm II (NSGA-II), a multi-objective evolutionary algorithm (MOEA), to perform the optimization, along with the aforementioned Rayleigh-Ritz solver to calculate the modal loss factors. The optimization was run for a total of 153 generations to produce the optimal profile. Counterintuitively, the final optimal design in the 153rd generation was found to be almost identical to that of the first generation, despite performing the optimization twice

with randomized starting generations. Nearly identical results were produced even when varying the material properties of the ABH. This may be an indication of either an error in the solver, which could be exploited by NSGA-II, or of a defect in the formulation of the problem or its parameters. Assuming the optimal design to be valid, it is an interesting annular shape with a sharp spike in the center, somewhat similar to the ring-shaped ABH designs proposed in 2019 by Deng et al. [38]. In both [37] and [38], the advantage of an annular ABH seems to be that the minimum thickness region is spread out into a ring, rather than only existing at the center.

1.3 Theory

1.3.1 Analytical ABH theory

The dynamic Euler-Bernoulli (thin) beam equation can be written as[†]

$$\rho A \frac{\partial^2 u}{\partial t^2} + \frac{\partial^2}{\partial x^2} \left(D \frac{\partial^2 u}{\partial x^2} \right) = f(x, t), \quad (1.5)$$

where ρ is the density of the beam material, $A(x)$ is the cross-sectional area of the beam, and $D(x)$ is the bending stiffness, equal to $EI(x)$, where E is the material Young's modulus and $I(x)$ is the second moment of area of the beam's cross-section. Assuming the beam's midplane displacement, $u(x, t)$, is time-harmonic and has the form $u(x, t) = U(x)e^{i\omega t}$, then the steady-state equation of motion becomes

$$-\omega^2 \rho A U + \frac{d^2}{dx^2} \left(D \frac{d^2 U}{dx^2} \right) = 0, \quad (1.6)$$

where it has also been assumed that there is no external applied force—that is, free vibration conditions. The WKB approximation is now used in a fashion similar to that used in [39]. This approximation is valid in the case that variations in geometric or material parameters are small on the scale of a wavelength. In the ABH literature, a more explicit formulation of this condition is given as [6, 40]

$$\left| \frac{1}{k^2} \frac{dk}{dx} \right| \ll 1, \quad \forall x, \quad (1.7)$$

[†]As mentioned in Section 1.1, historically ABH theory was developed using thin plate theory. However, the results from using thin beam theory are functionally equivalent.

where $k(x)$ is the local wavenumber given by Equation (1.1). Assuming that the WKB approximation is valid, the procedure starts by defining a small parameter, $\delta = \omega^{-\frac{1}{2}}$, so that the highest-order derivative in Equation (1.6) is considered a perturbation in the ordinary differential equation. It should be noted that the assumption that δ is small also means that the approximation is really only valid for high frequencies. However, Euler-Bernoulli beam theory is only valid for low frequencies, where the wavelength is much larger than the beam's thickness and inertial effects are negligible. This dichotomy means that the ABH theory of Mironov exists in some mid-range frequency—a fact that is not always made clear. With the parameter δ defined, Equation (1.6) becomes

$$-\rho AU + \delta^4 \frac{d^2}{dx^2} \left(D \frac{d^2 U}{dx^2} \right) = 0. \quad (1.8)$$

The WKB approximation now seeks a series solution of the form

$$U(x) = e^{i\phi(x)}, \quad (1.9)$$

$$\phi(x) = \frac{1}{\delta} \sum_{n=0}^{\infty} \delta^n S_n(x), \quad (1.10)$$

where the $S_n(x)$ are (generally complex) functions that are determined recursively. By plugging this solution into Equation (1.8), it becomes

$$\begin{aligned} \rho A = \delta^4 \left(D \left[i\phi^{(4)} - 4\phi''' \phi' - 3(\phi'')^2 - 6i\phi'' (\phi')^2 + (\phi')^4 \right] \right. \\ \left. + 2D' \left[i\phi''' - 3\phi'' \phi' - i(\phi')^3 \right] \right. \\ \left. + D'' \left[i\phi'' - (\phi')^2 \right] \right). \end{aligned} \quad (1.11)$$

Expanding this equation and collecting terms with the same powers of δ , this becomes

$$\rho A = D (S'_0)^4 + \delta \left[4D (S'_0)^3 S'_1 - 6iD S''_0 (S'_0)^2 - 2iD' (S'_0)^3 \right] + O(\delta^2), \quad (1.12)$$

where ‘big O’ notation has been used to denote the remaining powers of δ . The above equation must hold for all values of δ , which produces a system of equations. The zeroth-order WKB approximation is found by solving for the zeroth-order δ terms—that is, by solving

$$(S'_0)^4 = \frac{\rho A}{D} = \frac{k^4}{\omega^2}, \quad (1.13)$$

where the definition of the local wavenumber, $k(x)$, has been used. The solution to this equation gives

$$S_{0,p}(x) = \frac{i^p}{\sqrt{\omega}} \int_a^x k(y) dy, \quad p = 0, 1, 2, 3 \quad (1.14)$$

for some reference point, a . Note that there are actually four possible solutions, depending on the choice of p , since $S_0(x)$ is generally complex. This zeroth-order WKB approximation is thus a linear combination of them:

$$U_0(x) = C_0 e^{\int_a^x k(y) dy} + C_1 e^{i \int_a^x k(y) dy} + C_2 e^{-\int_a^x k(y) dy} + C_3 e^{-i \int_a^x k(y) dy}. \quad (1.15)$$

This clearly shows that the phase of the wave varies over space, according to the accumulated local phase. However, it does not include any variation in wave amplitude, which is expected if conservation of energy is to hold. Extending the solution to the first-order WKB approximation, one must now solve for the first-order δ terms

$$\begin{aligned} S'_{1,p} &= \frac{i}{2} \frac{D'}{D} + \frac{3i}{2} \frac{S''_{0,p}}{S'_{0,p}} \\ &= \frac{i}{2} \frac{D'}{D} + \frac{3i}{2} \frac{k'}{k}. \end{aligned} \quad (1.16)$$

Note that the above differential equation is actually independent of the choice of power, p . As such, there is only one solution, which will be denoted $S_1(x)$. To solve for $S_1(x)$, one must know how the bending stiffness, $D(x)$, and the wavenumber, $k(x)$, vary with x . For a rectangular beam with constant width of b and a varying thickness, $h(x)$, the relations are

$$D(x) = E \frac{h^3(x)b}{12}, \quad (1.17)$$

$$D' = E \frac{h^2 b}{4} h' = 3D \frac{h'}{h}, \quad (1.18)$$

$$k(x) = \left(\frac{12\omega^2 \rho}{E h^2(x)} \right)^{\frac{1}{4}} = \left(\frac{12\omega^2 \rho}{E} \right)^{\frac{1}{4}} h^{-\frac{1}{2}}(x), \quad (1.19)$$

and

$$k' = -\frac{1}{2} \left(\frac{12\omega^2 \rho}{E} \right)^{\frac{1}{4}} h^{-\frac{3}{2}} h' = -\frac{1}{2} k \frac{h'}{h}. \quad (1.20)$$

Plugging these into Equation (1.16) gives

$$S'_1 = \frac{3i}{2} \frac{h'}{h} - \frac{3i}{4} \frac{h'}{h} = \frac{3i}{4} \frac{h'}{h}. \quad (1.21)$$

Therefore,

$$S_1(x) = \frac{3i}{4} \ln(h(x)) + c. \quad (1.22)$$

Combining this with the zeroth-order WKB approximation, the first-order WKB approximation is thus

$$U_1(x) = \frac{C_0}{h(x)^{\frac{3}{4}}} e^{\int_a^x k(y) dy} + \frac{C_1}{h(x)^{\frac{3}{4}}} e^{i \int_a^x k(y) dy} + \frac{C_2}{h(x)^{\frac{3}{4}}} e^{-\int_a^x k(y) dy} + \frac{C_3}{h(x)^{\frac{3}{4}}} e^{-i \int_a^x k(y) dy}. \quad (1.23)$$

This shows that the displacement amplitude also increases with decreasing thickness, as expected. This increase in local amplitude and decrease in local bending wavelength is the key to the ABH effect and the efficacy of imperfect ABH vibration absorbers.

1.3.2 Reflection from an ABH termination

Reconsider now the zeroth-order WKB approximation of Equation (1.15). For a beam that is free at $x = 0$ and extends to infinity in the positive x -direction, it is clear that $C_0 = 0$, since for it to be otherwise would imply that $\lim_{x \rightarrow \infty} U_0(x) = \infty$. For all $x > a$, it is clear that the C_1 term represents an incoming (left-going) wave. As such, it is advantageous to normalize everything by C_1 to get

$$\frac{U_0(x)}{C_1} = e^{i\psi(x)} + R_n e^{-\psi(x)} + R e^{-i\psi(x)}, \quad (1.24)$$

where $\psi(x) = \int_a^x k(y) dy$ simply for ease of notation. R is the reflection amplitude coefficient and R_n is the amplitude of the nearfield component of the reflected wave, which disappears as x goes to infinity.

Two additional equations are required to solve for these two remaining coefficients. These two equations come from the boundary conditions at the free end, which are $U_0''(0) = U_0'''(0) = 0$. After normalizing by C_1 , the derivatives of $U_0(x)$ are

$$\begin{aligned} \frac{U_0''(x)}{C_1} &= i e^{i\psi(x)} (k'(x) + i k^2(x)) \\ &\quad - R_n e^{-\psi(x)} (k'(x) - k^2(x)) \\ &\quad - i R e^{-i\psi(x)} (k'(x) - i k^2(x)) \end{aligned} \quad (1.25)$$

and

$$\begin{aligned}\frac{U_0'''(x)}{C_1} = & \quad \text{ie}^{i\psi(x)} \left(k''(x) + 3ik'(x)k(x) - k^3(x) \right) \\ & - R_n e^{-\psi(x)} \left(k''(x) - 3k'(x)k(x) + k^3(x) \right) \\ & - iR e^{-i\psi(x)} \left(k''(x) - 3ik'(x)k(x) - k^3(x) \right).\end{aligned}\quad (1.26)$$

While it *is* possible to determine R from these equations, additional approximations can be made to simplify the solution. Firstly, if it is assumed that Equation (1.7) holds, then $|k'| \ll |k^2|$. This means that Equation (1.25) can be rewritten as

$$\frac{U_0''(x)}{C_1} \approx -e^{i\psi(x)} k^2(x) + R_n e^{-\psi(x)} k^2(x) - R e^{-i\psi(x)} k^2(x).\quad (1.27)$$

Secondly, Equation (1.7) also implies that $|k'k| \ll |k^3|$, so that Equation (1.26) can be rewritten as

$$\begin{aligned}\frac{U_0'''(x)}{C_1} \approx & \quad \text{ie}^{i\psi(x)} \left(k''(x) - k^3(x) \right) \\ & - R_n e^{-\psi(x)} \left(k''(x) + k^3(x) \right) \\ & - iR e^{-i\psi(x)} \left(k''(x) - k^3(x) \right).\end{aligned}\quad (1.28)$$

Setting Equations (1.27) and (1.28) equal to zero at $x = 0$, one can rearrange and substitute to solve for R to get

$$R = \frac{(-1 + i)k''(0) + (-1 - i)k^3(0)}{(1 + i)k''(0) + (1 - i)k^3(0)} e^{2i\psi(0)}.\quad (1.29)$$

This expression can be simplified further by assuming that $|k''| \ll |k^3|$, along the same lines as Equation (1.7). In such a case, the expression for R simplifies to

$$R \approx \frac{(-1 - i)k^3(0)}{(1 - i)k^3(0)} e^{2i\psi(0)} = -ie^{2i\psi(0)},\quad (1.30)$$

where the magnitude and phase, $|R|$ and $\angle R$, are

$$|R| = e^{-2\text{Im}\{\psi(0)\}}\quad (1.31)$$

and

$$\angle R = \frac{3\pi}{2} + 2\text{Re}\{\psi(0)\}.\quad (1.32)$$

At this point, it is worthwhile to point out two things. First, these results are consistent with, and form a complementary view of, the acoustic black hole phenomenon described in Section 1.1 and Equation (1.4). Namely, as $\psi(0)$ approaches infinity, the phase of R approaches infinity and theoretically perfect absorption occurs as the wave becomes ‘trapped’ in the termination. Second, for finite $\psi(0)$, the magnitude of R depends on the damping in the beam, which has thus far not been explicitly considered. Consider the following alteration of Equation (1.1)

$$\tilde{k}^4 = \frac{\omega^2 \rho A}{\widetilde{D}} = \frac{\omega^2 \rho A}{D(1+i\eta)}, \quad (1.33)$$

where the tilde denotes a complex quantity, and η is some material loss factor[‡]. In most practical cases, $\eta \ll 1$, so that the following approximations can be made:

$$(1+i\eta)^{-\frac{1}{p}} = r^{-\frac{1}{p}} e^{-i\frac{\theta}{p}}, \quad (1.34)$$

$$r = \sqrt{1+\eta^2} \approx 1, \quad (1.35)$$

$$\frac{\theta}{p} = \frac{1}{p} \text{atan}(\eta) \approx \frac{1}{p} \text{asin}(\eta) \approx \frac{\eta}{p}, \quad (1.36)$$

and

$$(1+i\eta)^{-\frac{1}{p}} = r^{-\frac{1}{p}} \cos\left(\frac{\theta}{p}\right) - i r^{-\frac{1}{p}} \sin\left(\frac{\theta}{p}\right) \approx \cos\left(\frac{\eta}{p}\right) - i \sin\left(\frac{\eta}{p}\right) \approx 1 - i\frac{\eta}{p}, \quad (1.37)$$

where small argument approximations for tangent, sine, and cosine have been used. Applying these approximations to Equation (1.33) gives

$$\tilde{k} \approx \sqrt[4]{\frac{\omega^2 \rho A}{D}} \left(1 - i\frac{\eta}{4}\right) = k \left(1 - i\frac{\eta}{4}\right). \quad (1.38)$$

Thus,

$$-2\text{Im}\{\psi(x)\} \approx \frac{\eta}{2} \int_a^x k(y) \, dy. \quad (1.39)$$

In much of the literature, the reference point, $x = a$, is chosen to be the point at which the ABH profile reaches some maximum thickness, $x = x_1$, such that $h(x_1) = h_1$. With

[‡]The plus sign before $i\eta$ is a result of the $e^{i\omega t}$ time convention used. In the case of an $e^{-i\omega t}$ time convention, the sign is flipped.

this convention, the reflection magnitude becomes

$$|R| = e^{\frac{\eta}{2} \int_{x_1}^0 k(y) dy} = e^{-\frac{\eta}{2} \int_0^{x_1} k(y) dy} = e^{-\eta\omega\tau}, \quad (1.40)$$

where Equation (1.4) has been used to write the reflection magnitude in a more compact and intuitive form. Under this convention, the above result is equivalent to that found in, e.g., references [6] and [39].

1.3.3 Multi-objective evolutionary optimization and the Borg algorithm

In the process of optimization, the goal is to find an optimal solution to a given problem within a feasible range of independent variables and perhaps within some constraints on the variables or solution. Provided an analytic objective function, $f(x)$, is available to relate the independent variables to the fitness of the solution, a common method is to use the mathematical properties of the gradient, ∇f , to find a solution that minimizes the function. In particular, the direction of $-\nabla f$ gives the direction of greatest decrease, while the norm of ∇f (which in practice is often approximated by finite differences) gives a measure of how quickly the function decreases in that direction. However, many real-world optimization problems arise where this procedure is either impractical or impossible. For example, gradient-based methods may break down for functions whose gradient is not continuous everywhere, and in combinatorial optimization the gradient is generally not well-defined [41]. It is common for multiple optimal solutions to exist, and gradient-based methods may converge to one solution while failing to locate others. Even in the case where a global optimum exists, gradient-based methods may converge to a non-optimal local minimum. This is where stochastic, heuristic-based optimization algorithms are more robust, as they make no assumptions about the properties of the objective space. In this sense, they act as so-called ‘black box’ optimizers. A particularly useful class of these algorithms, inspired by the processes of natural selection, is evolutionary algorithms (EAs) [41, 42]. By using numerical analogs to mutation, recombination, and selection, an EA takes an initial ‘population’ of solutions and produces ‘offspring’ that are more ‘fit’ than previous generations in terms of the objective function. In the limit that the number of generations goes to infinity, the EA should produce offspring that perform optimally with respect to the objectives.

EAs are more robust than gradient-based optimization algorithms insofar as they do not require the objective function to be smooth and can perform well on ‘noisy’ objective




Figure 1.2. Example Pareto front (solid curve) for a two-objective optimization problem [1]. The two objectives are f_1 and f_2 . The dashed curve represents the boundary of the realizable objective space, \mathbf{F} .

spaces—a situation in which gradient-based algorithms can perform particularly poorly. Although the stochastic nature of EAs make it difficult to analyze their convergence properties, there are a number of test functions on which they reliably outperform gradient-based methods [43]. Moreover, the population-based architecture of EAs means that they are inherently adept at identifying sets of optimal solutions when the design space is multimodal or there does not exist a unique global optimum. The EA selected for this dissertation is a powerful multi-objective evolutionary algorithm (MOEA) named ‘Borg’ and developed by Hadka and Reed [44]. Borg is part of a class of MOEAs that find a set of optimal solutions by approximating the so-called ‘Pareto front’. An example Pareto front for two objectives is shown in Figure 1.2.

First formalized by Vilfredo Pareto in 1906, Pareto optimality captures the inherent trade-offs that can exist when multiple objectives compete [45]. Given a set of M design variables, $\mathbf{x} = [x_1, x_2, \dots, x_M]^T$, one can conceptualize all possible combinations of N objectives, $\mathbf{J} = [J_1(\mathbf{x}), J_2(\mathbf{x}), \dots, J_N(\mathbf{x})]^T$, as an N -dimensional space. One design, \mathbf{x}_A , is said to ‘Pareto dominate’ another, \mathbf{x}_B , when $J_i(\mathbf{x}_A) \leq J_i(\mathbf{x}_B)$ for $i = 1, 2, 3, \dots, N$ and there is at least one objective, $J_j(\mathbf{x})$, such that $J_j(\mathbf{x}_A) < J_j(\mathbf{x}_B)$. In the case that neither design Pareto dominates the other, the two are said to be ‘non-dominated’ [1]. The Pareto

optimal set, then, is the set of all non-dominated designs within the M -dimensional design space. Borg and other Pareto-based MOEAs employ evolution-inspired heuristics to approximate the Pareto optimal set over a sufficiently large number of iterations.

Among other things, Borg is set apart by three features: its discretization of the objective space, its use of automatic restarts, and its auto-adaptive use of multiple recombination methods. Borg implements a version of the ϵ -dominance archive first proposed by Rudolph [46] and improved upon by Laumanns et al. [47]. Conceptually, the objective space is divided into discrete N -dimensional hypercubes or ‘boxes’ of side lengths $\epsilon_1, \epsilon_2, \dots, \epsilon_N$, and it is these ϵ -boxes, rather than the continuous objective function values, that are used to determine Pareto dominance. It can be shown that the ϵ -dominance scheme guarantees both convergence to the Pareto optimal set and a diverse set of designs [47]. By breaking up the objective space into discrete units, ϵ -dominance also provides an intuitive user-defined control of the resolution. This can be important in practical applications where improvements smaller than a certain increment may not be considered significantly better.

The ϵ -box implementation also facilitates a natural measure of search progress. Each time one design dominates another, the dominating design is archived. The archive size is used in Borg to determine the population size of candidate designs as well as the tournament size of competing designs. Similar to gradient-based methods, MOEAs can sometimes waste effort converging to local optima when the algorithm consistently finds small improvements near a local optimum, causing search to stagnate. Borg counteracts search stagnation by using ϵ as the minimum threshold for improvement and requiring the algorithm to periodically produce at least one solution whose improvement exceeds this threshold. If, after a number of iterations, the new archive does not include at least one solution that is outside the ϵ -boxes of the previous archive, then the search has stagnated and is restarted. During restart, the population and tournament sizes are adapted, the population is emptied and repopulated with the archive, and the remaining slots in the population are filled by random mutations of the archive designs.

Like other MOEAs, Borg searches the design space using recombination, the process of selecting traits from current designs and applying random mutation and crossover to produce new designs. Because it is difficult to know *a priori* which recombination operator will be most effective for a given problem, Borg uses several methods that are adaptively selected based on their ability to produce optimal solutions, an idea introduced by Vrugt and Robinson [48] and improved upon by Hadka and Reed [44]. The idea is to generate a feedback loop in which recombination operators that take a parent population

and produce more successful offspring designs are rewarded by increasing the number of offspring produced by those operators in future recombinations. As a result, Borg is not a single algorithm, but instead an algorithm template, whose operators are adaptively selected based on the problem. This also means that Borg can adapt its algorithm to better evaluate different regions of the design space for the same problem. The particular recombination operators that Borg chooses from are simulated binary crossover [49], differential evolution [50], parent-centric crossover [51], unimodal normal distributed crossover [52], simplex crossover [53], uniform mutation, and polynomial mutation [49].

Borg has been shown to be a robust MOEA that performs better than other MOEAs on a number of standard test functions [44], and has been successfully applied to optimization problems in several fields [54–58].

1.3.4 Optimization framework

A flowchart outlining the steps to formulate an optimization problem and develop an appropriate optimization framework is shown in Figure 1.3. The generalized process for a given optimization study is as follows: First, objectives are designed that distill physical behavior of interest or correspond to practical outcomes of the particular ABH implementation. Examples include overall mass, spatially-averaged square velocity response, and total dissipated power. Then, design variables and constraints relevant to those objectives are identified. The design variables are next related to the objectives through an objective function, which in some cases will not be an analytic expression but will instead involve the modelling of physical behavior using techniques such as the transfer matrix method (TMM) and the finite element method (FEM). To this end, the design variables are passed to a meshing algorithm that creates a mesh for the particular ABH implementation. This mesh is analyzed using the appropriate solution method (TMM or FEM), and the results are passed to a post-processor. The post-processor finally extracts the results and performs whatever necessary operations to calculate the objectives. This process of taking in the design variables and outputting the objectives as real scalars constitutes the objective function. As far as Borg is concerned, the objective function is a black box. This objective function is linked with Borg, which adaptively explores the solution space based on the output of the objective function, until a prescribed number of function evaluations has occurred or another termination condition is met. Borg then prints out the set of optimal solutions that approximate the Pareto front, thus concluding the optimization process. Because the Pareto front inherently captures the tradeoffs between objectives, analysis of trends in the solution

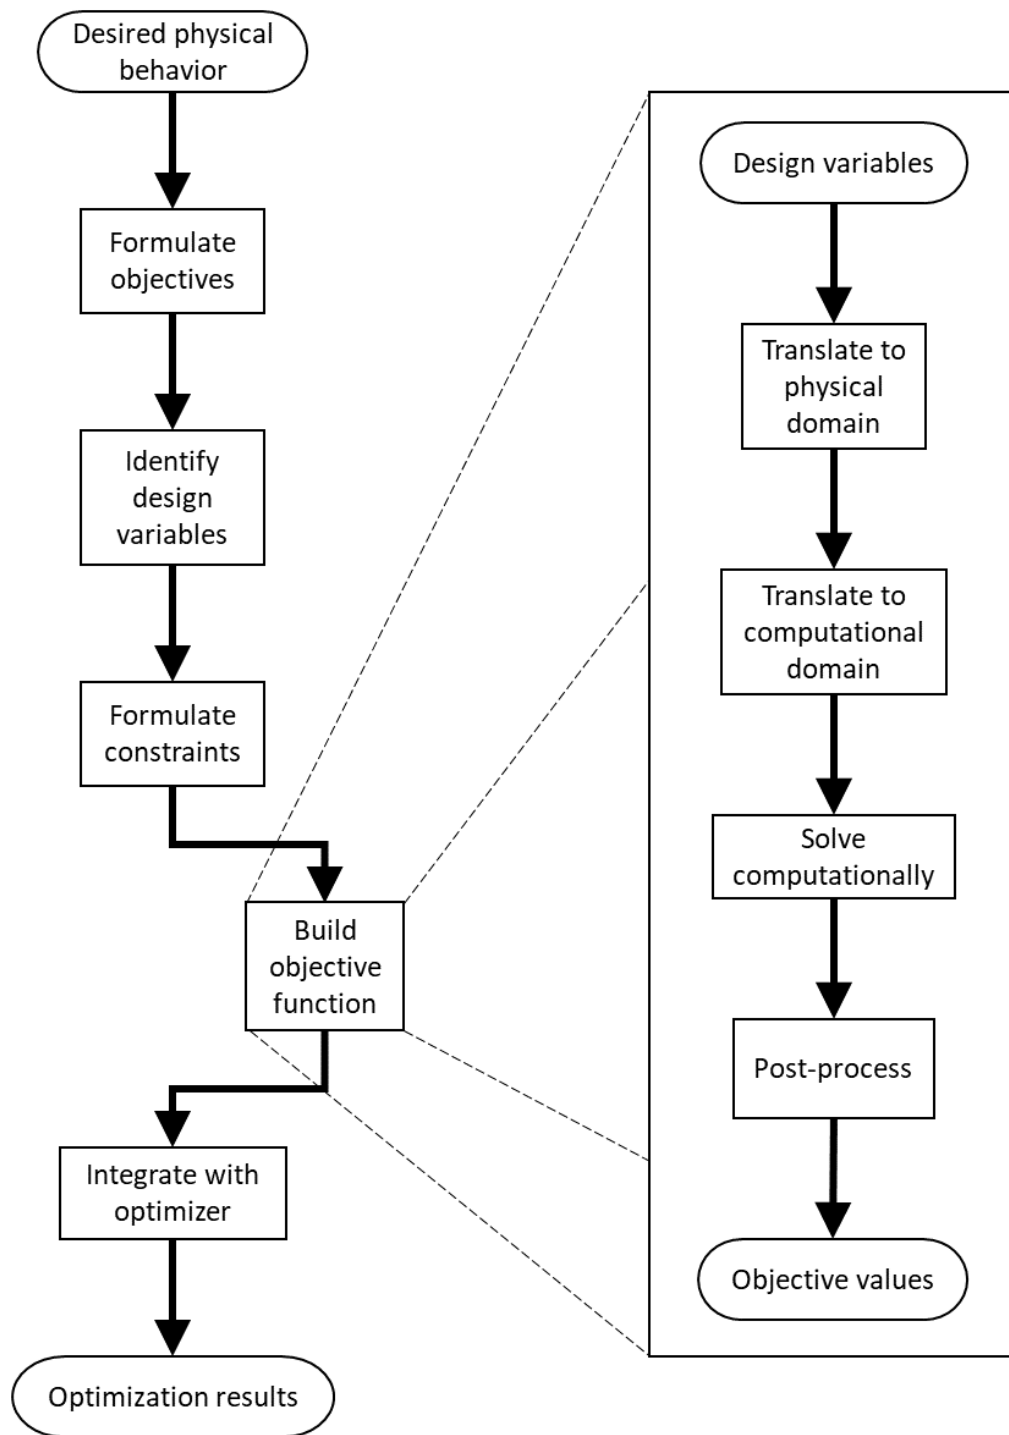


Figure 1.3. Flowchart representation of the process for formulating an optimization problem. Note that this process is not prescriptive and is only a formalization of the author’s own process.

set can illuminate how considerations for ABH design interact to produce tradeoffs in performance.

1.4 Outline of the dissertation

The objectives of this dissertation are threefold. Firstly, an optimization framework must be established for use in the structural optimization studies related to ABH design. This includes developing the methods and computer code necessary to perform the optimization, but also requires the selection of an appropriate optimizer. Secondly, the optimization framework will be used to perform several structural optimization studies related to practical ABH implementations. Of particular interest will be the evaluation of ABH vibration absorbers as part of a larger, finite structure. In contrast to evaluating the ABH alone, the dynamic response of the composite structure will be considered in evaluating the performance of the ABH vibration absorber. Some factors that will be considered include the thickness profile of the ABH, the particular way a given thickness profile is implemented, and the damping configuration included with the ABH vibration absorber. General design guidance will then be sought from the results of these structural optimization studies. Finally, throughout the work of this dissertation it has become clear that the optimal ABH design can be highly dependent on factors such as frequency range and the scale of the problem. Therefore, an additional objective of this dissertation is to develop a framework for ABH design that is independent of these factors, to the greatest extent possible.

With these objectives in mind, the structure of the dissertation is as follows:

- In Chapter 2 the Borg algorithm is tested against a standard MATLAB function for constrained minimization, `fmincon`. Both algorithms are tested on an example problem of optimizing the shape of a beam under harmonic excitation, from which parallels and comparisons are drawn.
- In Chapter 3 the Borg algorithm is applied to the problem of optimizing the shape of an ABH termination at the end of a cantilever beam. To model the dynamics of the system, a novel formulation of the transfer matrix method (TMM) is developed.
- In Chapter 4 the optimal design of three ‘styles’ of one-dimensional ABH are compared. For each style, the objective is to find the design that simultaneously minimizes vibration response and overall mass. Because these two objectives are at odds, a Pareto optimal set is found.

- In Chapter 5 the problem of optimizing a circular two-dimensional ABH in a plate is addressed. Vibration reduction performance is evaluated through a new energy-based measure. Along with deriving this new performance measure, a set of non-dimensional design variables is derived for use in the optimization. To model the dynamics of the system, a solver is developed using the finite element method (FEM) together with the deal.II library [59].
- In Chapter 6 the major findings of the dissertation are summarized and suggestions are made for possible future work.

Chapter 2 |

A comparison of optimization algorithms: Gradient-based versus evolutionary*

2.1 Introduction

In order to accomplish the goals of this dissertation, as they relate to acoustic black hole (ABH) optimization, it is of utmost importance to select a quality optimization algorithm. While there are many algorithms available, they can all be broadly described as one of two classifications: deterministic, gradient-based algorithms and stochastic, heuristic-based algorithms. Evolutionary algorithms (EAs) and the Borg algorithm in particular (detailed in Chapter 1) fall into the latter category. Additionally, the process of ABH optimization will require a robust structural optimization framework for evaluating ABH designs based on some objective(s).

The objectives of the work presented in this chapter are thus twofold. The first is to build a simple optimization framework, including the implementation of a structural-acoustic objective function, the evaluation of design variables, and the incorporation of an optimization algorithm. The second is to demonstrate the suitability of Borg for structural acoustic optimization by evaluating its performance on a test problem from the literature. To this end, Borg's performance is compared against that of a standard gradient-based algorithm to establish the advantages (and disadvantages) of a heuristic-based algorithm like Borg.

*A portion of the work presented in this chapter was also presented at the 179th Meeting of the Acoustical Society of America [60].




Figure 2.1. Graphical depiction of the current optimization problem. The beam geometry is represented by 100 beam segments, where each segment has constant thickness. The left-hand boundary condition is an imposed shear, while the right-hand boundary condition is zero displacement and zero rotation, i.e., fixed. The microphone region, (x_a, x_b) , has constant thickness h_{mic} throughout.

2.2 Problem description

The chosen test problem is taken with minimal modification from Berggren et al. [2], which was presented at the 10th World Congress on Structural and Multidisciplinary Optimization. Nominally, the problem is to minimize the coupling of vibration between a speakerphone’s speaker and its microphone through the speakerphone’s casing. For modeling purposes, the system is reduced to one dimension, with the casing modeled as a thin beam fixed at one end and the speaker modeled as a time-harmonic shear force acting on the other end. The design variable is the thickness distribution of the beam, $h(x)$, and the objective is to minimize the time-harmonic displacement amplitude[†], $U(x) = |\tilde{u}(x, t)|$, within the region representing the microphone, $0 < x_a \leq x \leq x_b < L_x$, where L_x is the length of the beam. To reduce the dimension of the design space, the beam is subdivided into 100 constant-thickness segments so that $h(x)$ is piecewise constant and can be represented by the vector $\mathbf{h} = [h_1 \ h_2 \ \cdots \ h_{100}]^T$. Additionally, the microphone region is enforced to have a constant thickness, h_{mic} , so that in practice the number of design variables, N_{vars} , is some number less than 100 and $h(x)$ can instead be represented by the vector $\mathbf{h} = [h_1 \ h_2 \ \cdots \ h_j \ h_{\text{mic}} \ \cdots \ h_{\text{mic}} \ h_{j+1} \ \cdots \ h_{N_{\text{vars}}}]^T$. A graphical depiction of the beam is shown in Figure 2.1.

Furthermore, the design space is restricted by constraints on the beam’s overall mass and its static compliance under a uniform load with both ends pinned. If m is the mass of

[†]Note that a tilde above a variable indicates that it is (in general) a complex quantity.

a given design and C is its static compliance, then the constraints take the form $m \leq \gamma_m$ and $C \leq \gamma_C$, respectively. With the objective and constraints so described, the problem can be formulated as follows:

$$\begin{aligned}
\min_{\mathbf{h} \in \mathcal{H}} J(\mathbf{h}) &= \sum_{\omega \in \mathcal{W}} J_\omega(\omega, \mathbf{h}) = \sum_{\omega \in \mathcal{W}} \int_{x_a}^{x_b} |U(x, \omega, \mathbf{h})|^2 dx \\
&\text{subject to} \\
\mathcal{H} &= \left\{ \mathbf{h} \in \mathbb{R}^{N_{\text{vars}}} : 0 < h^- \leq h_i \leq h^+, h_{\text{mic}} \text{ const.} \right\} \\
-\omega^2 \rho L_y h(x) U + \frac{\partial^2}{\partial x^2} \left(E \frac{L_y h^3(x)}{12} U_{,xx} \right) &= f(x, \omega), \quad \forall x \in (0, L_x) \\
\left. \begin{aligned} f(x) &= 0 \\ U_{,xx}(0) &= 0 \\ h^3(0) U_{,xxx}(0) &= F' \\ U(L_x) = U_{,x}(L_x) &= 0 \end{aligned} \right\} &\forall \omega \in \mathcal{W} \\
\left. \begin{aligned} f(x) &= 1 \\ U(0) = U_{,xx}(0) &= 0 \\ U(L_x) = U_{,xx}(L_x) &= 0 \end{aligned} \right\} &\text{for } \omega = 0 \\
\int_0^{L_x} \rho L_y h(x) dx &\leq \gamma_m \\
\frac{1}{L_x} \int_0^{L_x} |U(x, 0, \mathbf{h})| dx &\leq \gamma_C
\end{aligned}$$

In words, the above states that the goal of the problem is to find the vector of design variables, \mathbf{h} , that minimizes the objective function, $J(\mathbf{h})$. \mathbf{h} is restricted to be in the set of valid designs, \mathcal{H} , such that the microphone region has a constant thickness and the remaining portions of the beam have a thickness between some lower and upper limits. $J(\mathbf{h})$ is itself a summation of the squared displacement magnitude, $|U(x)|^2$, in the microphone region, (x_a, x_b) , across a set of frequencies, \mathcal{W} . Note the bin width is the same for all of the discrete frequencies and so it is not necessary to include it in the definition of $J(\mathbf{h})$. The dynamic displacement along the beam, $U(x)$, is determined by the time-harmonic solution of the dynamic Euler-Bernoulli beam equation, the solution of which depends on the particular evaluation frequency, ω , the thickness profile of the beam, $h(x)$, the beam's width, L_y , and the beam's material properties, namely its density, ρ , and its Young's modulus, E . Depending on the value of ω , there are different boundary conditions on the zeroth, first, second, and third derivatives of the displacement, $U(x)$,

$U_{,x}(x)$, $U_{,xx}(x)$, and $U_{,xxx}(x)$, respectively. In this sense, the dynamic equation acts as a PDE constraint on the optimization problem, with its boundary conditions depending on whether the evaluation frequency is in the set \mathcal{W} or is equal to zero—i.e., static beam deflection. The latter case applies only when evaluating the final two constraints, which as mentioned above are constraints on the beam’s total mass and its static compliance under a uniform load.

2.3 Theory

2.3.1 Transfer matrix method

Berggren et al. solve the dynamic Euler-Bernoulli beam equation using the finite element method (FEM) with cubic Hermite polynomials, as is standard for beam problems [61]. The use of cubic Hermite polynomials means the solution is continuous in both displacement and in rotation. Rather than use the FEM to solve the dynamic beam equation, the solution function here uses the transfer matrix method (TMM), which is described in detail in the next section. The advantage of the TMM is that the resulting solution is based on harmonic functions rather than polynomials, and is continuous in displacement, rotation, torque, and shear. In this sense, it is a ‘strong’ solution to the dynamic beam equation, as opposed to the ‘weak’ solution employed by the FEM.

To further explain the TMM, note that the dynamic equation for an Euler-Bernoulli beam can be written as

$$\rho A \frac{\partial^2 \tilde{u}}{\partial t^2} + \frac{\partial^2}{\partial x^2} \left(EI \frac{\partial^2 \tilde{u}}{\partial x^2} \right) = f(x, t), \quad (2.1)$$

where $f(x, t)$ is the external force per unit length, ρ is the density of the beam material, E is its Young’s modulus, A is the beam’s cross-sectional area, and I is the second moment of inertia about the beam’s neutral axis. In the case that E and I are independent of position, $f(x, t) = 0$, and assuming a steady-state, time-harmonic solution of the form $\tilde{u}(x, t) = U(x)e^{i\omega t}$, one form of the solution to Equation (2.1) is

$$U(x) = U_0 c_1(x) + \frac{\theta_0}{k} c_2(x) - \frac{M_0}{EI k^2} c_3(x) - \frac{Q_0}{EI k^3} c_4(x) \quad (2.2)$$

with $c_1(x)$, $c_2(x)$, $c_3(x)$, and $c_4(x)$ defined as

$$c_1(x) = \frac{1}{2} \left(\cosh(kx) + \cos(kx) \right), \quad (2.3)$$

$$c_2(x) = \frac{1}{2} \left(\sinh(kx) + \sin(kx) \right), \quad (2.4)$$

$$c_3(x) = \frac{1}{2} \left(\cosh(kx) - \cos(kx) \right), \quad (2.5)$$

and

$$c_4(x) = \frac{1}{2} \left(\sinh(kx) - \sin(kx) \right), \quad (2.6)$$

respectively. U_0 is the vertical displacement of the beam's neutral axis at $x = 0$, θ_0 is its rotation at $x = 0$, M_0 is the torque about the neutral axis at $x = 0$, Q_0 is the shear force at $x = 0$, and $k = \sqrt[4]{\frac{\omega^2 \rho A}{EI}}$ is the wavenumber. For a beam of length L , it is possible to relate the state variables at $x = 0$, $\mathbf{u}_1 = [U_0 \ \theta_0 \ M_0 \ Q_0]^\top$, to the state variables at $x = L$, $\mathbf{u}_2 = [U_L \ \theta_L \ M_L \ Q_L]^\top$, through the relation $\mathbf{u}_2 = \mathbf{Z}\mathbf{u}_1$, where

$$\mathbf{Z} = \begin{bmatrix} c_1(L) & \frac{c_2(L)}{k} & -\frac{c_3(L)}{EI k^2} & -\frac{c_4(L)}{EI k^3} \\ k c_4(L) & c_1(L) & -\frac{c_2(L)}{EI k} & -\frac{c_3(L)}{EI k^2} \\ -EI k^2 c_3(L) & -EI k c_4(L) & c_1(L) & \frac{c_2(L)}{k} \\ -EI k^3 c_2(L) & -EI k^2 c_3(L) & k c_4(L) & c_1(L) \end{bmatrix}. \quad (2.7)$$

This transfer matrix method can also be used to relate the state variables at any intermediate point along the beam, so long as each segment is uniform and homogeneous. In this way, the dynamics of complex beam geometries can be analyzed by partitioning them into segments that are approximately uniform and homogeneous. For such a beam partitioned into N segments and $N + 1$ nodes, the two ends of the beam can be related by

$$\mathbf{u}_{N+1} = \left(\prod_{i=1}^N \mathbf{Z}_i \right) \mathbf{u}_1 = \mathbf{T}\mathbf{u}_1. \quad (2.8)$$

The Π notation is used in Equation (2.8) and subsequently to denote sequential *left* multiplication of matrices. Given a set of known boundary conditions, the remaining unknown state variables in \mathbf{u}_1 and \mathbf{u}_{N+1} can be solved by explicitly calculating \mathbf{T} and carrying out the appropriate algebra. For example, if $\mathbf{u}_1 = [U_1 \ \theta_1 \ 0 \ F]^\top$ and $\mathbf{u}_{N+1} = [0 \ 0 \ M_{N+1} \ Q_{N+1}]^\top$ (i.e., a cantilever beam), then the following system of

equations results:

$$A_{11}U_1 + A_{12}\theta_1 = -A_{14}F, \quad (2.9)$$

$$A_{21}U_1 + A_{22}\theta_1 = -A_{24}F. \quad (2.10)$$

Because \mathbf{T} and F are known, U_1 and θ_1 can be solved for and \mathbf{u}_1 is therefore completely known. State variables at intermediate points can then be determined from \mathbf{u}_1 using sequential application of transfer matrices as $\mathbf{u}_{i+1} = \mathbf{Z}_i\mathbf{u}_i$.

If the beam is excited by a harmonic point force somewhere on the interior of the beam, then the transfer matrix formulation of Equation (2.8) is modified as follows: suppose a point force with magnitude Q is applied at node $1 < j < N + 1$. This would correspond to $f(x, t) = Qe^{i\omega t}\delta(x - x_j)$ in Equation (2.1). The state variables at this node would then be calculated as $\mathbf{u}_j = \mathbf{Z}_{j-1}\mathbf{u}_{j-1} + \mathbf{d}$, where $\mathbf{d} = [0 \ 0 \ 0 \ Q]^T$ is an external perturbation of the state variables in \mathbf{u} . The two ends of the beam are now related by

$$\mathbf{u}_{N+1} = \left(\prod_{i=1}^N \mathbf{Z}_i \right) \mathbf{u}_1 + \left(\prod_{i=j}^N \mathbf{Z}_i \right) \mathbf{d} = \mathbf{T}\mathbf{u}_1 + \mathbf{V}\mathbf{d} \quad (2.11)$$

and the unknown state variables can be solved for just as before. Note that damping can be easily incorporated by using a complex E in the definition of \mathbf{Z} , thus resulting in a complex definition of \mathbf{u} .

The situation is complicated somewhat for the case of static deflection. In such a case, $k = 0$ and so Equation (2.2) cannot be used as the general solution to Equation (2.1). Consider instead the static equation for an Euler-Bernoulli beam:

$$\frac{\partial^2}{\partial x^2} \left(EI \frac{\partial^2 u}{\partial x^2} \right) = f(x). \quad (2.12)$$

If it is still assumed that E and I are independent of position, then the solution for $f(x) = F = \text{const.}$ is

$$u(x) = U_0 + \theta_0 x - \frac{M_0}{2EI}x^2 - \frac{Q_0}{6EI}x^3 + \frac{F}{24EI}x^4. \quad (2.13)$$

One can still use the relation $\mathbf{u}_{i+1} = \mathbf{Z}_i \mathbf{u}_i + \mathbf{d}_i$, except now

$$\mathbf{Z}_i = \begin{bmatrix} 1 & L_i & -\frac{L_i^2}{2EI} & -\frac{L_i^3}{6EI} \\ 0 & 1 & -\frac{L_i}{EI} & -\frac{L_i^2}{2EI} \\ 0 & 0 & 1 & L_i \\ 0 & 0 & 0 & 1 \end{bmatrix} \quad (2.14)$$

and $\mathbf{d}_i = \left[F \frac{L_i^4}{24EI} \quad F \frac{L_i^3}{6EI} \quad -F \frac{L_i^2}{2} \quad -FL_i \right]^T$. In this form, the transfer matrix method can be used just as before.

2.3.2 Smoothing filters

Before describing the details of the optimization procedure, it is worth noting that Berggren et al. found it advantageous to use a spatial smoothing filter for two of the three sets of frequencies [2]. The smoothing filter is applied to the raw thickness profile such that the thickness profile, $h(x)$, is changed to $(h * K)(x)$. This has the effect of taking potentially noisy thickness profiles with many sharp jumps and generating smoother profiles that vary more gradually over space. The filter function, as given by Bendsøe and Sigmund, is [62]

$$(h * K)(x) = \frac{1}{\langle K \rangle} \int_0^{L_x} h(y) K(x - y) dy, \quad (2.15)$$

where

$$\langle K \rangle = \int_{\mathbb{R}} K(y) dy \quad (2.16)$$

and

$$K(x) = \begin{cases} 1 - \frac{|x|}{r} & \text{if } |x| \leq r \\ 0 & \text{otherwise.} \end{cases} \quad (2.17)$$

In other words, $(h * K)(x)$ is the convolution of $h(x)$ with a triangular window of width $2r$. The discrete analog to the above convolution is

$$(h * K)[n] = \frac{1}{\langle K \rangle} \sum_{m=0}^{N-1} h[m] K[n - m] \Delta x, \quad (2.18)$$

where

$$\langle K \rangle = \sum_{m \in \mathbb{Z}} K[m] \Delta x \quad (2.19)$$

and

$$K[n] = \begin{cases} 1 - \frac{|n|}{d} & \text{if } |n| \leq d \\ 0 & \text{otherwise.} \end{cases} \quad (2.20)$$

Note that $x = n\Delta x$. This discrete filter is implemented by Berggren et al. for the low-frequency and high-frequency cases, but not for the broadband case. As such, the same approach is taken here.

2.4 Procedure

The gradient-based optimization algorithm chosen to compare against Borg is MATLAB's `fmincon`. The `fmincon` function is actually a collection of algorithms that are used to minimize nonlinear scalar objective functions with multiple (possibly) nonlinear constraints. It is a default choice in MATLAB's Optimization Toolbox because it is generally effective for smooth objective functions with smooth constraints. Further details can be found in MATLAB's documentation. Both optimization algorithms call the same design evaluation function. A flowchart representation of this function is given in Figure 2.2. In short, given a vector of design variables, the overall mass and compliance of the design are first calculated. If either of these values violates the constraints, then the function returns these values along with a nominal objective value of 100. If the constraints are satisfied, then the function continues with the solution of the dynamic response of the design and ultimately returns the corresponding objective value. The constraints are thus enforced using a fixed penalty method.

The optimization procedure was carried out using the parameters given in Table 2.1. As in [2], three different optimizations were carried out, one for each of three sets of frequencies. The first case is broadband optimization, $\omega \in \mathcal{W}_{\text{BB}}$. Berggren et al. selected a frequency range from 300 Hz to 3400 Hz, with 50 equally-spaced frequencies within this range comprising the set \mathcal{W}_{BB} . The authors also conducted six additional narrowband test cases. However, [2] only presents results for two of these: 300 Hz to 800 Hz and 2300 Hz to 2800 Hz. From each of these frequency ranges, the authors again chose 50 equally-spaced evaluation frequencies to constitute the low-frequency set, \mathcal{W}_{LF} , and the high-frequency set, \mathcal{W}_{HF} .

For all three frequency sets, both Borg and `fmincon` were run for 500,000 objective function evaluations. In the case of `fmincon`, the gradient is estimated at each iteration using finite differences unless the user supplies a gradient function. As such, the actual number of function evaluations may be many times more than the number of iterations.

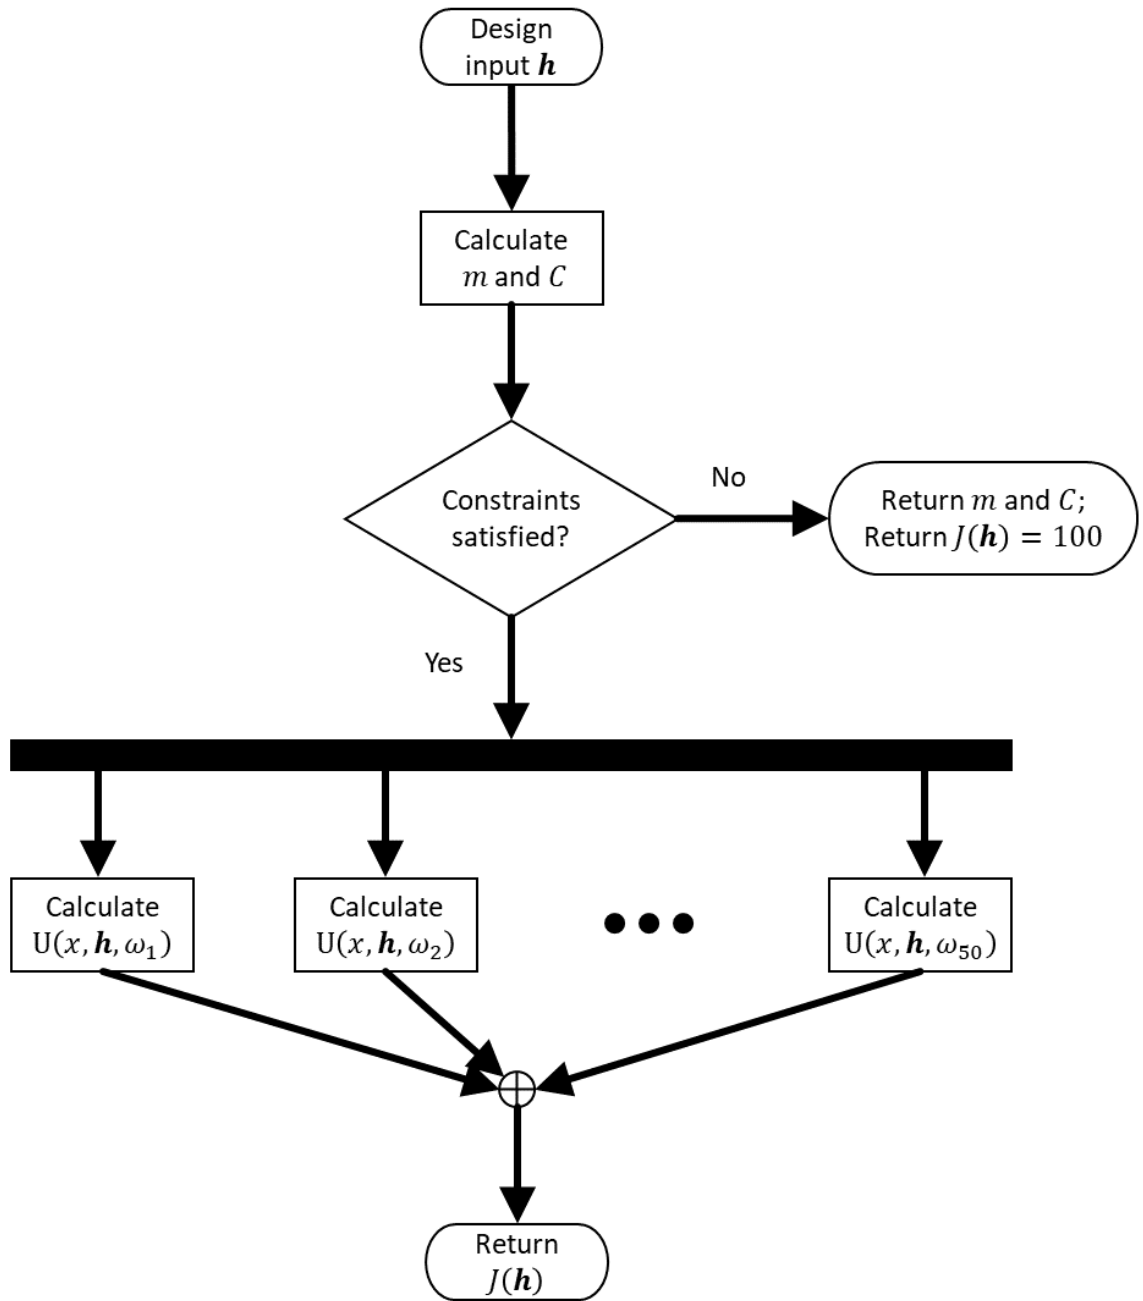


Figure 2.2. Flowchart representation of the design evaluation function called by both optimization algorithms.

Table 2.1. Parameters used in the optimization problem.

Material Parameters	
Young’s modulus, E	360 MPa
Density, ρ	$1100 \frac{\text{kg}}{\text{m}^3}$
Geometric Parameters	
Beam length, L_x	28 cm
Beam width, L_y	7 cm
Thickness limits, (h^-, h^+)	(2, 15) mm
Microphone thickness, h_{mic}	6 mm
Microphone region, (x_a, x_b)	(14, 17.92) cm
Runtime Parameters	
Broadband frequencies, \mathcal{W}_{BB}	$2\pi \times \left\{ 300, 300 + \frac{3100}{49}, \dots, 3400 \right\}$ Hz
Low frequencies, \mathcal{W}_{LF}	$2\pi \times \left\{ 300, 300 + \frac{500}{49}, \dots, 800 \right\}$ Hz
High frequencies, \mathcal{W}_{HF}	$2\pi \times \left\{ 2300, 2300 + \frac{500}{49}, \dots, 2800 \right\}$ Hz
Low-frequency filter radius, r_{LF}	4×2.8 mm
High-frequency filter radius, r_{HF}	5×2.8 mm
Mass constraint, γ_m	$1.0 \times m_{\text{ref}}$
Compliance constraint, γ_C	$1.6 \times C_{\text{ref}}$
Scaled driving force magnitude, F'	$\frac{12}{L_y E}$

Note: The reference mass and compliance, m_{ref} and C_{ref} , are those for a beam with a uniform thickness of 5 mm.

However, `fmincon` converged in fewer than 500,000 function evaluations for the present problem[‡], and so it was restarted several times with random starting designs until a total of 500,000 function evaluations was reached. The ‘optimal’ design shown below is thus the best of the set of designs converged upon by `fmincon`.

2.5 Results and discussion

This section presents the results of both Borg and `fmincon` for each of the three frequency sets: broadband frequencies ($\omega \in \mathcal{W}_{\text{BB}}$), low frequencies ($\omega \in \mathcal{W}_{\text{LF}}$), and high frequencies ($\omega \in \mathcal{W}_{\text{HF}}$). Finally, a general summary of the results, comparing Borg and `fmincon`, is also presented. Along with the results from Borg and `fmincon`, figures taken directly from the paper of Berggren et al. [2] are presented, both for reference and to compare

[‡]Note that in the high-frequency case, `fmincon`’s default tolerances were too low to produce convergence in 500,000 function evaluations or fewer, and so the minimum step size was increased from 10^{-10} to 10^{-8} .

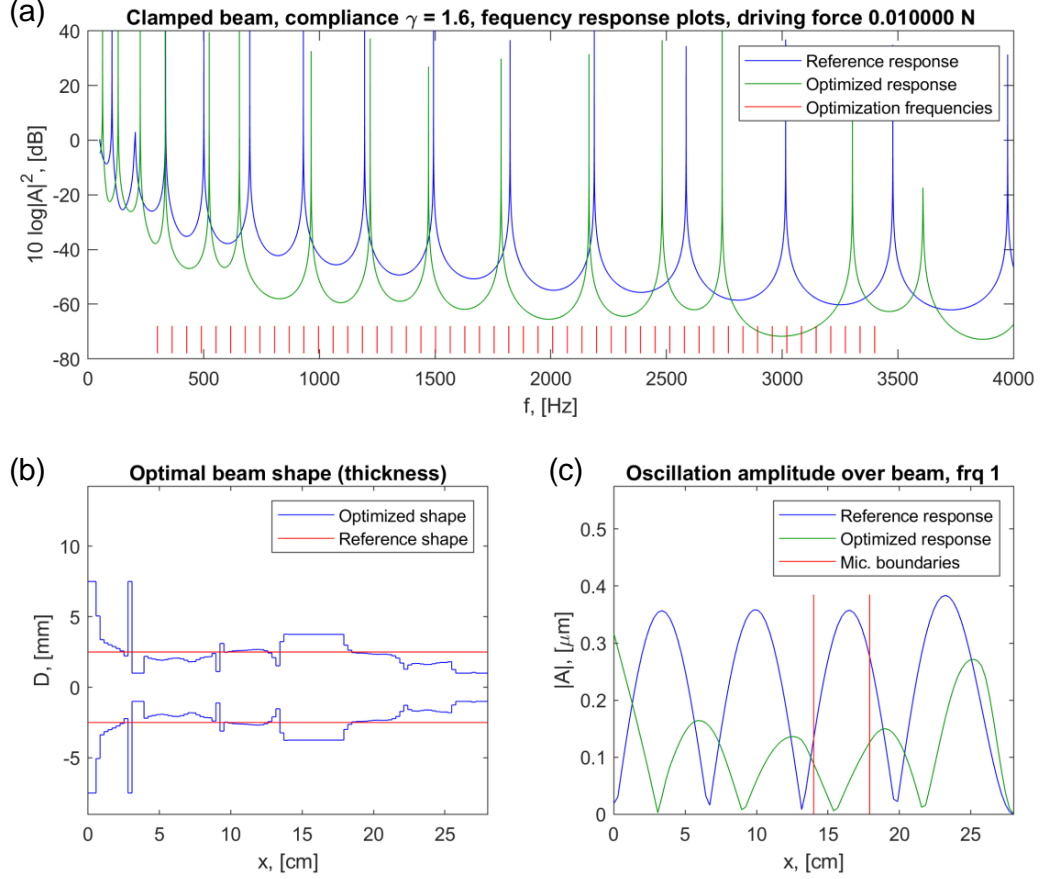


Figure 2.3. Optimization results for the broadband case, $\omega \in \mathcal{W}_{\text{BB}}$, using `fmincon`. The top plot (a) shows the spectrum of $J_\omega(\omega, \mathbf{h})$ in decibels for the optimal beam design (green) and a reference beam with a uniform thickness of 5 mm (blue). The evaluation frequencies are shown as red hatch marks. The bottom left plot (b) shows the optimal thickness distribution (blue) compared to the reference beam (red). The bottom right plot (c) shows the displacement amplitude, $|U|$, along the beam at $\omega = 2\pi \times 300$ Hz for the optimal beam design (green) and the reference beam (blue). The bounds of the microphone region are indicated by two red vertical lines.

against the results of this current work.

2.5.1 $\omega \in \mathcal{W}_{\text{BB}}$

Figures 2.3 and 2.4 show the results of the structural optimization using `fmincon` and `Borg`, respectively, for the case of broadband frequencies, $\omega \in \mathcal{W}_{\text{BB}}$. The corresponding plot from [2] is reproduced in Figure 2.5 for reference. To facilitate direct comparison, the optimization results are formatted in the same way as in [2] for each of the three

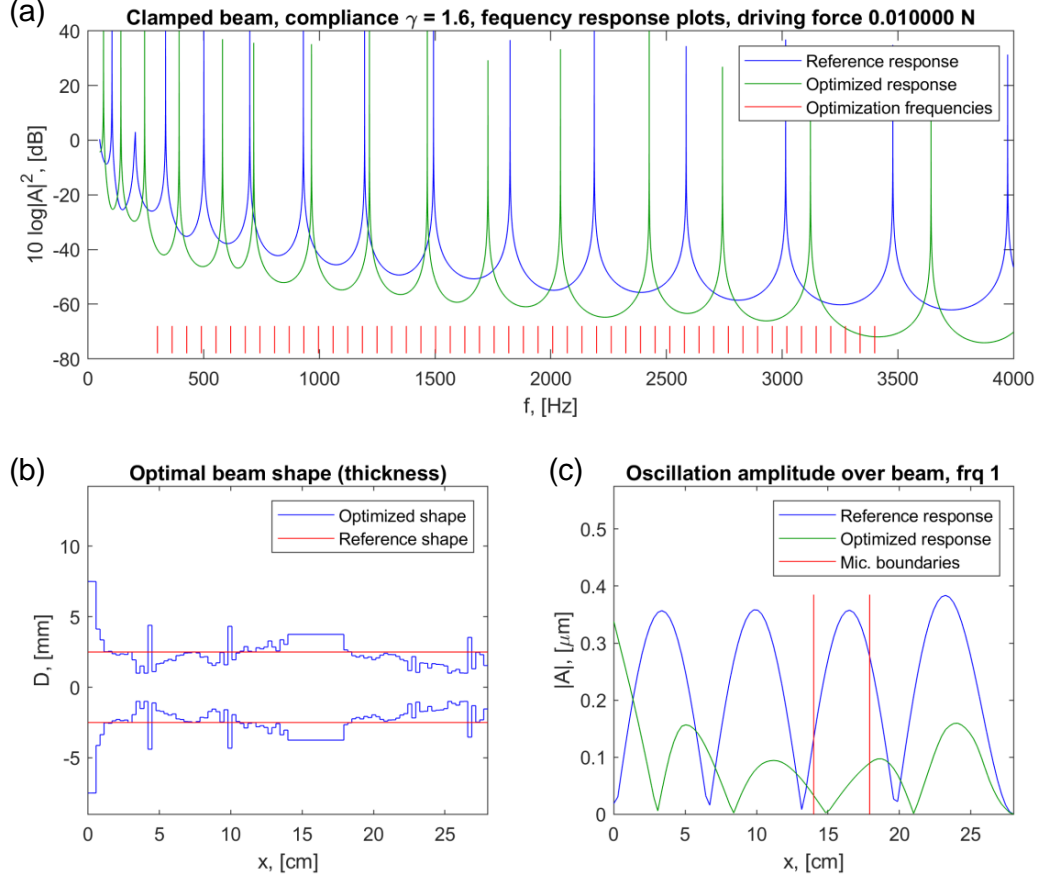


Figure 2.4. Optimization results for the broadband case, $\omega \in \mathcal{W}_{\text{BB}}$, using Borg. The top plot (a) shows the spectrum of $J_\omega(\omega, \mathbf{h})$ in decibels for the optimal beam design (green) and a reference beam with a uniform thickness of 5 mm (blue). The evaluation frequencies are shown as red hatch marks. The bottom left plot (b) shows the optimal thickness distribution (blue) compared to the reference beam (red). The bottom right plot (c) shows the displacement amplitude, $|U|$, along the beam at $\omega = 2\pi \times 300$ Hz for the optimal beam design (green) and the reference beam (blue). The bounds of the microphone region are indicated by two red vertical lines.

frequency sets. Taking Figure 2.3 as an example, the top section shows the full response spectrum for the optimal design (shown in green) as compared to a reference beam with a uniform thickness of 5 mm (shown in blue); the evaluation frequency set, \mathcal{W} , is shown as red hatch marks. The spectrum values represent the integral of the displacement amplitude across the microphone region. That is, the spectrum is calculated by evaluating the expression for $J_\omega(\omega, \mathbf{h})$ defined in Section 2.2. Adding together the spectrum values at the red hatch marks would give $\sum J_\omega = J(\mathbf{h})$. Note that although the peak heights may appear to differ between, e.g., Figures 2.3 and 2.5, the problem does not include

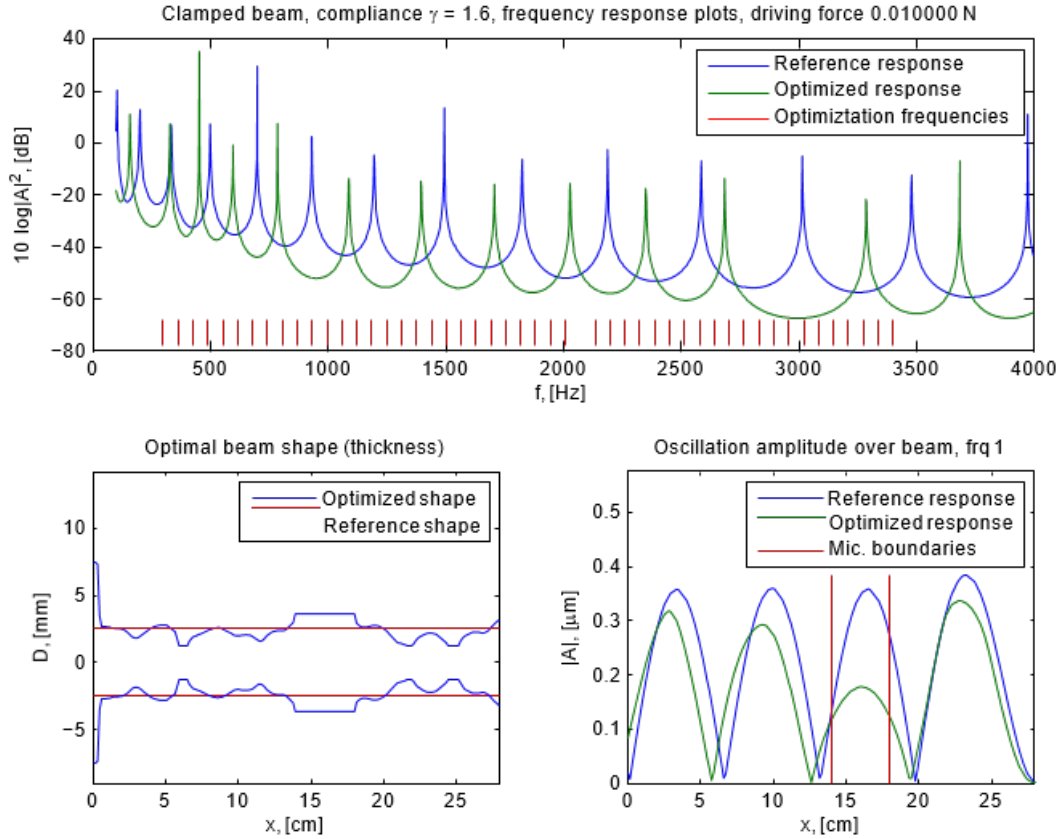


Figure 2.5. Figure 3 from [2], corresponding to Figures 2.3 and 2.4. Note that Berggren et al. use the variables D and A in place of h and U , respectively.

damping and so the displacement at resonance should theoretically be infinite. Therefore, the peak heights in the figures reproduced from [2] are artificially low, possibly due to undersampling. Looking again at Figure 2.3, a representation of the beam’s thickness profile is shown in the lower left corner, with the optimal design shown in blue and the reference beam shown in red. The lower right corner shows the optimal design’s displacement amplitude across the beam (shown in green) as compared to the reference beam (shown in blue) for the first frequency of the set. The microphone region, (x_a, x_b) , is indicated between two red vertical lines. As will be discussed below, altering the thickness profile acts to shift peaks in the response spectrum. As such, the two response plots in the lower right corner may have differing numbers of nodes and antinodes, in addition to different amplitude, because for the same evaluation frequency the effective wavenumber is different.

Comparing Figures 2.3 and 2.4, it is not immediately clear that the results from Borg and from `fmincon` share any similarities. It is worth pointing out that the optimal design

from Borg is much less ‘smooth’ than the `fmincon` design, insofar as it has more jumps in thickness from segment to segment. This is indicative of the nonlinear, stochastic search strategy of Borg as opposed to the linear gradient-based search strategy of `fmincon`. While the two designs differ noticeably, their mass distribution is similar, particularly at the excited end of the beam, which is substantially thicker than the rest of the beam. This is likely a mechanism to increase the input impedance by increasing mass at the drive point. Because the force is independent of the end thickness, the input power is inversely proportional to the impedance and so increasing the impedance has the effect of reducing the power transferred to the rest of the beam. The same strategy of increased mass at the excitation end is also seen in the results of Berggren et al. in Figure 2.5. The authors’ results are more similar to the results of `fmincon`, with the notable exception that the results of Berggren et al. shows a thickening of the beam near the clamped end, while the `fmincon` favored distributing the mass to other locations.

2.5.2 $\omega \in \mathcal{W}_{\text{LF}}$

The results paint a clearer picture in the narrowband cases. The respective plots for the low-frequency case, $\omega \in \mathcal{W}_{\text{LF}}$, are shown in Figures 2.6, 2.7, and 2.8. Similar to the broadband case, the forced end of the beam is significantly thicker than the rest of the beam for all three optimal designs—that from Borg, that from `fmincon`, and that from Berggren et al. All three also have distinct ‘lobes’ of increased thickness at regular intervals along the beam, although these lobes are not all in the same location nor the same size. In this sense, the results from `fmincon` and Berggren et al. are similar to one another. However, both the optimal design from Borg and the optimal design from Berggren et al. show a thickening around the microphone region, which may work in a similar way to the thickening of the excited end—namely, to reduce the response amplitude in that region given a certain input energy.

2.5.3 $\omega \in \mathcal{W}_{\text{HF}}$

Finally, the optimization results for the high-frequency case, $\omega \in \mathcal{W}_{\text{HF}}$, are shown in Figures 2.9 and 2.10 for `fmincon` and Borg, respectively, with the corresponding plot from [2] reproduced in Figure 2.11. This case is unique from the other two frequency ranges insofar as thickening of the excited end is not necessarily optimal, or is not as important. Rather, all three optimal designs show a distinct periodic structure, and there is effectively band gap behavior in the range $\omega \in \mathcal{W}_{\text{HF}}$ for the optimal designs from

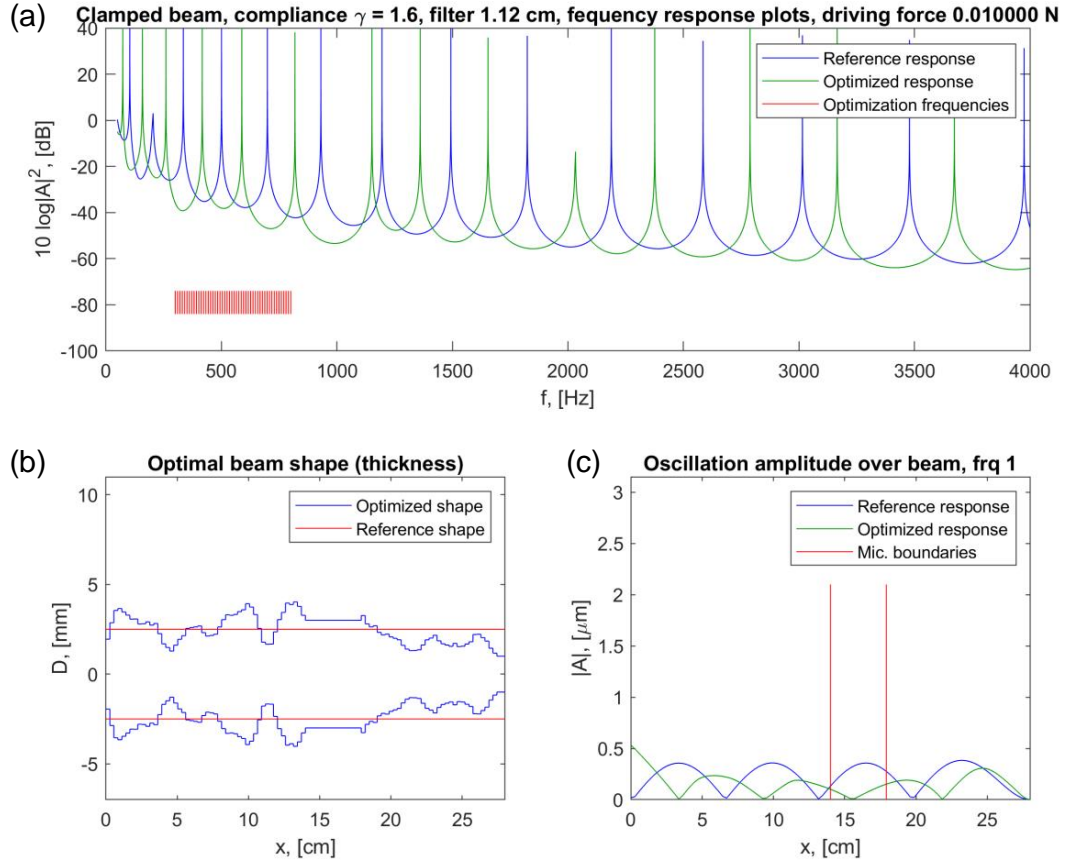


Figure 2.6. Optimization results for the low-frequency case, $\omega \in \mathcal{W}_{\text{LF}}$, using `fmincon`. The top plot (a) shows the spectrum of $J_\omega(\omega, \mathbf{h})$ in decibels for the optimal beam design (green) and a reference beam with a uniform thickness of 5 mm (blue). The evaluation frequencies are shown as red hatch marks. The bottom left plot (b) shows the optimal thickness distribution (blue) compared to the reference beam (red). The bottom right plot (c) shows the displacement amplitude, $|U|$, along the beam at $\omega = 2\pi \times 300$ Hz for the optimal beam design (green) and the reference beam (blue). The bounds of the microphone region are indicated by two red vertical lines.

Borg and `fmincon`. The mechanism of this behavior can be understood by looking at the dynamic displacement at several frequencies of interest. First, note that for the reference beam there is a resonance within the region of \mathcal{W}_{HF} , but that this resonance is absent for the optimized shapes. Figure 2.12 gives the dynamic response near this frequency, which shows that the vibration energy is concentrated away from the microphone region. Instead, there is greatly increased displacement concentrated near the first ‘lobe’ of the periodic structure. This behavior is akin to the antiresonance behavior of a discrete mass-spring system, wherein a natural mode shape has one degree of freedom unmoving.

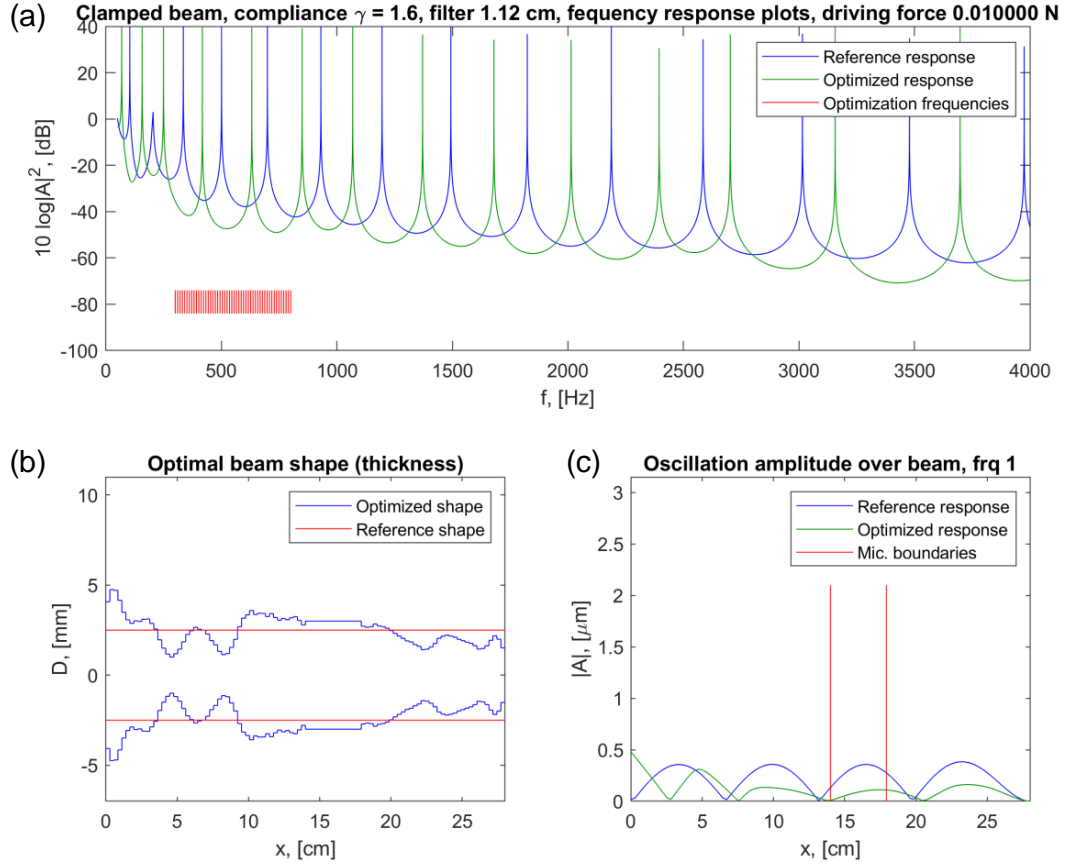


Figure 2.7. Optimization results for the low-frequency case, $\omega \in \mathcal{W}_{\text{LF}}$, using Borg. The top plot (a) shows the spectrum of $J_\omega(\omega, \mathbf{h})$ in decibels for the optimal beam design (green) and a reference beam with a uniform thickness of 5 mm (blue). The evaluation frequencies are shown as red hatch marks. The bottom left plot (b) shows the optimal thickness distribution (blue) compared to the reference beam (red). The bottom right plot (c) shows the displacement amplitude, $|U|$, along the beam at $\omega = 2\pi \times 300$ Hz for the optimal beam design (green) and the reference beam (blue). The bounds of the microphone region are indicated by two red vertical lines.

As a result, one would expect two new resonance peaks flanking the one original resonance frequency. Indeed, inspection of the spectra in Figure 2.10 shows this very splitting effect, with two new peaks surrounding the one peak in the reference spectrum. If the dynamic displacement is analyzed at these two new frequencies, as in Figure 2.13, it is clear that they correspond to the cases where vibration energy is moved away from the first lobe and focused at the other end of the beam. Effectively, the search strategies of both Borg and `fmincon` have resulted in the design of a vibroacoustic metamaterial, with a periodic structure tuned to push the dynamic response of the beam towards the lobed region of

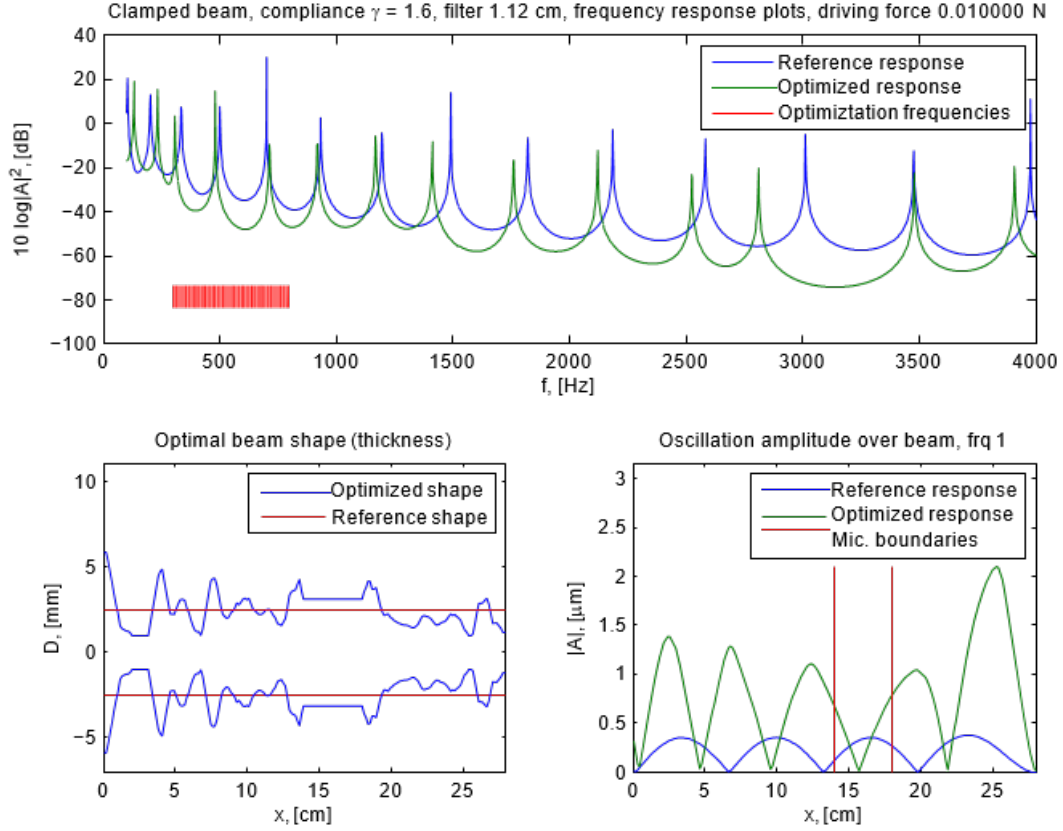


Figure 2.8. Figure 5 from [2], corresponding to Figures 2.6 and 2.7. Note that Berggren et al. use the variables D and A in place of h and U , respectively.

the beam for the analysis frequencies in \mathcal{W}_{HF} . Even the results of Berggren et al. show this periodic metamaterial design, although the period is significantly shorter and the lobes smaller in the results from Berggren et al.

2.5.4 Comparing `fmincon` and `Borg`

A summary of the optimization results for all three frequency sets is given in Table 2.2. Included in the table are the objective function value, J , the total mass normalized by the mass constraint, m/γ_m , and the static compliance normalized by the corresponding compliance constraint, C/γ_C . That is, the closer the value is to 1.00, the closer the design is to the limit of the respective constraint. It is clear from this summary that `Borg` produces an overall better design compared to that of `fmincon`. While the difference is not large in the broadband case, it is particularly stark in the high-frequency case, in which the best `Borg` design performs nearly 1000 times better than the best design of `fmincon`. It should be pointed out that all of the `Borg` designs lie at the very edge of the

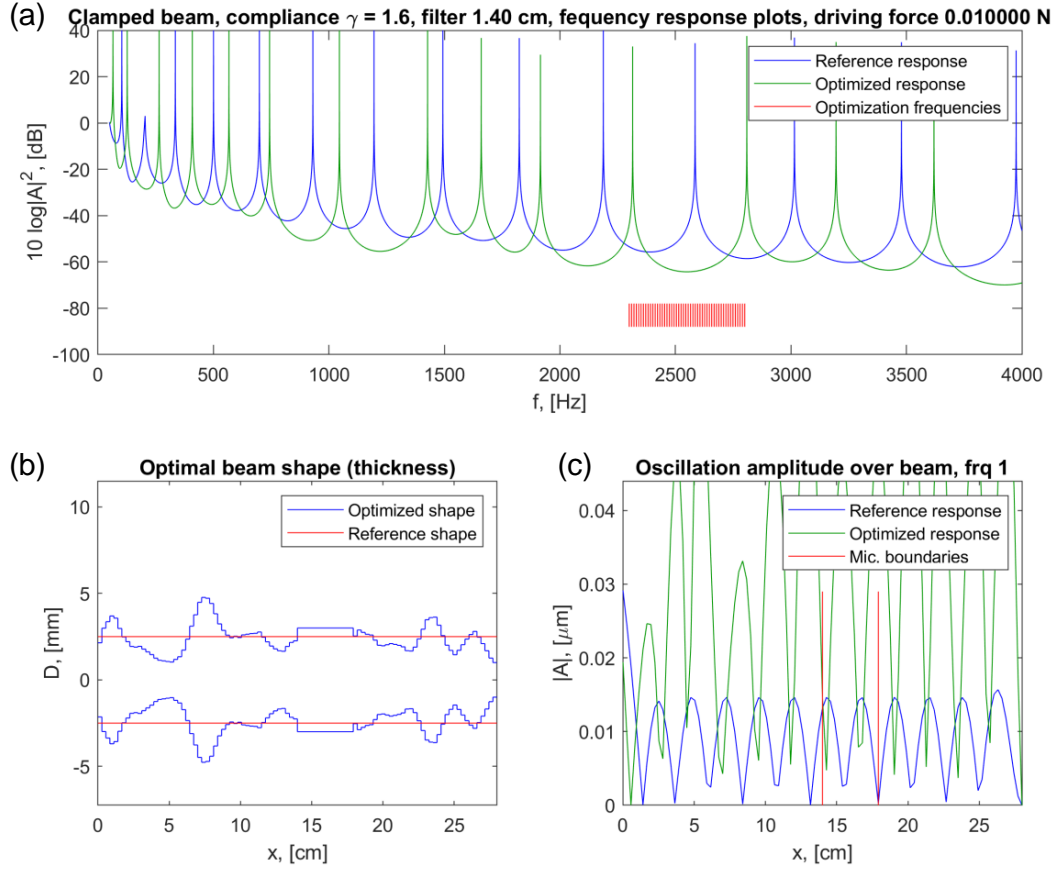


Figure 2.9. Optimization results for the high-frequency case, $\omega \in \mathcal{W}_{\text{HF}}$, using `fmincon`. The top plot (a) shows the spectrum of $J_\omega(\omega, \mathbf{h})$ in decibels for the optimal beam design (green) and a reference beam with a uniform thickness of 5 mm (blue). The evaluation frequencies are shown as red hatch marks. The bottom left plot (b) shows the optimal thickness distribution (blue) compared to the reference beam (red). The bottom right plot (c) shows the displacement amplitude, $|U|$, along the beam at $\omega = 2\pi \times 2300$ Hz for the optimal beam design (green) and the reference beam (blue). The bounds of the microphone region are indicated by two red vertical lines.

constrained design space, while the designs of `fmincon` are close but not always at the very edge.

It is not only the best designs that show this trend. Table 2.3 gives some statistical measures of the performance of each algorithm on these particular problems. Among these measures are the objective value and the number of function evaluations (NFEs) at each restart. That is, every time `fmincon` or Borg restarts, the objective function and NFEs are recorded; the measures are then calculated from these two sets. Because of the fundamentally different natures of the two algorithms, it is difficult to compare them

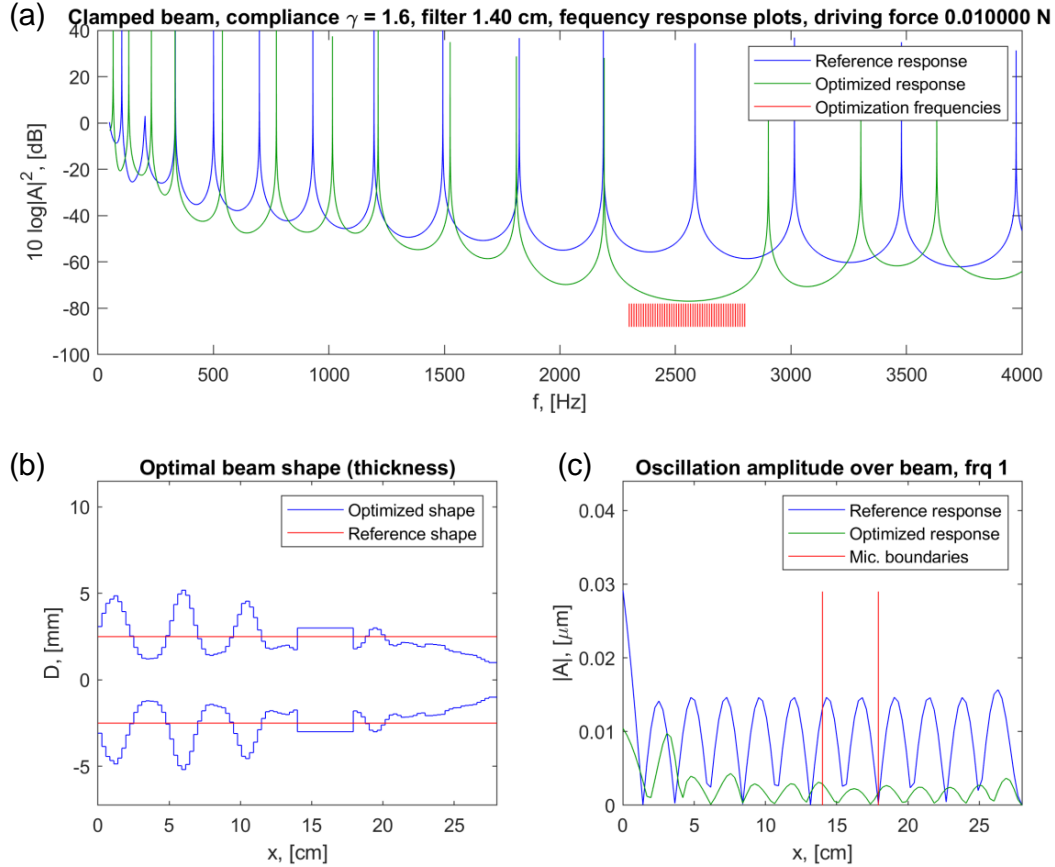


Figure 2.10. Optimization results for the high-frequency case, $\omega \in \mathcal{W}_{\text{HF}}$, using Borg. The top plot (a) shows the spectrum of $J_\omega(\omega, \mathbf{h})$ in decibels for the optimal beam design (green) and a reference beam with a uniform thickness of 5 mm (blue). The evaluation frequencies are shown as red hatch marks. The bottom left plot (b) shows the optimal thickness distribution (blue) compared to the reference beam (red). The bottom right plot (c) shows the displacement amplitude, $|U|$, along the beam at $\omega = 2\pi \times 2300$ Hz for the optimal beam design (green) and the reference beam (blue). The bounds of the microphone region are indicated by two red vertical lines.

one-to-one. Because Borg will restart more frequently after the first restart, the objective value, J , and the NFEs are given at the first restart. In contrast, the J and NFEs reported for `fmincon` are the averages for those two sets of recorded values. In this sense, Table 2.3 shows an approximation of the expected performance of a single run of each algorithm with no restarting. Under this interpretation, two things become apparent: 1) `fmincon` takes significantly fewer function evaluations to converge; and 2) Borg is more robust against the multimodal search space, reaching significantly better-performing designs before restarting. Indeed, Borg may take orders of magnitude more function




Figure 2.11. Figure 6 from [2], corresponding to Figures 2.9 and 2.10. Note that Berggren et al. use the variables D and A in place of h and U , respectively.




Figure 2.12. Optimal thickness distribution (blue) and displacement amplitude at $\omega = 2\pi \times 2557.32$ Hz (red) for the high-frequency case, $\omega \in \mathcal{W}_{\text{HF}}$, using Borg. Clearly, at 2557.32 Hz the vibration energy is concentrated at the left end away from the microphone region, (x_a, x_b) .




Figure 2.13. Optimal thickness distribution (blue) and displacement amplitude (red) for the high-frequency case, $\omega \in \mathcal{W}_{\text{HF}}$, using Borg. Analysis frequencies are $\omega = 2\pi \times 2191.23$ Hz (left) and $\omega = 2\pi \times 2902.02$ Hz (right). Whereas at 2557.32 Hz the vibration energy is concentrated at the drive point, here there is very little movement at the left end.

Table 2.2. Summary of optimization results for the three frequency sets. Most optimal objective value, J , along with the mass, m , and compliance, C , of the corresponding design, normalized by their respective constraints.

		fmincon	Borg
\mathcal{W}_{BB}	J	4.06×10^{-4}	2.69×10^{-4}
	m/γ_m	1.00	1.00
	C/γ_C	1.00	1.00
\mathcal{W}_{LF}	J	4.82×10^{-2}	7.42×10^{-3}
	m/γ_m	0.994	1.00
	C/γ_C	0.839	1.00
\mathcal{W}_{HF}	J	5.42×10^{-4}	9.18×10^{-7}
	m/γ_m	0.987	1.00
	C/γ_C	0.999	1.00

Table 2.3. Performance comparison of the two algorithms. Because Borg will restart more frequently after the first restart, the objective value, J , and the number of function evaluations (NFEs) are given at the first restart. In contrast, the J and NFEs reported for fmincon are the averages for all restarts.

		fmincon (average)	Borg (1 st restart)
\mathcal{W}_{BB}	J	5.44×10^{-3}	5.12×10^{-4}
	NFEs	415	100,982
\mathcal{W}_{LF}	J	7.73×10^0	1.75×10^{-2}
	NFEs	390	4,502
\mathcal{W}_{HF}	J	1.30×10^{-2}	2.01×10^{-4}
	NFEs	350	1,801

evaluations to converge, while `fmincon` may converge to designs that perform orders of magnitude worse.

In all three cases, \mathcal{W}_{BB} , \mathcal{W}_{LF} , and \mathcal{W}_{HF} , the primary outcome of the optimization was to shift peaks in the beam’s response to be outside the frequency range of interest and/or in between the analysis frequencies. By shifting the peaks away from the analysis frequencies, the peaks effectively ‘disappear’ from the point of view of the objective function. This phenomenon can be partially alleviated by integrating across the frequency range rather than evaluating it at discrete points. However, any analysis that involves a finite frequency range will be susceptible to this phenomenon to some degree because the peaks at the edge of the range can be shifted to be just outside of it.

2.6 Conclusions

As mentioned in the introduction, the two goals of this chapter were to build an effective optimization framework and to use that framework in the context of an example structural optimization problem and thereby compare the performance of Borg against that of a standard gradient-based algorithm. The selected test problem was taken from Berggren et al. [2] and involved tailoring the thickness profile of a cantilever beam to isolate a portion of it from vibration. As the results of both optimization algorithms show, the framework was successful in determining designs that achieve isolation that is orders of magnitude better than a reference uniform beam. Moreover, designs from both algorithms share commonalities with the results from Berggren et al., adding further credibility to the optimization framework.

In terms of comparing the two optimization algorithms, the results of this structural optimization study illustrate several important things. First, `fmincon` converges significantly faster than Borg. On average, `fmincon` converges in about 400 function evaluations, while Borg may take upwards of 100,000 function evaluations before it restarts. In this sense, `fmincon` is the preferred algorithm, especially when function evaluations are expensive. Second, the objective space is multimodal, as indicated by the multiple designs converged upon by `fmincon`. In this sense, Borg is the preferred algorithm with its global search strategy, including the use of automatic restarts. This is as opposed to a gradient-based algorithm like `fmincon`, which converges only to local minima by the nature of its design. Indeed, `fmincon` is expected to converge upon designs that perform on average orders of magnitude worse than those of Borg. Even after multiple runs with random seeding, the best design of Borg outperforms the best

design of `fmincon`.

Moreover, Borg naturally extends to multi-objective problems in a way that `fmincon` does not. For example, although inequality constraints were used in this test case, it may be advantageous to instead incorporate constraints like mass and compliance as additional objectives. One can then imagine a set of optimal designs that trade off one objective for another—for example, reducing the beam’s compliance but increasing its vibration response. It should be noted that there do exist multi-objective variants of common gradient-based methods, such as Newton’s method [63] and steepest descent [64]. However, as mentioned in the introduction, gradient-based methods are not robust against noisy and/or discontinuous objective spaces, which frequency-dependent objectives often lead to. This is one way in which Borg is unequivocally better suited than gradient-based algorithms like `fmincon`. This distinct advantage, along with the superior robustness demonstrated by the results of this chapter, mean that Borg was selected as the optimization algorithm for all subsequent work in this dissertation.

Chapter 3 | Optimization of an ABH at the end of a cantilever beam*

3.1 Introduction

As established in Chapter 1, structures whose thickness decreases according to a power law exhibit the ‘acoustic black hole’ (ABH) effect, which can be used to effectively reduce vibration. Such ABH vibration absorbers work by reducing the wavelength within the ABH region while simultaneously increasing displacement, thus focusing strain energy and more effectively dissipating energy. However, because the early ABH theory was based on the assumption that the thickness eventually tapers to zero, the only prescription on the power law exponent was that it be greater than or equal to two [6]. Similarly, limits on the length of the ABH region were only necessary to ensure a gradual change in thickness and thus minimize back scattering (see Equation (1.7)). Once a nonzero truncation thickness is introduced, the parameters of the power law will have a much greater effect on the ABH’s performance. While it is possible to derive analytical expressions for reflection in the ABH taper, predicting the performance of an ABH vibration absorber as part of a larger structure can become more complicated. This is where a rigorous optimization framework can help elucidate.

An optimization framework was established in Chapter 2 using beam displacement at discrete frequencies as a metric with the Borg evolutionary algorithm. In particular, the transfer matrix method (TMM) was used to analyze the dynamics of a beam with a complex thickness distribution and an objective function was developed that condensed the dynamic response of this beam to a format usable by Borg. This chapter applies

*A portion of the work presented in this chapter was also published in the Journal of the Acoustical Society of America [65].




Figure 3.1. Graphical illustration of the optimization problem. The beam has overall length L_x and the base beam has thickness h_1 , while the ABH taper is determined by its length, L_{ABH} , its minimum thickness, h_0 , and its taper power, m . The composite structure is driven at a point, x_f , with a point force of magnitude F . The boundary at $x = 0$ is modelled as free and the boundary at $x = L_x$ is modelled as fixed.

a similar framework to the problem of a one-dimensional ABH vibration absorber at the end of a beam. Notable differences are an improved TMM described in Section 3.3 and a carefully formulated quadrature to better capture the broadband dynamics in the objective function.

3.2 Problem description

The problem considered is that of a one-dimensional acoustic black hole (ABH) vibration absorber at the end of a cantilever beam. The taper profile of the ABH is of the form $h(x) = \varepsilon x^m + h_0$, where the coefficient ε can be uniquely determined by the ABH's length, L_{ABH} , its minimum thickness, h_0 , its maximum thickness, h_1 , and its taper power, m , by the relation

$$\varepsilon = \frac{h_1 - h_0}{(L_{ABH})^m}. \quad (3.1)$$

The length and thickness of the beam will be held fixed, so the only design variables left are the ABH length, minimum thickness, and taper power, which can be represented as a three-element vector, $\mathbf{h} = [L_{ABH} \ h_0 \ m]^T$. A graphical depiction of the problem and design variables is shown in Figure 3.1. When the beam is driven by a harmonic force at the point x_f , the goal is to minimize the overall beam response in some frequency band,

$[\omega_a, \omega_b]$. Formally, the problem can be written as:

$$\begin{aligned} \min_{\mathbf{h} \in \mathcal{H}} J(\mathbf{h}) &= C \int_{\omega_a}^{\omega_b} \langle \dot{U}^2(\omega, \mathbf{h}) \rangle d\omega = \frac{C}{L_x} \int_{\omega_a}^{\omega_b} \int_0^{L_x} |i\omega U(x, \omega, \mathbf{h})|^2 dx d\omega \\ &\text{subject to} \\ \mathcal{H} &= \left\{ \mathbf{h} \in \mathbb{R}^3 : 0 \leq L^- \leq L_{\text{ABH}} \leq L^+, 0 \leq h^- \leq h_0 \leq h^+, 0 \leq m^- \leq m \leq m^+ \right\} \\ &\quad -\omega^2 \rho L_y h(x) U + \frac{\partial^2}{\partial x^2} \left(E \frac{L_y h^3(x)}{12} U_{,xx} \right) = f(x, \omega), \quad \forall x \in (0, L_x) \\ &\quad \left. \begin{aligned} f(x) &= F \delta(x - x_f) \\ U_{,xx}(0) &= U_{,xxx}(0) = 0 \\ U(L_x) &= U_{,x}(L_x) = 0 \end{aligned} \right\} \forall \omega \in [\omega_a, \omega_b] \end{aligned}$$

where

$$C \int_{\omega_a}^{\omega_b} \langle \dot{U}^2(\omega, \mathbf{0}) \rangle d\omega = 1$$

and

$$h(x) = \begin{cases} \varepsilon x^m + h_0 & \text{for } 0 \leq x \leq L_{\text{ABH}} \\ h_1 & \text{for } L_{\text{ABH}} < x \leq L_x. \end{cases}$$

In words, the above states that the goal of the problem is to find the vector of design variables, \mathbf{h} , that minimizes the objective function, $J(\mathbf{h})$. The design variables, \mathbf{h} , are restricted to be in the set \mathcal{H} of valid designs, which is defined by the minimum and maximum allowable value of each design variable. The objective function, $J(\mathbf{h})$ is calculated by integrating the spatially-averaged squared velocity response, $\langle \dot{U}^2 \rangle$, within the frequency range $[\omega_a, \omega_b]$. A normalization factor, C , is applied to the objective function so that $J(\mathbf{0}) = 1$, where $\mathbf{h} = \mathbf{0}$ denotes a reference beam of uniform thickness h_1 . The velocity response, \dot{U} , is calculated from the dynamic displacement response, U , by multiplication of $i\omega$. This is because it is assumed that the displacement is time-harmonic—that is, $u(x, t) = U(x)e^{i\omega t}$. Therefore, the time derivative of the displacement will be $\dot{u}(x, t) = i\omega U(x)e^{i\omega t}$. The displacement is calculated by solving the dynamic Euler-Bernoulli beam equation, which acts as a PDE constraint on the optimization. Because the solution to the equation will depend on the value of ω , and the thickness profile, $h(x)$, U is also dependent on frequency and the design variables contained in \mathbf{h} . When solving the beam equation, the external force is a harmonic point force with magnitude F applied at $x = x_f$. The boundary conditions are modelled as free at $x = 0$ and fixed at $x = L_x$. That is, the displacement at the minimum thickness tip of the

ABH is unconstrained. This means that the displacement could approach infinity if the beam is excited at resonance. As such, the thinnest 25% of the ABH was modelled as having a viscoelastic damping layer attached to one side. This added damping layer was incorporated into the model as a complex bending stiffness, as will be described in the next section. A small amount of material damping (see Table 3.1 on page 61) was also modelled in the bulk beam material for the case that $L_{\text{ABH}} = 0$.

3.3 Theory

3.3.1 Damping model

The damping model used is that of Oberst [5], which is applicable to thin beams and plates with attached, unconstrained, extensional damping layers of homogeneous viscoelastic material, when the damping material has a much greater loss factor (which is the case in any practical application). The model of Ross, Ungar, and Kerwin [28] appears to be more popular in the ABH literature. However, the primary concern of [28] is the modelling of constrained layer damping, and the authors note that the model of Oberst is more exact for the case of extensional damping.

Recall that the bending stiffness of an Euler-Bernoulli beam is $D = EI$, where E is the Young's modulus of elasticity of the material and I is the second moment of area of the beam's cross section. If the bending stiffness of the base beam is $\widetilde{D} = D(1 + i\eta)$ and the bending stiffness of the damping material is $\widetilde{D}_d = D_d(1 + i\eta_d)$, then the effective complex bending stiffness, $\widetilde{D}_{\text{eff}} = D_{\text{eff}}(1 + i\eta_{\text{eff}})$, can be found analytically by integrating the combined strains and comparing the resultant torques to the Euler-Bernoulli beam model.

Assumptions of the Euler-Bernoulli beam model include: 1) deflections are much smaller than the beam thickness; 2) lines orthogonal to the neutral axis remain orthogonal; and 3) the neutral axis does not change length. Assumption 1 means deflection angles can be approximated as $\theta \approx \frac{du}{dx}$. Assumptions 2 and 3 mean that an infinitesimal segment of length dx , after deformation, forms part of an annulus. If r is defined as the distance from the center of the annulus to the neutral axis, then after deformation the infinitesimal length of the neutral axis—which remains unchanged according to assumption 3—is related to an infinitesimal angle as $dx = rd\theta$. Away from the neutral axis, the new length




Figure 3.2. Graphic of a two-layer beam. The neutral axis is denoted by a dashed line.

is $dx' = (r - z)d\theta$, according to assumption 2. Therefore, the strain in the x -direction is

$$\epsilon_x(z) = \frac{\Delta l}{l} = \frac{dx' - dx}{dx} = -z \frac{d\theta}{dx} = -z \frac{d^2 u}{dx^2}. \quad (3.2)$$

The stress in the x -direction is given simply by $\sigma_x = E\epsilon_x$.

In the case of two layers, it is not yet explicit where the neutral axis lies. However, the distance from the neutral axis to the interface between the two materials can be denoted δ , as shown in Figure 3.2. If the beam is in static deflection, then at any point along the beam the net stress is zero. That is,

$$\begin{aligned} \int_A \sigma_x \, dA &= \int_{\delta-h}^{\delta+h_d} \int_0^{L_y} \sigma_x \, dy dz \\ &= -EL_y \int_{\delta-h}^{\delta} z \frac{d^2 u}{dx^2} \, dz - E_d L_y \int_{\delta}^{\delta+h_d} z \frac{d^2 u}{dx^2} \, dz = 0, \end{aligned} \quad (3.3)$$

where L_y is the width of the composite beam, h is the thickness of the base beam, h_d is the thickness of the damping layer, E is the Young's modulus of the base beam, and E_d is the Young's modulus of the damping layer. Evaluating the integral and rearranging gives the following expression for the z -coordinate of the neutral axis:

$$\delta = \frac{1}{2} \cdot \frac{Eh^2 - E_d h_d^2}{Eh + E_d h_d}. \quad (3.4)$$

Clearly, in the case of no damping layer the neutral axis is at the center of the base beam. In the case that $E_d \ll E$ and h and h_d are comparable—as is often the case in practice—the neutral axis will not be far from that point.

To find the effective complex bending stiffness, it is sufficient to compare the torque

in the composite beam with that for a completely homogeneous beam. At any point x along the beam, the net torque about the neutral axis is[†]

$$\widetilde{M}(x) = \int_A z \widetilde{\sigma}_x \, dA. \quad (3.5)$$

Note the tilde over the characters, indicating complex quantities. This is because we now consider the dynamic quantities, whereas before the forces were static, and so the quantities were entirely real. In practice, this means that the entirely real Young's modulus, E , becomes a complex value, $\widetilde{E} = E(1 + i\eta)$, where η models the viscous losses in the material.

In the case of a homogeneous beam, the dynamic torque reduces to $\widetilde{M}(x) = -\widetilde{D} \frac{d^2 u}{dx^2}$. For the current case of two materials of comparable thickness, the full integral is

$$\int_{\delta-h}^{\delta+h_d} \int_0^{L_y} z \widetilde{\sigma}_x \, dy dz = -\widetilde{E} L_y \int_{\delta-h}^{\delta} z^2 \frac{d^2 u}{dx^2} \, dz - \widetilde{E}_d L_y \int_{\delta}^{\delta+h_d} z^2 \frac{d^2 u}{dx^2} \, dz. \quad (3.6)$$

Comparing this to the torque for a simple homogeneous beam, it is clear that the effective bending stiffness is

$$\begin{aligned} \widetilde{D}_{\text{eff}} &= \widetilde{E} L_y \int_{\delta-h}^{\delta} z^2 \, dz + \widetilde{E}_d L_y \int_{\delta}^{\delta+h_d} z^2 \, dz \\ &= L_y (\widetilde{E} h + \widetilde{E}_d h_d) \delta^2 - L_y (E h^2 - E_d h_d^2) \delta + \frac{L_y}{3} (E h^3 + E_d h_d^3). \end{aligned} \quad (3.7)$$

Substituting the z -coordinate of the neutral axis found from Equation (3.4), and using the fact that $\widetilde{E} = E(1 + i\eta)$ and $\widetilde{E}_d = E_d(1 + i\eta_d)$, the full expression becomes

$$\begin{aligned} \widetilde{D}_{\text{eff}} &= \frac{L_y}{12} \left(\frac{E^2 h^4 [(1 - \eta^2) p + 2\eta q]}{p^2 + q^2} + \frac{s [(1 - \eta \eta_d) p + (\eta + \eta_d) q]}{p^2 + q^2} + \frac{E_d^2 h_d^4 [(1 - \eta_d^2) p + 2\eta_d q]}{p^2 + q^2} \right) \\ &+ \frac{i L_y}{12} \left(\frac{E^2 h^4 [2\eta p - (1 - \eta^2) q]}{p^2 + q^2} + \frac{s [(\eta + \eta_d) p - (1 - \eta \eta_d) q]}{p^2 + q^2} + \frac{E_d^2 h_d^4 [2\eta_d p - (1 - \eta_d^2) q]}{p^2 + q^2} \right), \end{aligned} \quad (3.8)$$

where

$$p = E h + E_d h_d, \quad (3.9)$$

$$q = \eta E h + \eta_d E_d h_d, \quad (3.10)$$

[†]Many include a negative sign in front of the integral so that a positive torque induces a positive curvature. The opposite convention is used here, so that a positive torque induces positive power flow in the positive x -direction.

and

$$s = 2EE_d (2h^3h_d + 3h^2h_d^2 + 2hh_d^3). \quad (3.11)$$

After significant rearranging, the effective bending stiffness can be put in the form $\widetilde{D}_{\text{eff}} = D_{\text{eff}}(1 + i\eta_{\text{eff}})$ with

$$D_{\text{eff}} = D \cdot \frac{1 + 2a(2\xi + 3\xi^2 + 2\xi^3) + a^2\xi^4}{1 + a\xi} \quad (3.12)$$

and

$$\eta_{\text{eff}} = \eta_d \cdot \frac{(a\xi)(3 + 6\xi + 4\xi^2 + 2a\xi^3 + a^2\xi^4)}{(1 + a\xi)(1 + 2a(2\xi + 3\xi^2 + 2\xi^3) + a^2\xi^4)}, \quad (3.13)$$

where $a = \frac{E_d}{E}$, $\xi = \frac{h_d}{h}$, and D is the real bending stiffness of the base beam, $D = \text{Re}\{\widetilde{D}\} = E\frac{L_y h^3}{12}$. Note that the effective bending stiffness, D_{eff} , will generally be larger than D , meaning the composite structure will be stiffer than the base beam alone. Furthermore, to properly model the added damping layer, one should account for the added mass by using an effective density, $\rho_{\text{eff}} = \rho + \rho_d \xi$, where ρ_d is the density of the damping material.

3.3.2 Block Riccati transfer matrix method

For evaluation of the beam's dynamic response, the transfer matrix method (TMM) detailed in Chapter 2 is used again. Recall that a complex beam geometry can be divided into N segments that are each approximated as having uniform thickness. The dynamic response at the two ends of segment i can be related by a (generally complex) transfer matrix, \mathbf{Z}_i , such that $\mathbf{u}_{i+1} = \mathbf{Z}_i \mathbf{u}_i$. Here, \mathbf{u}_i is a vector of state variables (displacement, rotation, torque, and shear) at node i and similarly for \mathbf{u}_{i+1} . Once the state variables at the beam's boundaries are completely determined, all state variables at intermediate points can be calculated. In the case of an external force at node i , the transfer relation changes to

$$\mathbf{u}_{i+1} = \mathbf{Z}_i \mathbf{u}_i + \mathbf{d}_i, \quad (3.14)$$

where \mathbf{d}_i is a vector of external perturbations in the state variables.

For the accurate representation of smooth geometries, it is necessary to have many small segments, especially at higher frequencies. Unfortunately, many segments means numerical error can propagate and cause numerical instabilities in the standard TMM. An example of such instability is shown in Figure 3.3, which shows the drive-point mobility at two locations along a simply-supported beam. In this particular test case, the instability




Figure 3.3. Illustration of the instabilities that arise from the basic transfer matrix method. Pictured are drive-point mobilities for a simply-supported beam of length L .

occurs above 4,500 Hz, though the particular point of instability will of course vary from case to case. This phenomenon has been reported as far back as 1965 in the context of compression waves through layered elastic media [66]. In addition to elastic waves, the TMM is often employed in the solution of problems involving electron-transport [67], optics [68], and quantum physics [69]. As such, there have been multiple approaches to controlling or eliminating the instability inherent in the TMM, including the scattering matrix method [70], the enhanced transmittance method [71], and the global matrix method [72].

The approach used here will be a modified formulation of the TMM developed by Horner and Pilkey, called the Riccati transfer matrix method (RTMM) [73]. The advantage of the RTMM is that, like the generalized Riccati transformation, it converts a numerically unstable two-point boundary value problem into a stable initial value problem [74]. Moreover, the RTMM requires no numerical integration to solve the transformed system and the matrix components can be determined analytically.

A Riccati equation, in the general sense, is a first-order differential equation of the form

$$\frac{dy}{dx} = a(x) + b(x)y + c(x)y^2. \quad (3.15)$$

This form of equation is named after Jacopo F. Riccati, who studied several particular equations of this form in the early 18th century [75]. Solutions to the Riccati equation

are now well understood, and so it is particularly useful that any second-order differential equation of the form

$$y'' = \alpha(x)y' + \beta(x)y \quad (3.16)$$

can be transformed into an equivalent Riccati equation. In particular, using the ‘Riccati transformation’ $y = R(x)y'$ allows the substitution

$$\begin{aligned} y'' &= \alpha y' + \beta y \\ \implies Ry'' &= \alpha Ry' + \beta Ry \\ \implies y' - R'y' &= \alpha Ry' + \beta R^2 y' \\ \implies R' &= 1 - \alpha R - \beta R^2. \end{aligned} \quad (3.17)$$

Thus, a second-order equation in $y(x)$ has been transformed into a Riccati equation in $R(x)$, with $a(x) = 1$, $b(x) = -\alpha(x)$, and $c(x) = -\beta(x)$. In many physical systems, this reduces a second-order boundary value problem to a first-order initial value problem. Rybicki and Usher [74] used this Riccati transformation as inspiration for the name of their own transformation, termed the “generalized Riccati transformation.”

Any system of 1st-order ordinary differential equations can be written in the form

$$\frac{d}{dx} \begin{bmatrix} \phi_1(x) \\ \phi_2(x) \end{bmatrix} = \begin{bmatrix} \Gamma_{11}(x) & \Gamma_{12}(x) \\ \Gamma_{21}(x) & \Gamma_{22}(x) \end{bmatrix} \begin{bmatrix} \phi_1(x) \\ \phi_2(x) \end{bmatrix} + \begin{bmatrix} \mathbf{g}_1(x) \\ \mathbf{g}_2(x) \end{bmatrix}, \quad (3.18)$$

with boundary conditions

$$\phi_1(a) = \mathbf{c}_1 \quad (3.19)$$

and

$$\phi_2(b) = \mathbf{c}_2, \quad (3.20)$$

where ϕ_1 , \mathbf{g}_1 , and \mathbf{c}_1 are $n_1 \times 1$ vectors; and ϕ_2 , \mathbf{g}_2 , and \mathbf{c}_2 are $n_2 \times 1$ vectors. The generalized Riccati transformation defines two new variables, \mathbf{R}_{12} and ψ_1 , such that $\phi_1(x) = \mathbf{R}_{12}(x)\phi_2(x) + \psi_1(x)$. The similarities to the Riccati transformation are clear, especially when Equation (3.16) is written in the equivalent form

$$\frac{d}{dx} \begin{bmatrix} y \\ y' \end{bmatrix} = \begin{bmatrix} 0 & 1 \\ \beta(x) & \alpha(x) \end{bmatrix} \begin{bmatrix} y \\ y' \end{bmatrix}. \quad (3.21)$$

From the generalized Riccati transformation, one can derive the following relations:

$$\phi_1' = \mathbf{R}'_{12}\phi_2 + \mathbf{R}_{12}\phi_2' + \psi_1' \quad (3.22)$$

and

$$\phi_2' = \mathbf{\Gamma}_{21}\phi_1 + \mathbf{\Gamma}_{22}\phi_2 + \mathbf{g}_2 = (\mathbf{\Gamma}_{22} + \mathbf{\Gamma}_{21}\mathbf{R}_{12})\phi_2 + \mathbf{\Gamma}_{21}\psi_1 + \mathbf{g}_2. \quad (3.23)$$

These relations can be used to write a differential equation in terms of ψ_1 :

$$\begin{aligned} \psi_1' &= (\mathbf{\Gamma}_{11} - \mathbf{R}_{12}\mathbf{\Gamma}_{21})\psi_1 \\ &\quad - (\mathbf{R}'_{12} - \mathbf{\Gamma}_{12} + \mathbf{R}_{12}\mathbf{\Gamma}_{22} - \mathbf{\Gamma}_{11}\mathbf{R}_{12} + \mathbf{R}_{12}\mathbf{\Gamma}_{21}\mathbf{R}_{12})\phi_2 \\ &\quad + (\mathbf{g}_1 - \mathbf{R}_{12}\mathbf{g}_2), \end{aligned} \quad (3.24)$$

with

$$\psi_1(a) = \mathbf{c}_1 - \mathbf{R}_{12}(a)\phi_2(a). \quad (3.25)$$

The dependence on ϕ_2 can be eliminated from Equations (3.24) and (3.25) if

$$\mathbf{R}'_{12} = \mathbf{\Gamma}_{12} - (\mathbf{R}_{12}\mathbf{\Gamma}_{22} - \mathbf{\Gamma}_{11}\mathbf{R}_{12}) - \mathbf{R}_{12}\mathbf{\Gamma}_{21}\mathbf{R}_{12} \quad (3.26)$$

and

$$\mathbf{R}_{12}(a) = \mathbf{0}. \quad (3.27)$$

It should be evident Equation (3.26) is a Riccati equation, albeit in matrix form. In the case that \mathbf{R}_{12} satisfies the above Riccati equation, Equations (3.24) and (3.25) simplify to

$$\psi_1' = (\mathbf{\Gamma}_{11} - \mathbf{R}_{12}\mathbf{\Gamma}_{21})\psi_1 + (\mathbf{g}_1 - \mathbf{R}_{12}\mathbf{g}_2), \quad (3.28)$$

with

$$\psi_1(a) = \mathbf{c}_1. \quad (3.29)$$

In summary, the procedure to solve the original system, Equations (3.18)–(3.20), is

1. Solve Equations (3.26) and (3.27) for $\mathbf{R}_{12}(x)$
2. Solve Equations (3.28) and (3.29) for $\psi_1(x)$
3. Solve Equations (3.23) and (3.20) for $\phi_2(x)$
4. Use the generalized Riccati transformation, $\phi_1(x) = \mathbf{R}_{12}(x)\phi_2(x) + \psi_1(x)$, to solve for $\phi_1(x)$

With this in mind, we move on to the derivation of the Riccati transfer matrix method. Firstly, Equation (3.14) is modified to be of the form

$$\begin{bmatrix} \mathbf{v} \\ \mathbf{w} \end{bmatrix}_{i+1} = \begin{bmatrix} \mathbf{Z}^{vv} & \mathbf{Z}^{vw} \\ \mathbf{Z}^{wv} & \mathbf{Z}^{ww} \end{bmatrix}_i \begin{bmatrix} \mathbf{v} \\ \mathbf{w} \end{bmatrix}_i + \begin{bmatrix} \mathbf{d}^v \\ \mathbf{d}^w \end{bmatrix}_i, \quad (3.30)$$

where \mathbf{v}_i is a 2×1 subvector containing the state variables that are homogeneous at $x = 0$; that is, $\mathbf{v}_1 = [0 \ 0]^T$. \mathbf{w}_i is a 2×1 subvector containing the complementary nonhomogeneous state variables. For the cantilever beam free at $x = 0$, the homogeneous state variables are shear, Q , and torque, M . The complimentary state variables are then displacement, U , and rotation, θ , respectively. In such a case, the subvector $\mathbf{v}_i = [Q_i \ M_i]^T$ and the subvector $\mathbf{w}_i = [U_i \ \theta_i]^T$. The subvectors \mathbf{d}^v and \mathbf{d}^w then represent external perturbations in the respective state variables, and the elements of \mathbf{Z}_i are rearranged accordingly.

The generalized Riccati transformation that relates \mathbf{v}_i to \mathbf{w}_i , as seen earlier, is

$$\mathbf{v}_i = \mathbf{R}_i \mathbf{w}_i + \mathbf{p}_i, \quad (3.31)$$

where \mathbf{R}_i is an as-of-yet undetermined 2×2 matrix and \mathbf{p}_i is an as-of-yet undetermined 2×1 vector. However, because $\mathbf{v}_1 = [0 \ 0]^T$, it is clear that $\mathbf{R}_1 = \mathbf{0}$ and $\mathbf{p}_1 = \mathbf{0}$. A recurrence relation for \mathbf{R}_{i+1} and \mathbf{p}_{i+1} is then sufficient to solve for all \mathbf{R}_i and \mathbf{p}_i . Using Equation (3.31) with Equation (3.30), the following relations can be derived:

$$\mathbf{v}_{i+1} = \left(\mathbf{Z}_i^{vv} \mathbf{R}_i + \mathbf{Z}_i^{vw} \right) \mathbf{w}_i + \left(\mathbf{Z}_i^{vv} \mathbf{p}_i + \mathbf{d}_i^v \right) \quad (3.32)$$

and

$$\mathbf{w}_{i+1} = \left(\mathbf{Z}_i^{wv} \mathbf{R}_i + \mathbf{Z}_i^{ww} \right) \mathbf{w}_i + \left(\mathbf{Z}_i^{wv} \mathbf{p}_i + \mathbf{d}_i^w \right) \quad (3.33)$$

Using Equation (3.33) to solve for \mathbf{w}_i and substituting that into Equation (3.32) gives

$$\begin{aligned} \mathbf{v}_{i+1} &= \left(\mathbf{Z}_i^{vv} \mathbf{R}_i + \mathbf{Z}_i^{vw} \right) \left(\mathbf{Z}_i^{wv} \mathbf{R}_i + \mathbf{Z}_i^{ww} \right)^{-1} \mathbf{w}_{i+1} \\ &+ \left(\mathbf{Z}_i^{vv} \mathbf{p}_i + \mathbf{d}_i^v \right) - \left(\mathbf{Z}_i^{vv} \mathbf{R}_i + \mathbf{Z}_i^{vw} \right) \left(\mathbf{Z}_i^{wv} \mathbf{R}_i + \mathbf{Z}_i^{ww} \right)^{-1} \left(\mathbf{Z}_i^{wv} \mathbf{p}_i + \mathbf{d}_i^w \right). \end{aligned} \quad (3.34)$$

This can be written as $\mathbf{v}_{i+1} = \mathbf{R}_{i+1} \mathbf{w}_{i+1} + \mathbf{p}_{i+1}$, where

$$\mathbf{R}_{i+1} = \left(\mathbf{Z}_i^{vv} \mathbf{R}_i + \mathbf{Z}_i^{vw} \right) \left(\mathbf{Z}_i^{wv} \mathbf{R}_i + \mathbf{Z}_i^{ww} \right)^{-1} \quad (3.35)$$

and

$$\mathbf{p}_{i+1} = \left(\mathbf{Z}_i^{vv} \mathbf{p}_i + \mathbf{d}_i^v \right) - \mathbf{R}_{i+1} \left(\mathbf{Z}_i^{wv} \mathbf{p}_i + \mathbf{d}_i^w \right). \quad (3.36)$$

These are exactly the recurrence relations being sought. The RTMM also requires a backwards solve, which can be accomplished using either Equation (3.32) or (3.33). Horner and Pilkey use the latter, which gives the form

$$\mathbf{w}_i = \mathbf{T}_i \mathbf{w}_{i+1} + \mathbf{q}_i, \quad (3.37)$$

$$\mathbf{T}_i = \left(\mathbf{Z}_i^{wv} \mathbf{R}_i + \mathbf{Z}_i^{ww} \right)^{-1}, \quad (3.38)$$

and

$$\mathbf{q}_i = -\mathbf{T}_i \left(\mathbf{Z}_i^{wv} \mathbf{p}_i + \mathbf{d}_i^w \right). \quad (3.39)$$

Thus, like with the system of equations (3.18), the RTMM solves the problem in sequential steps:

1. Set $\mathbf{R}_1 = \mathbf{0}$ and $\mathbf{p}_1 = \mathbf{0}$
2. Use Equations (3.35), (3.36), (3.38), and (3.39) sequentially from node $i = 1$ to node $i = N$
3. Use the generalized Riccati transformation, $\mathbf{v}_i = \mathbf{R}_i \mathbf{w}_i + \mathbf{p}_i$, along with the right-hand boundary conditions to solve for the unknown elements in \mathbf{v}_{N+1} and \mathbf{w}_{N+1}
4. Use Equation (3.37) sequentially from node $i = N$ to node $i = 1$ to solve for the remaining \mathbf{w}_i and the generalized Riccati transformation to solve for \mathbf{v}_i

Note, however, that because \mathbf{Z}_i is frequency dependent, \mathbf{v}_i and \mathbf{w}_i would normally need to be recalculated for each frequency of interest. For additional speed-up, a block matrix approach can be used to handle the calculation of multiple frequencies at once. For K analysis frequencies, 2×1 vectors are converted to $2K \times 1$ vectors and 2×2 matrices are converted to $2K \times 2K$ block-diagonal matrices. For example, \mathbf{v}_i becomes

$$\begin{bmatrix} \mathbf{v}_i(\omega_1) \\ \mathbf{v}_i(\omega_2) \\ \vdots \\ \mathbf{v}_i(\omega_K) \end{bmatrix},$$

while \mathbf{R}_i becomes

$$\begin{bmatrix} \mathbf{R}_i(\omega_1) & \mathbf{0} & \cdots & \mathbf{0} \\ \mathbf{0} & \mathbf{R}_i(\omega_2) & \cdots & \mathbf{0} \\ \vdots & \vdots & \ddots & \vdots \\ \mathbf{0} & \mathbf{0} & \cdots & \mathbf{R}_i(\omega_K) \end{bmatrix}.$$

A distinct advantage of this block formulation is that analysis can be carried out in almost exactly the same way as the single frequency case. In particular, addition and multiplication operations remain unchanged and the inverse of a matrix like \mathbf{R}_i becomes

$$\begin{bmatrix} \mathbf{R}_i(\omega_1) & \mathbf{0} & \cdots & \mathbf{0} \\ \mathbf{0} & \mathbf{R}_i(\omega_2) & \cdots & \mathbf{0} \\ \vdots & \vdots & \ddots & \vdots \\ \mathbf{0} & \mathbf{0} & \cdots & \mathbf{R}_i(\omega_K) \end{bmatrix}^{-1} = \begin{bmatrix} \mathbf{R}_i^{-1}(\omega_1) & \mathbf{0} & \cdots & \mathbf{0} \\ \mathbf{0} & \mathbf{R}_i^{-1}(\omega_2) & \cdots & \mathbf{0} \\ \vdots & \vdots & \ddots & \vdots \\ \mathbf{0} & \mathbf{0} & \cdots & \mathbf{R}_i^{-1}(\omega_K) \end{bmatrix}.$$

Moreover, such matrices can be efficiently stored using sparse data formats, which also facilitates efficient addition, multiplication, and inversion operations. At the time of writing, the author's work is the first known proposition of this sparse data formulation of the RTMM.

3.4 Procedure

Table 3.1 shows the material and geometric parameters used in the optimization. A routine was written to calculate $J(\mathbf{h})$ given the design variables L_{ABH} , h_0 , and m . A flowchart representation of the routine is shown in Figure 3.4. Given the design input vector, $\mathbf{h} = [L_{ABH} \ h_0 \ m]^T$, the geometry is first recursively divided into a minimum of 50 segments per wavelength, which was shown to produce a relative error less than 1% without significantly increasing calculation time. Figure 3.5 shows a convergence study carried out in preparation for the optimization. The 'true' value for the convergence study was taken to be the spatially-averaged squared velocity response of a representative ABH cantilever beam, calculated using 100 segments per wavelength. The error was calculated as the L^2 norm of the relative difference between the response with fewer segments and the 'true' response. 50 segments per wavelength was chosen as a reasonable compromise between accuracy and runtime.

Once the geometry is segmented, the displacement response, $U(x, \omega)$, due to the point force, $f(x, \omega) = F\delta(x - x_f)$, is calculated using the block RTMM described in

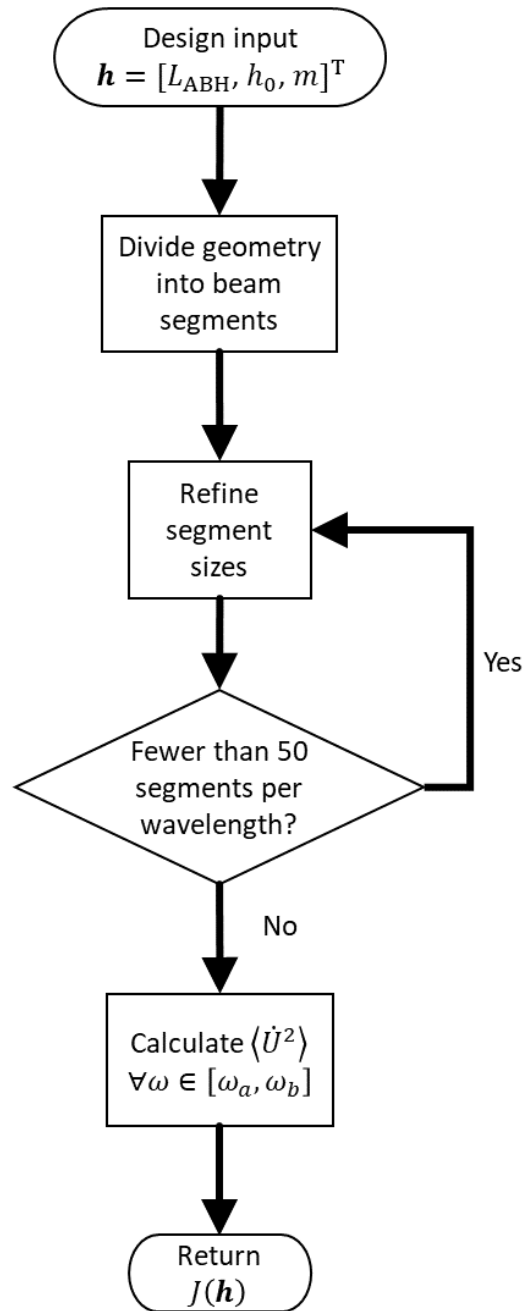


Figure 3.4. Flowchart representation of the design evaluation function used in the optimization.

Table 3.1. Parameters used in the optimization problem.

Material Parameters	
Beam Young's modulus, E	70 GPa
Beam density, ρ	2700 $\frac{\text{kg}}{\text{m}^3}$
Beam loss factor, η	0.0001
Damping Young's modulus, E_d	9 MPa
Damping density, ρ_d	1812 $\frac{\text{kg}}{\text{m}^3}$
Damping loss factor, η_d	0.2
Geometric Parameters	
Beam length, L_x	30 cm
Beam width, L_y	6.35 mm
Beam thickness, h_1	6.35 mm
Damping thickness, h_d	3.81 mm
ABH length limits, $[L^-, L^+]$	[0, 22.26] cm
Minimum thickness limits, $[h^-, h^+]$	[0.635, 6.35] mm
Taper power limits, $[m^-, m^+]$	[2, 12]
Runtime Parameters	
Frequency range, $[\omega_a, \omega_b]$	$2\pi \times [50, 2000]$ Hz
Driving force magnitude, F	1 N
Drive point, x_f	22.26 cm

Section 3.3.2, with the appropriate free-fixed boundary conditions. Finally, the squared velocity response, \dot{U}^2 , is calculated and integrated across x and ω using the trapezoidal rule, and multiplied by the appropriate factors to obtain $J(\mathbf{h})$. Analysis frequencies were chosen to range from 50 Hz to 2,000 Hz with eleven points per 3 dB bandwidth, assuming a structural Q of 10^4 . That is, given any frequency, ω , that lies between ω_i and ω_{i+1} , the span $\omega_{i+1} - \omega_i$ is less than or equal to $11\omega \times 10^{-4}$. This structural Q of 10^4 comes from the modal Q of the reference beam with uniform thickness h_1 having loss factor $\eta = 0.0001$.

3.5 Results and discussion

The routine described in Section 3.4 was used with the Borg multi-objective evolutionary algorithm for 2,000 function evaluations to determine the optimal design. Figure 3.6 shows the progress of the Borg algorithm across these 2,000 function evaluations in terms of the three design variables and the objective value. From this, it is clear that Borg quickly determines that the optimal length, L_{ABH} , is the greatest length allowed and that the optimal minimum thickness, h_0 , is the smallest thickness allowed. The remainder of




Figure 3.5. Convergence study used to determine the number of segments per wavelength to use with the transfer matrix method. The relative error between the ‘true’ response of an example geometry and the response using the given number of segments per wavelength is shown as a solid line. The corresponding runtime of the solver in seconds is given as a dashed line.

Table 3.2. Optimal design variables for the problem described in Section 3.2, and the resultant objective function value.

	L_{ABH}	h_0	m	$J(\mathbf{h})$
Reference beam	0 cm	0 mm	0	1.000
Optimal design	22.26 cm	0.635 mm	3.06	0.2987
Classically optimal design	22.26 cm	0.635 mm	10	0.4866

the search is focused on dialing in the optimal taper power at around 3 or 4. Table 3.2 shows the final, optimal design variables, as well as the value of $J(\mathbf{h})$ for the optimal design compared to the reference beam of uniform thickness. A graphical depiction of the optimal ABH profile is shown in Figure 3.7.

From the results in Table 3.2, it is clear that the inclusion of an ABH reduces the objective function by up to a factor of 3.35 within the frequency band considered. As

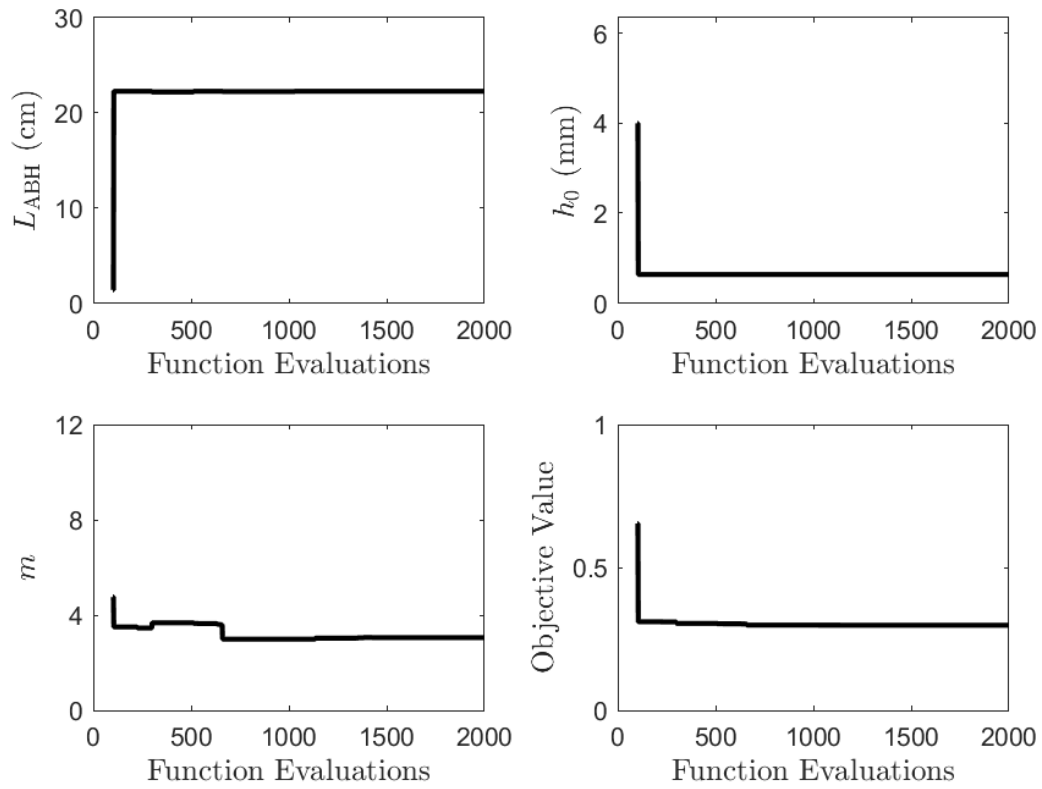


Figure 3.6. Evolution of the Borg algorithm’s archive during the optimization, in terms of the design variables and objective value, $J(\mathbf{h})$, after each function evaluation.




Figure 3.7. Reference beam of uniform thickness (top) and optimal ABH thickness profile (bottom) with added damping layer (black). The left-hand edge of each beam is modelled as free and the right-hand edge is modelled as fixed in the calculation of U . Dimensions are not to scale.




Figure 3.8. Spatially-averaged squared velocity, $\langle \dot{U}^2 \rangle$, of a reference beam with uniform thickness (dashed) and of the optimal ABH thickness profile found by Borg (solid). Note that the scale of the x-axis is logarithmic.

demonstrated by previous authors, a longer taper and smaller minimum thickness is optimal for vibration reduction [7,8]. However, a larger taper power, m , was not found to be optimal, which is contrary to others’ findings. This may be due to the inherently local nature of the optimization formulation. That is, because the frequency range $[\omega_a, \omega_b]$ is fixed, independent of the modal characteristics of each design, some designs will be penalized for having structural modes that fall within the fixed frequency range. This behavior might not be captured when considering only reflection from the ABH as in, e.g., references [7] and [3].

Looking at Figure 3.8—which shows the optimal design’s spatially-averaged squared velocity response, $\langle \dot{U}^2 \rangle$, as a function of frequency—there is clearly a greater number of resonances compared to the reference beam. Note, however, that while the beam with an ABH vibration absorber shows an increased number of resonances, the average value of $\langle \dot{U}^2 \rangle$ is still more than three times lower for the ABH design—a fact that may




Figure 3.9. Spatially-averaged squared velocity, $\langle \dot{U}^2 \rangle$, of a reference beam with uniform thickness (dashed) and of the ‘classically optimal’ ABH thickness profile that produces minimal reflection (solid). Note that the scale of the x-axis is logarithmic.

not be obvious with the logarithmic scale in Figure 3.8. Indeed, previous authors have demonstrated that for increasing taper power, ABH modes move closer together in frequency [33]. Considering that this increased modal density is correlated with increased taper power, there appears to be a point of diminishing returns for m greater than about 3, at which a higher taper power is suboptimal because it results in a greater number of structural modes within the frequency range $[\omega_a, \omega_b]$. Compare, for example, the ‘classically optimal’ design given in Table 3.2 and whose spectrum is shown in Figure 3.9. This design is optimal insofar as it minimizes the average reflection coefficient magnitude in the frequency range $[\omega_a, \omega_b]$. The term ‘classical’ is used here because early ABH theory (including the seminal work of Mironov) viewed the primary benefit of an ABH taper as reducing reflection of bending waves. While it was understood that increased amplitude occurred within the ABH taper, this side effect was not considered in the performance evaluation of ABH design until later, although this increased amplitude can




Figure 3.10. Average reflection coefficient magnitudes, $|R|_{\text{avg}}$, for taper designs having length $L_{\text{ABH}} = 22.26$ cm and minimum thickness $h_0 = 0.635$ mm, but with varying taper power, m . The reflection coefficient was calculated in the same manner as [3]. Note that although the average values are close to one, the reflection coefficient magnitudes were shown to reach as low as 0.6 at particular frequencies.

itself be considered undesirable behavior.

The classically optimal design in Table 3.2 was determined through a parametric study in the taper power, m . Using the impedance matrix method of Georgiev et al. [3], the reflection coefficient was calculated across the frequency range $[\omega_a, \omega_b]$ for each of 19 different taper designs. The results of this parametric study are shown in Figure 3.10. Each design had the same taper length and minimum thickness as the optimal design found by Borg, but their taper powers varied in increments of 0.5. A taper power of ten was found to produce the smallest average reflection coefficient magnitude, at about 0.9205. For reference, the optimal design found by Borg produces an average reflection coefficient magnitude of approximately 0.9247. Note that although the average values are rather large, the reflection coefficient magnitudes were shown to reach as low as 0.6 at particular frequencies.




Figure 3.11. Length of the ABH taper in which Equation (1.7) holds, relative to the total taper length, for the optimal design and the classically optimal design. Equation (1.7) is here considered satisfied when $\left| \frac{1}{k^2} \frac{dk}{dx} \right| \leq 0.3$.

When discussing classical ABH theory, it is important to consider its applicability to the current problem. Consider, for example, the condition of Equation (1.7) from Chapter 1. A benefit of using the particular taper profile $h(x) = \varepsilon x^m + h_0$ is that even for arbitrarily low frequencies or arbitrarily large variations in thickness, there exists a portion of the taper in which Equation (1.7) is satisfied. Even so, it is important to point out that neither the optimal design that minimizes the spatially-averaged squared velocity nor the classically optimal design that minimizes reflection satisfies Equation (1.7) for the entire taper at all frequencies. Figure 3.11 shows as a function of frequency the proportion of the ABH taper in which $\left| \frac{1}{k^2} \frac{dk}{dx} \right| \leq 0.3$. Clearly, there is no frequency for which 100% of the classically optimal taper satisfies Equation (1.7). It may be unfair, then, to call it the classically optimal design, since it does not satisfy completely the assumptions of Mironov’s original analysis. As mentioned above, however, there always exists a portion of the taper in which the assumptions *are* satisfied. Notably, the last

25% of both designs (corresponding to the damped portion) satisfy Equation (1.7) at all frequencies.

Furthermore, as described on page 65, the term ‘classical’ ABH theory as used in this dissertation applies beyond the original analysis of Mironov. The assumption in Equation (1.7), which stems from the validity conditions of the WKB approximation, is not required by the impedance matrix method used by Georgiev et al. [3] and subsequently used here to determine the classically optimal design. Rather, the classical viewpoint of an ABH is that of an isolated feature which admits some incident bending wave and greatly reduces the returned reflected wave. This classically optimal design is therefore optimal insofar as it minimizes the average reflection coefficient magnitude. Yet, as illustrated by the objective values in Table 3.2, the classically optimal design performs worse than that found by Borg when the composite structure (beam with ABH termination) is considered.

This is an important distinction for practical applications. At the time this optimization study was carried out, it was the first to minimize the velocity response of a finite structure with a one-dimensional ABH termination by varying all the ABH parameters. Previous authors had treated the reflection coefficient from the ABH taper as a surrogate measure for the performance of an ABH termination [3, 7, 8]. Although the two measures (velocity response and reflection) are expected to produce the same optimal designs as the analysis frequency bandwidth increases, for finite bandwidths—and especially for lower frequencies—the two measures produce significantly different results. This fact was alluded to by Conlon, Fahline, and Semperlotti, in which the authors found that when the dimension of the ABH is much smaller than the bending wavelength in the base structure, the modal characteristics of the ABH become the determining factors for its performance as a vibration absorber [12]. For reference, below about 1,200 Hz the ratio of L_{ABH} to bending wavelength is less than one. At 50 Hz the ratio is approximately 0.2, which is significantly less than one. Using these ABH modes to modulate the modal characteristics of the composite structure then becomes the optimal strategy, as illustrated by the spectra in Figures 3.8 and 3.9.

Finally, because another benefit of ABH vibration absorbers is a reduction in overall mass, an auxiliary measure, $J_{\text{M}}(\mathbf{h})$, can be defined to compare the masses of different designs. Similarly, a measure of the average kinetic energy of the structure, $J_{\text{KE}}(\mathbf{h})$, can be defined that is dependent on both the mass and the velocity response of the structure. These two measures are defined as

$$J_{\text{M}}(\mathbf{h}) = C_{\text{M}} \int_0^{L_x} \mu(x) \, dx \quad (3.40)$$

Table 3.3. Objective function values for the objective used in the optimization, $J(\mathbf{h})$, and two auxiliary objectives, $J_M(\mathbf{h})$ and $J_{KE}(\mathbf{h})$, measuring overall mass and average kinetic energy, respectively.

	$J(\mathbf{h})$	$J_M(\mathbf{h})$	$J_{KE}(\mathbf{h})$
Reference beam	1.000	1.000	1.000
Optimal design	0.2987	0.5714	0.3982
Classically optimal design	0.4866	0.4676	1.419

and

$$J_{KE}(\mathbf{h}) = \frac{C_{KE}}{L_x} \int_{\omega_a}^{\omega_b} \int_0^{L_x} \mu(x) \left| i\omega U(x, \omega, \mathbf{h}) \right|^2 dx d\omega, \quad (3.41)$$

where $\mu(x) = L_y(\rho h(x) + \rho_d h_d(x))$ is the linear density of the beam and C_M and C_{KE} are normalization constants so that both measures are unity for the reference beam. Note that $h_d(x)$ equals h_d for the thinnest 25% of each ABH and zero otherwise. The numeric values of these measures are given in Table 3.3 for the reference beam, the optimal design found by Borg, and the classically optimal design. Comparing the three measures, it is clear that both the optimal design and the classically optimal design reduce the overall mass and the average velocity response when compared to the reference beam. However, the classically optimal design has a higher average kinetic energy than even the reference beam, due to a greater number of resonances in the analysis frequency band and a greater mobility. These results further underpin the fact that the classically optimal design is in fact suboptimal in the current application.

3.6 Conclusions

The work presented in this chapter sought to expand upon the work in Chapter 2. The optimization framework developed in Chapter 2 was further improved by deriving and implementing a novel version of the transfer matrix method (TMM) that is both faster and more stable than the standard TMM. This improved optimization framework was then applied to the shape optimization of an acoustic black hole (ABH) vibration absorber at the end of a cantilever beam. At the time this optimization study was carried out, it was the first to minimize the velocity response of a finite structure with a one-dimensional ABH termination by varying all the ABH parameters.

As the results of the shape optimization study corroborate, acoustic black hole (ABH) vibration absorbers are highly effective at reducing bending vibrations. However, the inclusion of an ABH vibration absorber also significantly alters the dynamics of the

composite structure. This may be detrimental if trying to control a discrete set of resonances without affecting other frequency regions, but is not necessarily detrimental for broadband vibration reduction, since the overall structural losses also increase with the inclusion of an ABH vibration absorber and the average response therefore decreases.

The results of this study also illustrate an important practical distinction for the design of ABH vibration absorbers. Previous authors have treated the reflection coefficient from an ABH taper as a surrogate measure for its performance, while this study explicitly used the dynamic response of a finite structure. As shown by comparing the optimal design of this study against the design that theoretically minimizes reflection, the two design approaches produce significantly different results. This is a significant finding and one of the key contributions from the work in this chapter. In engineering ABH features for real-world design goals, it may often be more important to consider the way that the ABH modulates the modal characteristics of the composite structure, rather than only considering the reflection characteristics of the ABH taper itself. This is of special importance when considering finite bandwidths and lower frequencies.

Chapter 4 | Optimization and comparison of three styles of ABH*

4.1 Introduction

In much of the literature, the thickness of a beam or plate with an embedded acoustic black hole (ABH) is tapered according to the power law relation $h(x) = \varepsilon x^m + h_0$, with $m \geq 2$ and ε defined by Equation (3.1). In general, as the thickness of a beam or plate decreases, the flexural wave speed also decreases. In the case that the minimum thickness, h_0 , goes to zero in the above power law, the wave's transit time to or from the point of zero thickness goes to infinity and there is theoretically no reflection – thus resulting in perfect absorption. However, as mentioned in Chapter 3, because the ABH effect depends explicitly only on the thickness profile of the beam, there remains an ambiguity as to the precise topology that such an ABH vibration absorber may take. Accordingly, there is a variety of designs that can be found in experimental and numerical work throughout the literature. The current study considers three common ‘styles’ of one-dimensional ABHs to investigate the effect that the choice of style has on performance. These three styles are termed here ‘standard symmetric’, ‘standard nonsymmetric’, and ‘double-leaf’, and are depicted graphically in Figure 4.1 to illustrate their differences. Example mode shapes of nominal designs for each of the three styles are also shown in Figure 4.2 and their modal densities are shown in Figure 4.3. Note that the standard symmetric and standard nonsymmetric styles have similar modal characteristics, while the double-leaf style has a notably higher modal density, as well as strong axial modes that the other two styles do not have.

*A portion of the work presented in this chapter was also published in the Journal of Sound and Vibration [76].



Figure 4.1. Three styles of one-dimensional acoustic black hole vibration absorbers (not to scale). The three styles are (a) Standard symmetric; (b) Standard nonsymmetric; and (c) Double-leaf. Added damping is indicated in black. Note that the colors used for each style in this figure will be used throughout the chapter to indicate each color’s respective style.

Throughout much of the early development of ABH theory, the standard symmetric style was the style explicitly or implicitly used, since the midplane remains parallel along the entire length of the ABH. Often authors illustrate the ABH graphically as the standard nonsymmetric style, which may be considered approximately the same as the standard symmetric style for lower frequencies. Krylov [7] was the first to consider the two styles as explicitly distinct from one another, although his analysis was restricted to differences in the contribution of the damping layers. In particular, Krylov modeled the contribution of damping in the standard symmetric ABH as twice that in the standard nonsymmetric ABH, assuming the same thickness and material properties. The remaining dynamics of the system were considered identical.

Inspired by the work of Bowyer and Krylov [77], Zhou et al. [78] compared the static and dynamic behavior of the standard symmetric and the double-leaf (the authors use the term ‘double-layered’) ABH styles using the finite element method (FEM). Results of their analysis showed that the double-leaf ABH had a consistently higher modal loss factor and more localized energy within the ABH portion of the beam than a comparably-designed

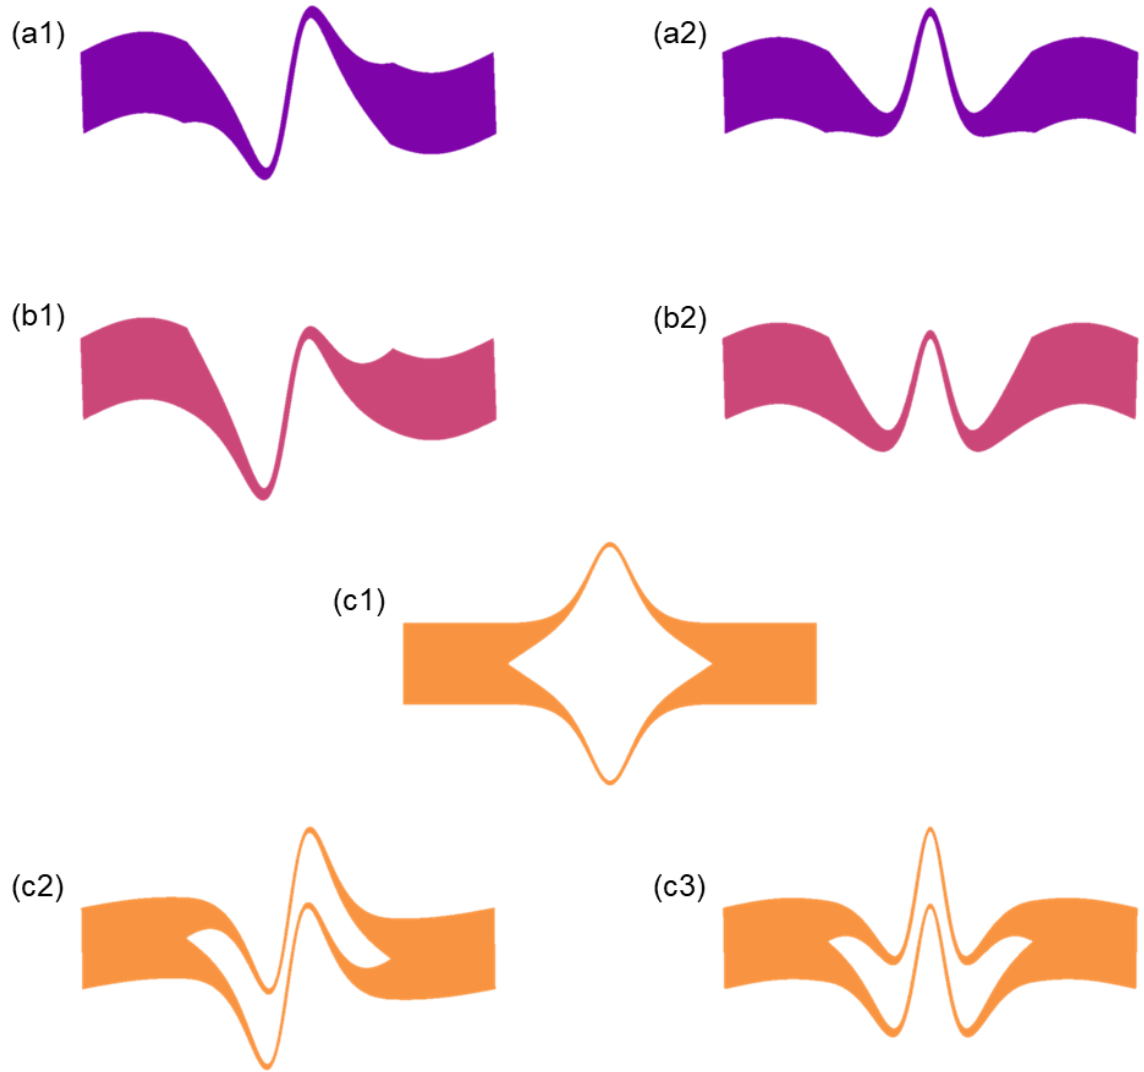


Figure 4.2. Example mode shapes for nominal ABH designs (not to scale). Modal frequencies are (a1) 405.8 Hz and (a2) 551.5 Hz for the standard symmetric style; (b1) 405.6 Hz and (b2) 551.7 Hz for the standard nonsymmetric style; and (c1) 43.5 Hz, (c2) 232.4 Hz, and (c3) 373.9 Hz for the double-leaf style.

standard symmetric ABH. As the authors note, the two halves of the double-leaf ABH act in parallel as two nonsymmetric ABHs, each with half the thickness of the symmetric ABH. A thinner taper and thicker relative damping are both understood to enhance the ABH effect, and so this ‘splitting’ effect of the double-leaf ABH can account for much of its improved performance over the standard symmetric ABH.

In addition to the above work, there is a plethora of studies in the literature concerned specifically with one-dimensional ABHs in beams. The bulk of these studies have focused




Figure 4.3. Modal densities for the nominal designs shown in Figure 4.2. (●) Standard symmetric. (■) Standard nonsymmetric. (▼) Double-leaf, axially symmetric modes. (▲) Double-leaf, axially anti-symmetric modes.

on examining the dynamics of various ABH designs to gain some insight into how ABH design effects performance [13, 79–81], with a subset corroborating their theoretical analyses through experiment [33, 82–84]. Some have investigated nonlinear effects in ABHs to improve their performance [85, 86], while others have looked at attaching piezoelectric patches to ABH features and tailoring them to improve energy harvesting [87–89] and vibration reduction [90, 91]. A number of authors have proposed numerical methods aimed at modeling ABHs in beams [65, 92–94], and a few authors have modeled how ABHs in beams might affect radiated sound [94, 95]. The ‘double-leaf’ style of ABH seems to be particularly popular as a candidate for one-dimensional metamaterial beams [96–99]. With the rise of additive manufacturing, a number of authors have investigated the additional capabilities and complexities that come from 3D printed beams with ABHs [100–102]. One group has even proposed and studied a spiral ABH termination [103–105]. There are some previous studies in the literature that have addressed the optimization of a

one-dimensional ABH feature [24, 26, 65, 106]. The current work builds upon previous findings by applying a rigorous multi-objective optimization scheme to each of the three above styles—‘standard symmetric’, ‘standard nonsymmetric’, and ‘double-leaf’. At the time of writing, a direct comparison of these three styles has not been carried out except for the author’s work, and certainly not in the context of optimal ABH design.

4.2 Problem description

For each style, we consider the problem of a one-dimensional ABH vibration absorber embedded in a thin beam of length L_x , having pinned boundary conditions at each end. The beam is excited at a location, x_f , by a harmonic force acting normal to the midplane and distributed across the beam’s width. The ABH shape parameters, L_{ABH} , h_0 , and m , which determine the beam’s profile, $h(x)$, may vary between some lower and upper bounds. Additionally, the location of the center of the ABH, x_{ABH} , may vary along the latter half of the beam, opposite the force location. The percentage of the ABH taper, P_d , that includes an additional free damping layer of thickness h_d and density ρ_d , may also vary. Each possible design can be uniquely defined by the vector $\mathbf{h} = [L_{\text{ABH}} \ h_0 \ m \ x_{\text{ABH}} \ P_d]^T$. The two objective functions used for this optimization are

$$J_1(\mathbf{h}) = C_1 \left[\rho_d h_d P_d L_{\text{ABH}} + \int_0^{L_x} \rho h(x) \, dx \right] \quad (4.1)$$

and

$$J_2(\mathbf{h}) = \frac{C_2}{L_x} \int_{\omega_a}^{\omega_b} \int_0^{L_x} |\dot{u}_z(x, \omega)|^2 \, dx d\omega = C_2 \int_{\omega_a}^{\omega_b} \langle \dot{u}_z^2 \rangle \, d\omega, \quad (4.2)$$

where \dot{u}_z is the complex velocity of the midplane in the vertical direction, $\langle \dot{u}_z^2 \rangle$ is the spatially-averaged square velocity, and C_1 and C_2 are chosen so that $J_1 = J_2 = 1$ for a uniform beam without added damping—i.e., when $\mathbf{h} = \mathbf{0}$. J_1 is proportional to the mass of the entire system, with the first term in the sum representing the mass of the damping layer and the second representing the mass of the beam with ABH. J_2 is proportional to the squared velocity response in (ω_a, ω_b) , averaged over the length of the beam. If it is assumed that the (complex) displacement of the beam’s midplane has the form $u_z(x, t) = u_z(x)e^{i\omega t}$, then the velocity is $\dot{u}_z(x, t) = i\omega u_z(x, t)$ with its magnitude depending on ω . The objectives having been defined, the optimization problem may be formulated as

$$\min_{\mathbf{h} \in \mathcal{H}} \mathbf{J}(\mathbf{h}) = [J_1(\mathbf{h}) \ J_2(\mathbf{h})]^T$$

subject to

$$\begin{aligned} \mathcal{H} = & \left\{ \mathbf{h} \in \mathbb{R}^5 : L^- \leq L_{\text{ABH}} \leq L^+, h^- \leq h_0 \leq h^+, \right. \\ & \left. m^- \leq m \leq m^+, x^- \leq x_{\text{ABH}} \leq x^+, P^- \leq P_d \leq P^+ \right\} \\ & -\omega^2 \rho \mathbf{u} - \text{div}(\boldsymbol{\sigma}(\mathbf{u})) = \mathbf{f}, \quad \forall \mathbf{x} \in \Omega \\ & \left. \begin{aligned} \mathbf{f} &= \left[0 \quad 0 \quad F \delta(x - x_f) \delta\left(z - \frac{1}{2}h(x_f)\right) \right]^T \\ \mathbf{u}\left(0, y, \frac{1}{2}h(0)\right) &= \mathbf{u}\left(L_x, y, \frac{1}{2}h(L_x)\right) = \mathbf{0} \\ u_y &= 0 \end{aligned} \right\} \forall \omega \in [\omega_a, \omega_b] \end{aligned}$$

In words, the above states that the goal is to simultaneously minimize J_1 and J_2 , where the possible designs, \mathbf{h} , are restricted to exist in the set \mathcal{H} , which effectively defines the bounds of the search space. The dynamic response of the structure obeys steady-state linear elastodynamics for all points in the domain of the beam, Ω . It should be noted that beam is modelled with complex material parameters (see Table 4.1), and so the displacement, \mathbf{u} , will generally be complex. The constraint on u_y means the elastodynamic problem considered follows plane strain assumptions, while the constraints on \mathbf{u} at $x = 0$ and $x = L_x$ indicate pinned boundary conditions. The definition of \mathbf{f} indicates a distributed force at $\mathbf{x} = \left[x_f \quad y \quad \frac{1}{2}h(x_f) \right]^T$, $\forall y \in \Omega$, which is the intersection of the beam's midplane and the plane defined by $x = x_f$, thus defining a line segment.

4.3 Procedure

An automatic meshing function was developed to take as input the design variables detailed in Section 4.2 and produce a geometric mesh and force and boundary conditions suitable for analysis using the finite element method (FEM). Again, FEM is more appropriate for this study than the transfer matrix method (TMM) because TMM does not incorporate modelling of the second dimension. The Euler-Bernoulli dynamic beam equation, as a one-dimensional equation, assumes that the thickness at a given point is distributed symmetrically along the midplane. However, most graphical depictions of ABH tapers show a nonsymmetric geometry as in Figure 4.1(b). This is also a common way in which physical ABHs are manufactured for experimental studies. In order to properly model this variation along the second dimension, an alternative method to TMM must be used. FEM is a standard method used in dynamic mechanical analysis and so it has been used here.

The meshing function meshes the beam using linear hexahedral elements, with two




Figure 4.4. Example of how a given design might be divided for the purposes of meshing. The leftmost region will be the same for all designs, as its purpose is to guarantee nodes exist at $x = 0$ for the constraints and nodes exist at $x = x_f$ for the applied forces.

elements across the beam's width, two elements through the beam's thickness, and at least twelve elements per wavelength along the beam's length, as determined by the frequency range of interest. Twelve elements per wavelength was chosen based on previous analysis of a one-dimensional ABH [25] and a mesh convergence study (see Appendix A), which indicated that when a one-dimensional ABH is meshed with twelve elements per wavelength, comparable results are achieved using either linear elements or quadratic elements. The ability to use linear elements, therefore, greatly reduces computation time.

In order to ensure at least twelve elements per wavelength, the beam is meshed piecewise in the following way. First, the beam is divided into three or four regions, depending on the value of x_{ABH} . One of these regions is the ABH taper, while the remaining regions are the uniform base beam with thickness h_1 . The leftmost region is always the same and is necessary to ensure nodes exist at x_f . An approximate example of how the beam might be divided is shown graphically in Figure 4.4. For a portion of the uniform beam with length L_u , the number of segments needed can be simply calculated in advance by noting that the bending wavelength is

$$\lambda = 2\pi \left(\frac{Eh_1^2}{12\omega^2\rho(1-\nu^2)} \right)^{\frac{1}{4}}. \quad (4.3)$$

The number of segments required is thus $N = \lceil 12L_u\lambda^{-1} \rceil$, and the length of each segment is L_u/N . In the case of the ABH taper, the calculation is not as simple. However, the same general approach still applies, except that instead of dealing with an entire length

at once, one considers a summation of differential lengths

$$\frac{dx}{\lambda(x)}. \quad (4.4)$$

To get the number of segments needed to adequately mesh a portion of the ABH from $x = a$ to $x = b$, the calculation is

$$N = \left\lceil 12 \int_a^b \frac{dx}{\lambda(x)} \right\rceil = \left\lceil \frac{12}{2\pi} \int_a^b \left(\frac{12\omega^2\rho(1-\nu^2)}{Eh^2(x)} \right)^{\frac{1}{4}} dx \right\rceil. \quad (4.5)$$

In order to mesh the ABH taper with sufficient resolution, but without unnecessary refinement, the taper region is meshed recursively. For a given portion of the taper, N is calculated by Equation (4.5). If that number is greater than one, then the segment is split into two equal segments and the process is repeated for each of those segments. A flowchart of the process is shown in Figure 4.5.

The added damping was modeled explicitly in the analysis, rather than incorporated as an additional loss factor as was done in Chapter 3. This explicit approach was chosen for several reasons. First, any increase in bending stiffness, though expected to be small, is accounted for and can be non-negligible in the minimum thickness region of the ABH taper. Second, the effect of mass loading is accounted for. This is particularly significant, as the density of the damping material is, in this case and in many practical applications, on the same order as that of the beam. Again, in the minimum thickness region of the taper, the effect of mass loading can be significant, even for thin damping layers. It should be noted that several authors [3, 33, 65, 85, 86, 94, 107, 108] have followed the work of Kyrlov [7] in using an analytical model for constrained layer damping developed by Ross, Kerwin, and Ungar [28], which accounts for changes in effective bending stiffness and surface mass density. However, by explicitly modeling the free damping layer, separately from the beam, these two effects of increased bending stiffness and mass loading are included in a more holistic and theoretically more accurate way.

Table 4.1 gives the range of design variables considered for the optimization problem described, as well as material parameters of the beam and damping layer. ω_a was chosen to be the critical frequency of the unmodified beam, while ω_b was found to have minimal response when a random sample of design variables was evaluated. The drive location, x_f , was chosen so as to excite all relevant modes within the frequency range of interest. For each ABH style, J_1 and J_2 were evaluated 6,000 times using the Borg MOEA to determine the set of designs that approximate the Pareto optimal set. This number of function

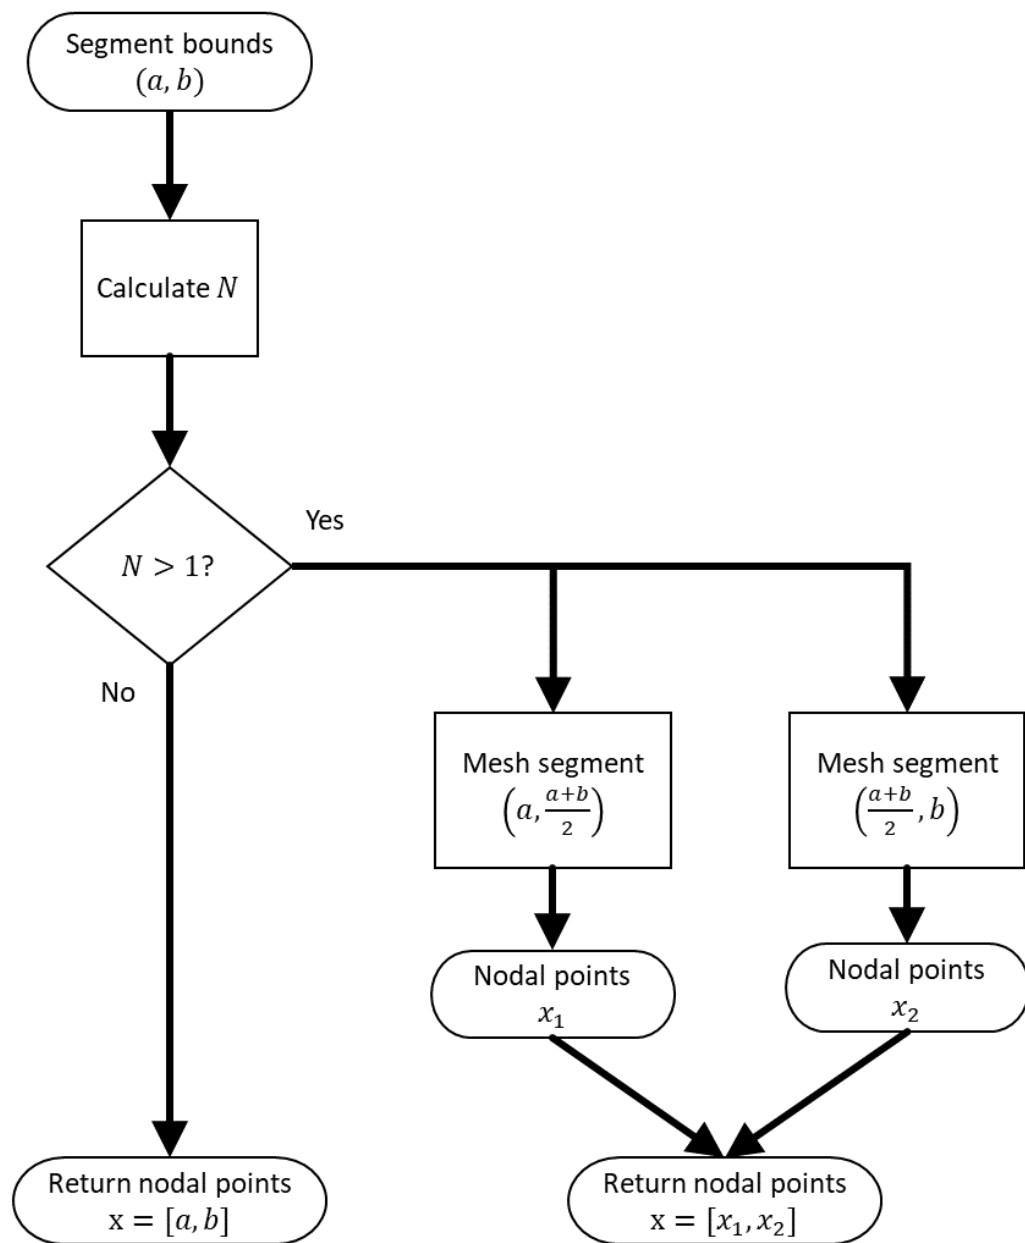


Figure 4.5. Flowchart representation of the recursive function used to mesh the ABH region. N is calculated according to Equation (4.5).

Table 4.1. Variable and parameter bounds of the design problem and geometric and material properties of the system. *For the standard nonsymmetric style, $h_d = 3$ mm.

Material Parameters	
Beam Young's modulus, E	70 GPa
Beam density, ρ	2,700 kg/m ³
Beam Poisson's ratio, ν	0.35
Beam loss factor, η	1×10^{-4}
Damping Young's modulus, E_d	9 MPa
Damping density, ρ_d	1,812 kg/m ³
Damping Poisson's ratio, ν_d	0.45
Damping loss factor, η_d	0.2
Geometric Parameters	
Beam length, L_x	61 cm
Beam thickness, h_1	6.35 mm
Damping length, L_d	$P_d L_{ABH}$
Damping thickness, h_d	1.50 mm*
ABH length limits, $[L^-, L^+]$	[1, 15] cm
Minimum thickness limits, $[h^-, h^+]$	[0.635, 6.35] mm
Taper power limits, $[m^-, m^+]$	[2, 8]
ABH position limits, $[x^-, x^+]$	[30.5, 61] cm
Percent damping limits, $[P^-, P^+]$	[0, 1]
Runtime Parameters	
Frequency range, $[\omega_a, \omega_b]$	$2\pi \times [2000, 7000]$ Hz
Driving force magnitude, F	1 N
Drive point, x_f	2.17857 cm

evaluations was chosen because the Pareto set appeared not to change significantly at this point in all three cases.

The above automatic meshing function was used with the commercial FEM solver NASTRAN to solve for the complex velocity of the midplane at a given frequency. Given a design vector, \mathbf{h} , the beam geometry is meshed by calculating the positions of nodes and the elements to which these nodes are assigned. Because the number of elements across the width and through the thickness remains unchanged, it is only a matter of determining the number of elements along the length, as described earlier and shown in Figure 4.5. The nodes and elements so defined, they, along with the constraints, forces, and analysis frequencies, are written in the appropriate format to be read by NASTRAN as .dat files. NASTRAN is then called with these .dat files as input, and upon completion it writes the output as a .pch file. In order to speed up analysis, NASTRAN is run using its distributed memory parallelization (DMP) feature, which

allows each processor to calculate the response at a frequency unique to that processor. Once complete, the `.pch` output file is next read to extract the response of the midplane in the vertical direction, u_z , from which the spatially-averaged square velocity, $\langle \dot{u}_z^2 \rangle$, is derived for each analysis frequency. J_1 is then calculated analytically according to Equation (4.1), while J_2 is calculated from Equation (4.2) using the trapezoidal rule for integration. A flowchart of the whole process is given in Figure 4.6. The function used to calculate J_1 and J_2 from the input design variables was linked to the Borg multiobjective evolutionary algorithm (MOEA) to find an optimal solution set for the optimization problem described in Section 4.2.

The analysis frequencies between ω_a and ω_b are determined through recursive piecewise linear interpolation of $\langle \dot{u}_z^2 \rangle$, such that the relative error at the midpoints of the line segments is less than or equal to 10^{-2} . It is precisely the same philosophy as the meshing function shown in Figure 4.5, where now (a, b) defines a frequency domain, and instead of calculating N , the relative error between a linear interpolation and the actual value is calculated. In the case this error is greater than 10^{-2} , the frequency domain is recursively divided. In fact, this recursive division is done in parallel by first dividing the domain $[\omega_a, \omega_b]$ according to the number of available processors, so as to make use of NASTRAN’s DMP capabilities described above.

4.4 Results and discussion

When the objective function values of the Pareto optimal set are plotted to form the Pareto fronts shown in Figure 4.7, they illustrate the inherent trade-off between the two objectives. The Pareto fronts of the three styles show a similar trade-off, indicating that the mechanism responsible for the trade-off between objectives is similar, if not the same. It is also worth noting that there exists a threshold above which there is no trade-off and the presence of a thoughtfully designed ABH vibration absorber is advantageous to both objectives. For reference, Table 4.2 details the designs at either end of the Pareto fronts, while Figure 4.8 shows graphical depictions of the designs that minimize J_2 (‘Min Response’ in Table 4.2).

Looking closely at Table 4.2, a few things become evident. First, the optimal set of values for L_{ABH} , h_0 , and m favor a longer ABH with a smaller minimum thickness and more extreme taper power for all three styles. This trend has been observed by previous authors in the context of minimizing reflection from a taper at the end of a beam or plate with free boundary conditions [3, 7, 8, 32, 109]. However, recall that the

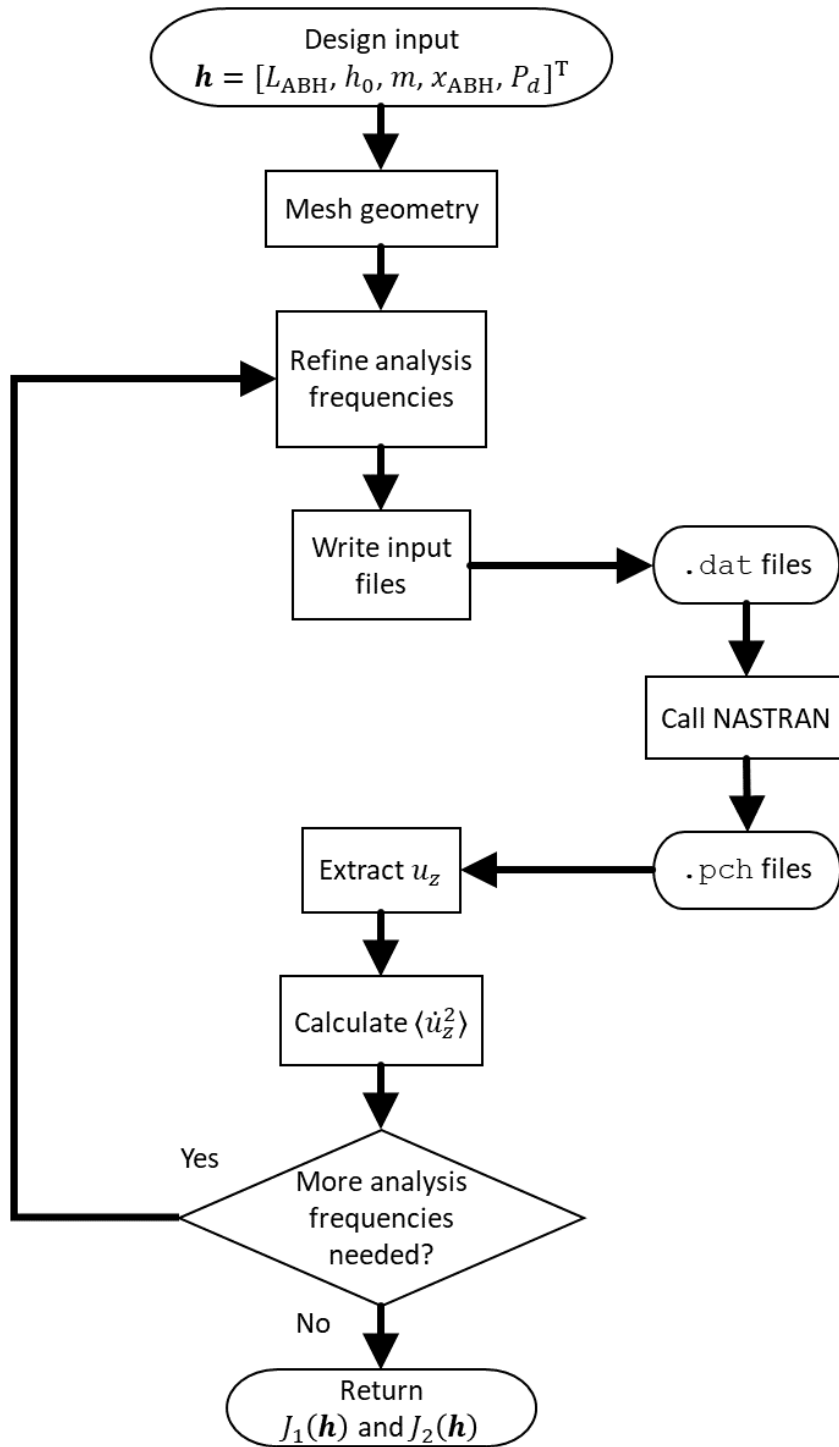


Figure 4.6. Flowchart representation of the design evaluation function used by Borg in the optimization process.

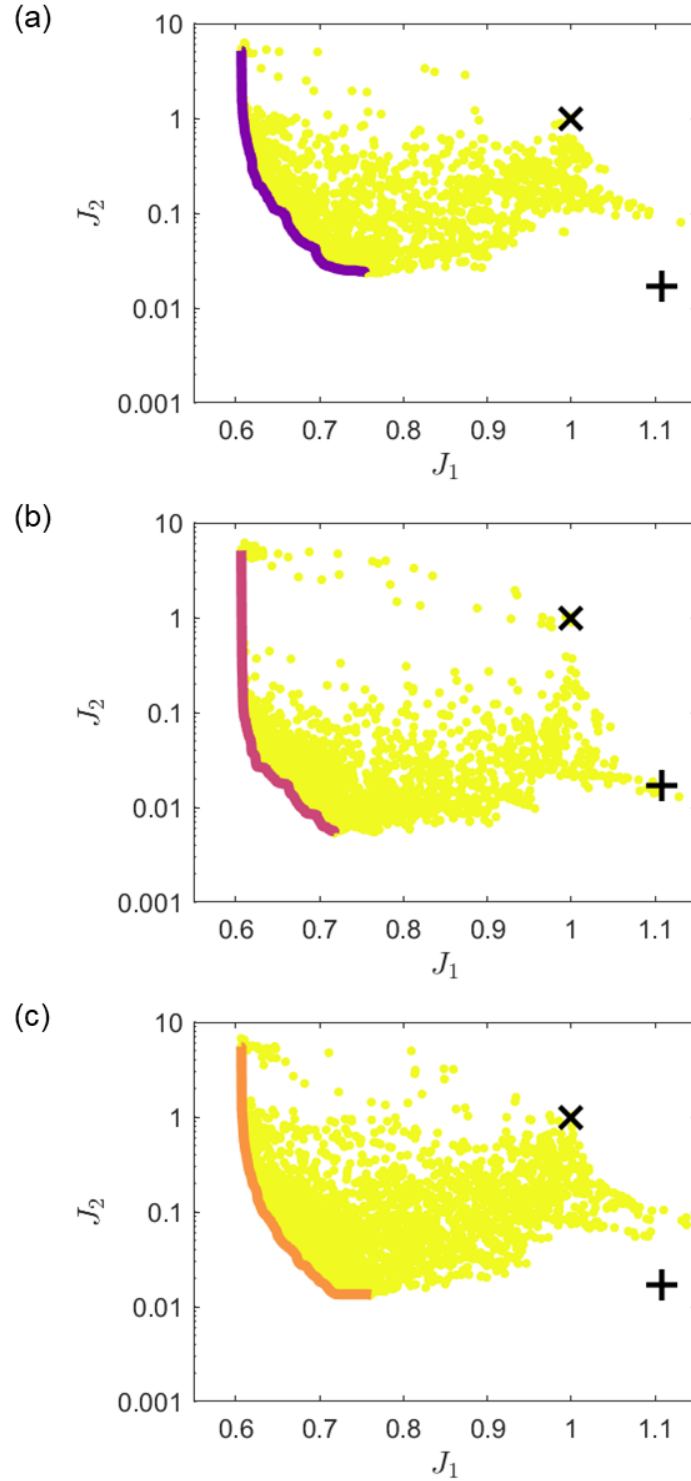


Figure 4.7. Approximate Pareto fronts for two objectives and three styles. (a) Standard symmetric. (b) Standard nonsymmetric. (c) Double-leaf. The objective values are also shown for an unmodified beam (x) and for an unmodified beam with the maximum damping allowed within the variable bounds (+).

Table 4.2. Designs from the three Pareto fronts in Figure 4.7 that minimize J_1 ('Min Mass') or minimize J_2 ('Min Response') for each of the three styles.

	Min Mass	Min Response
Standard Symmetric	$L_{ABH} = 15$ cm $h_0 = 0.635$ mm $m = 7.99$ $x_{ABH} = 35.2$ cm $L_d = 0$ cm	$L_{ABH} = 14.2$ cm $h_0 = 0.636$ mm $m = 7.82$ $x_{ABH} = 33.1$ cm $L_d = 12.1$ cm
Standard Nonsymmetric	$L_{ABH} = 15$ cm $h_0 = 0.635$ mm $m = 8$ $x_{ABH} = 34.1$ cm $L_d = 0$ cm	$L_{ABH} = 15$ cm $h_0 = 0.686$ mm $m = 7.98$ $x_{ABH} = 34.7$ cm $L_d = 10.3$ cm
Double-Leaf	$L_{ABH} = 15$ cm $h_0 = 0.635$ mm $m = 8$ $x_{ABH} = 34.7$ cm $L_d = 0$ cm	$L_{ABH} = 15$ cm $h_0 = 0.635$ mm $m = 8$ $x_{ABH} = 34.7$ cm $L_d = 14.9$ cm




Figure 4.8. Graphical depiction (not to scale) of the design of each style that minimizes J_2 ('Min Response' in Table 4.2). (a) Standard symmetric; (b) Standard nonsymmetric; and (c) Double-leaf. Also depicted is (d) An unmodified beam with the maximum damping allowed within the variable bounds (represented by a + symbol in Figure 4.7).

results of Chapter 3 indicated that optimal ABH designs in the two contexts do not always coincide. Often, when an ABH is embedded in a finite, rather than semi-infinite, structure, the performance of the ABH vibration absorber is dependent upon its discrete modal behavior [12, 110–112]. This apparent contradiction can be resolved by considering the scale of the ABH and the frequency bandwidth. At 2,000 Hz, the bending wavelength in the beam is approximately 18 cm, which is comparable to the size of the ABH. The ABH therefore begins to act more like a broadband absorber within the structure. It is

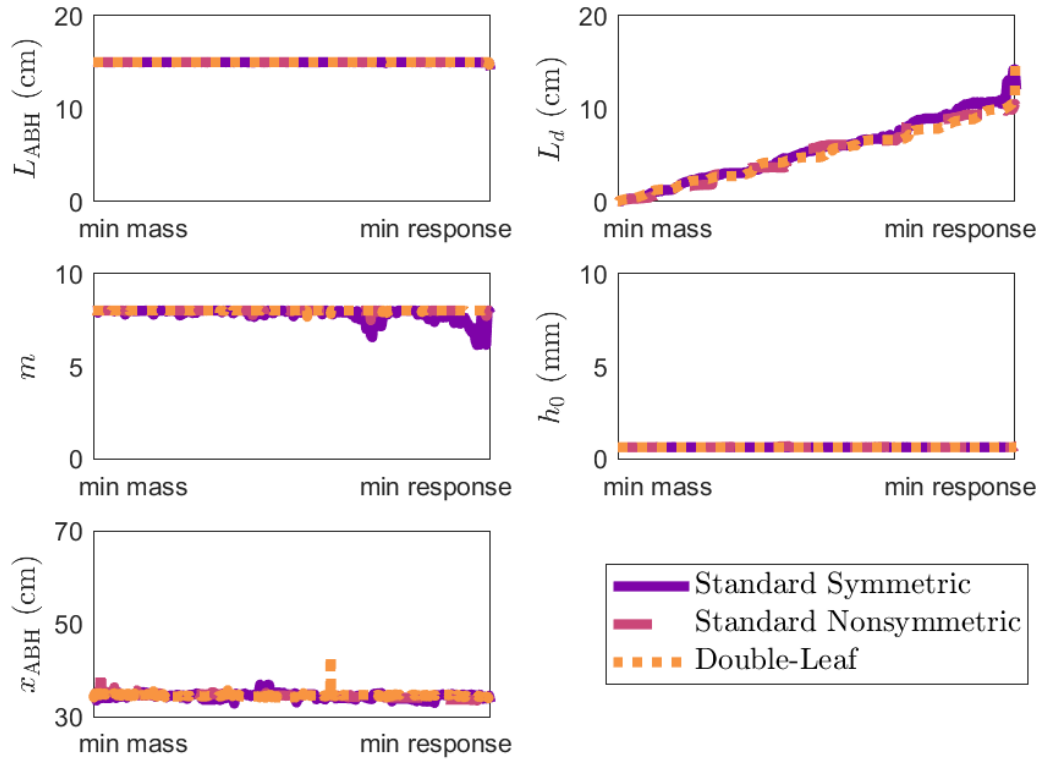


Figure 4.9. Trends for each design variable along the three Pareto fronts in Figure 4.7, from the design that minimizes J_1 ('min mass') to the design that minimizes J_2 ('min response').

worth noting that the three designs do not satisfy the condition of Equation (1.7) from Chapter 1 at all points for all frequencies. However, as mentioned in Chapter 3, a benefit of using the particular taper profile $h(x) = \varepsilon x^m + h_0$ is that there always exists a portion of the taper in which Equation (1.7) is satisfied. Moreover, the results of Chapter 3 demonstrated that designs that do not totally satisfy Equation (1.7) can still outperform designs that do.

The second thing evident from Table 4.2 is that the trade-off captured by the Pareto front is dominated by the effect of the damping layer, as seen in the variable trends of Figure 4.9. For the range of design variables considered, the trade-off between objectives is a function only of the length of damping, $L_d = P_d L_{ABH}$, for any given style of ABH. It is therefore worthwhile to consider the effect of damping in further detail. Previous authors [3, 27, 113] have noted that only a minimal increase in performance is achieved by applying damping to more than the thinnest 40% of the taper area. Figure 4.10 shows the dependence of J_2 on the amount of damping for the set of optimal designs.




Figure 4.10. Dependence of J_2 on damping along the Pareto front. The dashed lines represent the length of damping, L_d , required to achieve 99% of the possible reduction in J_2 for each ABH style.

Also plotted is the amount of damping at which 99% of the maximum reduction in J_2 is achieved. In line with [3,27,113], it is clear that only a fraction of the taper area needs to be damped to achieve the majority of the benefits of an ABH vibration absorber. For the standard nonsymmetric style with double-thickness damping, less than 7% of the taper requires added damping to achieve 99% of the vibration response reduction.

Third, not only do the three Pareto fronts share a similar shape, but the variable trends are even more closely matched. This indicates that the differences between the three Pareto fronts are a result of the choice of ABH style alone and that this choice has a significant effect on performance of the ABH vibration absorber. Vibration response spectra are shown in Figure 4.11 for the three ABH styles considered. For each style, the response spectrum shown is that of the design that minimizes J_2 (‘Min Response’ in Table 4.2). Because the design variables for each design are nearly the same, comparing the spectra can help clarify the effect due to the choice of ABH style. Comparing the

spectra of the standard symmetric and standard nonsymmetric styles, the spectra look very similar except that the nonsymmetric response is more heavily damped. This is in line with the analysis of Krylov [7], which found that for $E_d \ll E$ and reasonable damping layer thicknesses, the energy absorption from two layers of equal thickness will be less than the absorption from a single layer with twice the thickness. Conversely, the response spectrum of the double-leaf style shows a higher modal density, in line with the observation of Zhou et al. [78] that the two halves act in parallel as two tapers of half thickness each. This also explains why the double-leaf spectrum is more heavily damped than the standard symmetric spectrum.

4.5 Additional comparison

While the results of the optimization clearly indicate that the choice of ABH style has a significant effect on performance of the vibration absorber, it is worthwhile to also consider the effects that the choice of style has on other practical aspects. As such, several additional analyses were carried out in order to highlight further differences between the styles and to more holistically characterize the decision space for which style is best suited to a particular application.

4.5.1 Static compliance

The design for each style that minimizes J_2 ('Min Response' in Table 4.2) was subjected to two types of static forces. In the first case, a 1 N compressive force was applied in the axial direction at either end of each design. Note that the pinned boundary conditions were removed for this case, or else there would be no motion. In the second case, a 1 Nm torque was applied at each end so as to induce a negative curvature in the beam. That is, a positive torque was applied at $x = 0$ and a negative torque was applied at $x = L_x$. In both cases, the total absolute elemental strain (TAES) of the midplane was calculated as a measure of the resultant deformation. If the length of the n^{th} segment of the midplane is ℓ_n when no force is applied, and the length of this segment changes to ℓ'_n as a result of the force, then the strain is calculated as

$$s = \frac{\sum |\ell'_n - \ell_n|}{L_x}. \quad (4.6)$$

The results of these calculations are given in Table 4.3, including the corresponding strain values for an unmodified beam as reference. These values can be better understood

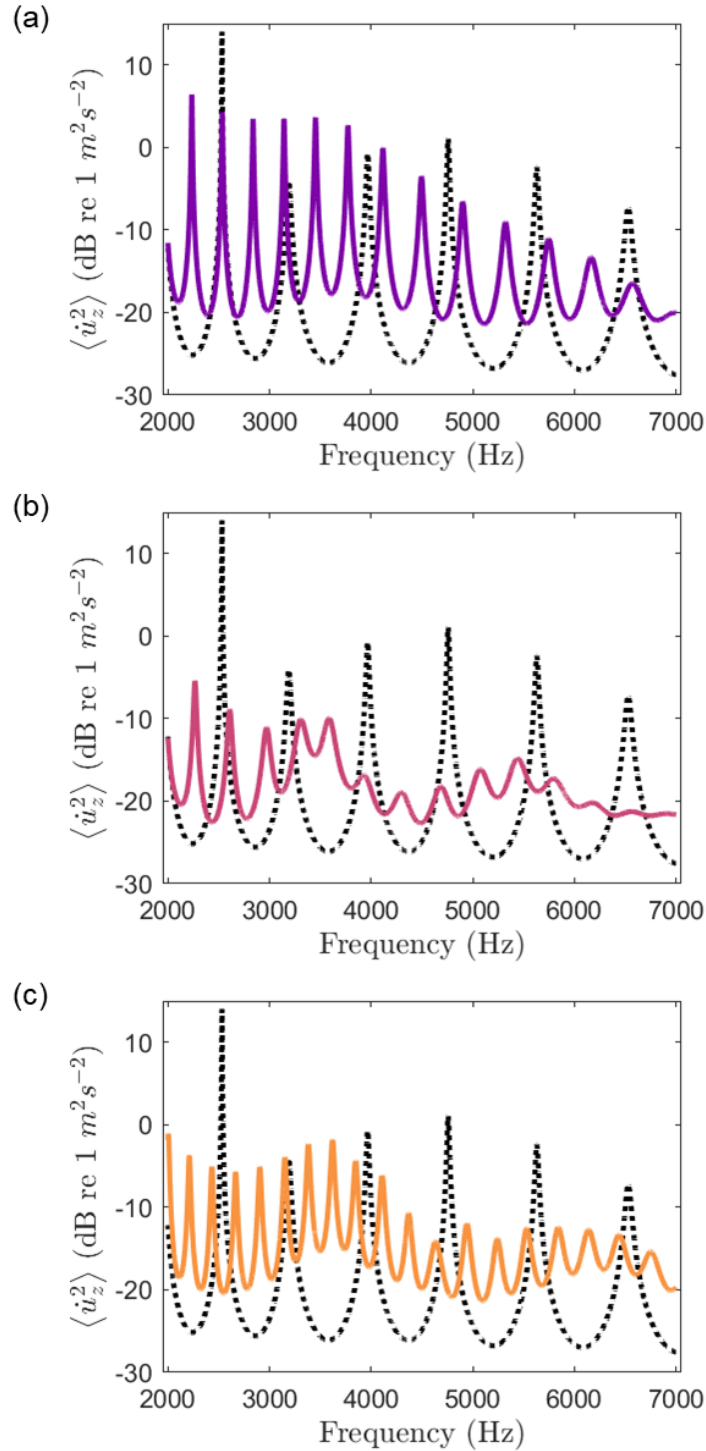


Figure 4.11. Spatially-averaged square velocity spectra for the design of each style that minimizes J_2 ('Min Response' in Table 4.2). (a) Standard symmetric. (b) Standard nonsymmetric. (c) Double-leaf. Also shown is the spectrum for an unmodified beam with the maximum damping allowed within the variable bounds (dashed).

Table 4.3. Total absolute elemental strain values calculated by Equation (4.6) as a result of an axial compression force and a torque applied at the ends of the beam. Note that these are different from normal strain and shear strain. The designs are those that minimize J_2 (‘Min Response’ in Table 4.2), as well as an unmodified beam for reference.

	s (compression)	s (torque)
Standard Symmetric	3.4×10^{-5}	8.0×10^{-2}
Standard Nonsymmetric	9.4×10^{-2}	2.2×10^{-5}
Double-Leaf	3.2×10^{-5}	2.3×10^{-5}
Unmodified	8.3×10^{-6}	4.3×10^{-7}

by looking at the deformations shown in Figure 4.12. In the case of compression, the standard symmetric style shows minimal deformation. Because the force is precisely axially symmetric, the only functional difference between the standard symmetric beam and the unmodified beam is that there is less material in the ABH region, thereby making it slightly more compliant. If the compression force were not perfectly parallel to the beam midplane, it is expected that deformations such as those of the standard nonsymmetric and double-leaf beams would occur. In each case, the power law taper of the ABH transforms the axial force and generates rotation at the edge of the ABH. Because the thickness of the beam is significantly decreased within the ABH, the bending stiffness is greatly decreased and double damage is done as the beam buckles. However, this effect is largely mitigated in the double-leaf design by its symmetry towards the top and bottom of the beam, which greatly increases its compression stiffness. Despite the buckling at the ABH, the TAES due to the compression force is comparable to that of the standard symmetric design.

In the case of an applied torque, the roles are reversed somewhat. Whereas the TAES in the case of compression is on the order of 10^{-5} for the standard symmetric design and on the order of 10^{-2} for the standard nonsymmetric design, the TAES due to an applied torque is on the order of 10^{-2} for the standard symmetric design and on the order of 10^{-5} for the standard nonsymmetric design. This again has to do with the distribution of the stresses, which are concentrated in the minimum thickness region of the standard symmetric ABH, while being more distributed in the standard nonsymmetric design, in which the minimum thickness region is located at the bottom of the beam. The double-leaf design does an even better job by having minimum thickness regions at both the top and bottom of the beam, increasing the overall bending stiffness. It is interesting to note that the double-leaf design appears to buckle slightly at the center of the ABH. However, this does not appear to affect its performance, according to Table 4.6.

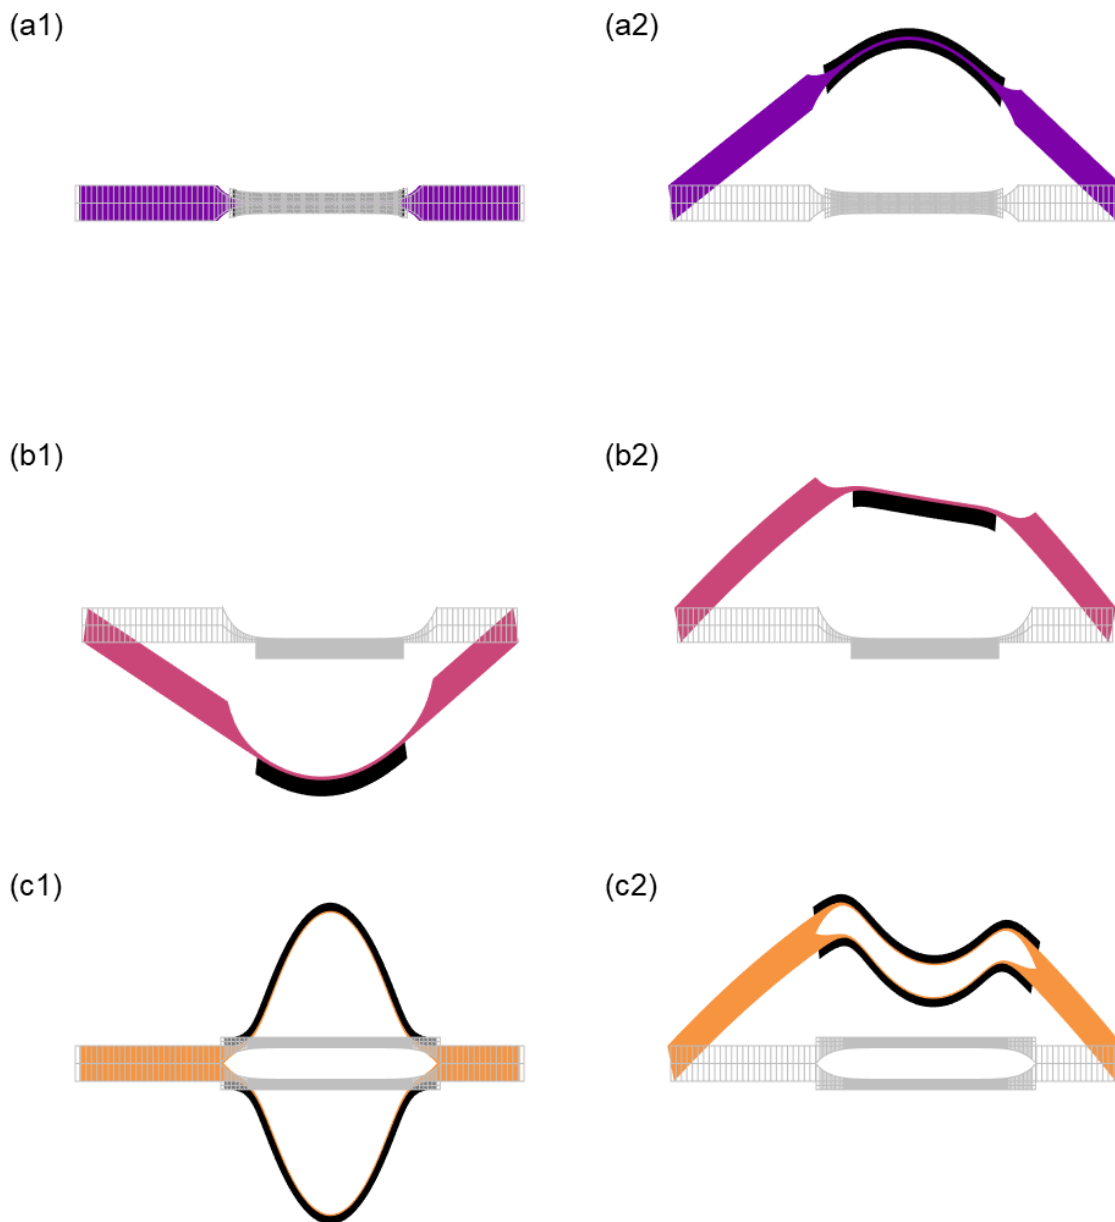


Figure 4.12. Deformation (not to scale) corresponding to the values in Table 4.3. Shown is deformation due to axial compression for (a1) Standard symmetric, (b1) Standard nonsymmetric, and (c1) Double-leaf designs, as well as deformation due to torque for (a2) Standard symmetric, (b2) Standard nonsymmetric, and (c2) Double-leaf designs. Skeletons of the undeformed geometries are also shown in light grey.

4.5.2 Radiated sound power

As another comparative measure, an approximation of the radiated sound power per unit width was calculated from the velocity data represented in Figure 4.11. An FFT implementation of the Rayleigh integral was used to calculate the far-field pressure as a function of angle, assuming that each beam design is baffled [114]. In order to increase resolution for later integration in k -space, the velocity response of the midplane was padded with zeros beyond the end of the beam. The pressure values at each angle was used to calculate the far-field intensity, assuming $p(\mathbf{r}) = \rho c \dot{u}_r(\mathbf{r})$, and the intensity as a function of angle was integrated over a hemispherical surface by dividing it into quadrilateral patches and taking the mean of the four vertices as the value for a given patch. The sound power was calculated in this way as a function of frequency, and this sound power spectrum was integrated according to one-third octave bins to get the total radiated sound power in each bin. The results are shown in Figure 4.13.

In almost all one-third octaves, the designs with an ABH radiate less sound power than an unmodified beam. This is to be expected, since the wavespeed throughout the length of the modified beam is supersonic for the entire frequency range. Conversely, the wavespeed at the minimum thickness portion of the ABH is subsonic for the entire frequency range. This effectively decouples the structure from the surrounding fluid and reduces far-field radiation. An exception appears to exist in the case of the standard symmetric design, which radiates greater sound power in the 2000 Hz one-third octave band and comparable sound power in the 3150 Hz one-third octave band. This can be explained by looking at the spectra in Figure 4.11. The upper bound of the 2000 Hz one-third octave is about 2250 Hz, which is below the modal frequency of the unmodified beam at about 2500 Hz. Conversely, there are multiple peaks in the standard symmetric spectrum, which also has a generally greater spatially-averaged square velocity in this part of the frequency range than at other frequencies. Even accounting for this exception, the standard nonsymmetric and double-leaf styles radiate notably less sound power than standard symmetric style—sometimes close to 10 dB less. With respect to each other, however, the standard nonsymmetric style and the double-leaf style perform comparably, with one style radiating less sound in one one-third octave band and the other style radiating less in another.




Figure 4.13. Radiated sound power spectrum in one-third octaves for each of the designs that minimize J_2 (‘Min Response’ in Table 4.2). Also shown is the one-third octave radiated sound power for an unmodified beam with the maximum damping allowed within the variable bounds.

4.5.3 Heterogeneous ABH parameters

Finally, a parametric study was also performed on the standard nonsymmetric ABH to illustrate the effect of a one-dimensional ABH with two different power-law tapers. The design that minimizes J_2 (‘Min Response’ in Table 4.2) was selected, and the taper power, m , of each half of the ABH was independently varied from 2 to 8 in steps of 0.5. The other design variables remained unchanged. J_2 was calculated for each unique design, the results of which are shown in Figure 4.14. As might be expected, the change in J_2 is approximately symmetric, with J_2 minimized for both taper powers equal to that of the original design. However, there appears to be a slight asymmetry, with the left taper having a greater effect on J_2 . This is likely due to the position of the ABH along the beam being just right of center, which means that the overall structure itself is asymmetric.

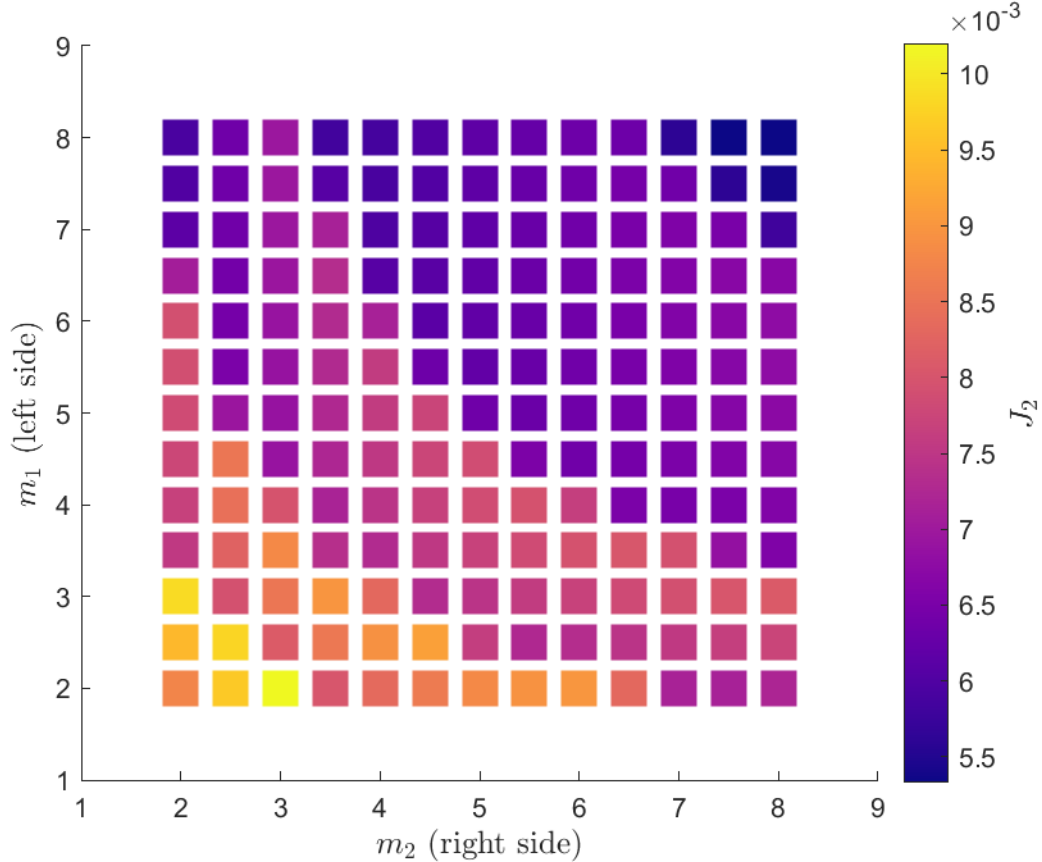


Figure 4.14. Approximate contour plot of J_2 as a function of the two taper powers for each half of the standard nonsymmetric ABH style. m_1 is the taper power of the left half of the ABH. m_2 is the taper power of the right half of the ABH. The remaining design variables are those under ‘Min Response’ in Table 4.2.

4.6 Conclusions

A formal optimization was carried out for three styles of one-dimensional ABH vibration absorbers, each embedded in a beam that is subject to a harmonic force. A multi-objective evolutionary optimization scheme was used to identify the set of ABH designs that optimally minimize the beam’s vibration response and its overall mass. The three styles show a similar trade-off between the two objectives, which for any given style is a function only of the amount of added damping. This indicates that the differences in trade-off are a result of the choice of ABH style alone and that this choice has a significant effect on performance of the ABH vibration absorber. Moreover, while the double-leaf style has preferable stiffness properties for static forces, the results of this work show that the standard nonsymmetric style is superior to either of the other styles

in terms of vibration and sound reduction. This fact is significant, given that this work is the first and only to directly compare the three ABH styles.

It should be noted that this study used a frequency-independent proportional loss factor in modeling both the beam and damping materials, as is common in the structural dynamics literature. In practice, both loss factor and Young’s modulus can vary substantially with frequency—especially in viscoelastic polymers such as those often employed for vibration damping. Had the loss factor and Young’s modulus been modeled as such, the results detailed in Section 4.4 would likely have differed somewhat. In particular, a frequency-dependent loss factor could make it advantageous to shift structural modes towards regions where damping is relatively high, thereby potentially altering the set of optimal designs and the shape of the Pareto fronts. However, the authors expect that the high-level results of this study would remain functionally the same. That is, for the design space considered the amount of added damping dominates the trade-off between the two objectives and any differences in trade-off are related to the choice of ABH style alone.

Indeed, the results of this study indicate that careful design of the damping layer can be important to ABH performance. A few authors have looked specifically at the design of damping layers in ABH vibration absorbers since this work was reported [34, 36, 115, 116], although as in this study they all use a frequency-independent loss factor and Young’s modulus. It is also worth noting that the results in Section 4.4 indicate there exists a threshold above which there is no trade-off and the presence of a thoughtfully designed ABH vibration absorber is advantageous to both objectives. That is to say, the inclusion of a thoughtfully designed ABH in the beam reduces the beam’s vibration response, even with minimal damping material. Furthermore, among the three styles, the vibration responses of the ‘standard symmetric’ and ‘standard nonsymmetric’ ABHs are very similar except that the nonsymmetric response is more heavily damped, while the response of the ‘double-leaf’ ABH shows a higher modal density. In practical applications, it may be important to consider other elements of design when implementing an ABH vibration absorber. Additional analysis was therefore conducted to characterize the optimal ABH designs in ways independent of their vibration reduction properties. It was demonstrated that an ABH design that is optimized for vibration reduction may perform poorly compared to other designs and other styles under these additional considerations. This further underpins the significance of directly comparing these three styles.

Chapter 5 | A generalized measure of ABH performance*

5.1 Introduction

The formulation of structural optimization problems requires a great deal of effort and forethought. This was covered in some detail in Chapter 1, but it bears repeating. The aim when developing some measure of optimality is to distill the relevant physics into a representation that is simple to manipulate, but also accurately represents the desired behavior. Although the spatially-averaged square velocity response has been used extensively in this dissertation as a measure of performance, there are many options, including point mobility and modal quantities. Choosing or designing an appropriate performance measure is a non-trivial task.

Furthermore, as demonstrated in each of the previous chapters, the optimal shape of the structure will depend not only on the measure of performance but perhaps just as much on the scale of the problem, i.e., the size of the structure and its corresponding dynamic response. In Chapter 2 use of discrete frequencies led to an optimal strategy that shifted peaks to between the analysis frequencies. Chapter 3 showed that the classical ABH approach of minimizing reflection can be suboptimal when considering a finite structure within a desired frequency band. Indeed, the ‘classically optimal’ design was shown to be suboptimal precisely because it introduced additional peaks within the analysis frequency band. This was not the case in Chapter 4, in part because great care was taken to tailor the frequency range such that designs would not be penalized for peaks that straddle the frequency window. It is impractical to put such effort into

*A portion of the work presented in this chapter has also been submitted for publication in JASA Express Letters.

choosing the frequency range each time an ABH vibration absorber consists of a different scale or different material properties.

As such, it is advantageous to be able to approach the problem of ABH design in a manner that is independent of the structure's size and material. There should be a criterion by which one can directly compare the performance of a 1 m ABH made out of steel to that of a 10 cm ABH manufactured out of aluminum. Additionally, it was demonstrated in Chapter 4 that the application of damping has a significant control over the performance of an ABH vibration absorber. It begs the question, then, whether one can design a measure of performance that incorporates the damping directly. In particular, is there a way to quantify the energy coupling between the ABH taper and the added damping?

The primary goal of the current chapter is to develop a generally applicable measure of ABH performance in finite structures and apply it to a test problem. To that end, a dissipated power ratio is derived that is directly related to the coupling between ABH and damping but is largely independent of scale. Rather than setting out to solve a particular design problem, this work aims to establish a foundation upon which future studies can be based and/or from which ABH vibration absorber performance can be given a universal language. As a test case, the one-dimensional ABH optimization covered in Chapters 3 and 4 is expanded to a two-dimensional ABH vibration absorber in a square plate. Most practical applications of ABH vibration absorbers involve two-dimensional ABHs, so this optimization problem is an important extension of the previous work. The objective of the optimization is to simultaneously maximize the new power-based ABH performance measure and to minimize the size of the ABH feature. To facilitate the structural analysis of each ABH design in three space dimensions, without a great loss of accuracy or great increase in analysis time, an auxiliary goal of the work is to derive and build an improved finite element solver. As such, the development of this solver is also documented in this chapter.

5.2 Problem description

The test problem under consideration is the shape optimization of a two-dimensional acoustic black hole (ABH) vibration absorber. The ABH is embedded in a square plate of side length a , and it has an attached free damping layer that covers the entire bottom of the ABH. A graphical representation of the geometry is shown in Figure 5.1. The results of Chapter 4 showed that of the three styles considered—standard symmetric, standard




Figure 5.1. Graphical representation of the optimization problem under consideration. Depicted are (a) Top; (b) Bottom; and (c) Cross-sectional views of the geometry. The geometry is a square plate of side length a and thickness h_1 , with an embedded two-dimensional ABH feature and an attached free damping layer. The damping layer covers the same area of the ABH, which is a circle with radius R_{ABH} . The taper profile of the ABH is controlled by the minimum thickness, h_0 , and the power law exponent, m . Not pictured is the parameter γ , which alters the gradient of the ABH.

nonsymmetric, and double-leaf—the standard nonsymmetric style was preferable to the other two styles in many ways. As such, it is the only style considered here. The ABH thickness profile is independent of angle, but its thickness varies with radius, r , as $h(r) = \varepsilon(r + \beta)^m + \gamma$. As derived in Appendix B, ε and β can be calculated as

$$\beta = \frac{R_{ABH}(h_0 - \gamma)^{\frac{1}{m}}}{(h_1 - \gamma)^{\frac{1}{m}} - (h_0 - \gamma)^{\frac{1}{m}}} \quad (5.1)$$

and

$$\varepsilon = \frac{h_1 - \gamma}{(R_{ABH} + \beta)^m}, \quad (5.2)$$

where R_{ABH} is the radius of the circular ABH, h_1 is the thickness of the host plate, h_0 is the thickness at the center of the ABH, and γ is a fourth parameter that roughly controls the gradient at the center of the ABH. There are therefore four variables that uniquely define the shape of the ABH. A vector, \mathbf{h} , of nondimensional design variables can be defined such that $\mathbf{h} = \left[2\frac{R_{\text{ABH}}}{a} \quad \sqrt{\frac{h_1}{h_0}} \quad m \quad \frac{\gamma}{h_0} \right]^T = [v_1 \quad v_2 \quad v_3 \quad v_4]^T$. This set of nondimensional variables still uniquely defines the ABH geometry, but has the added benefit that it is independent of the scale of the problem. In this way, the nondimensional variables describe a continuum of ABH designs, differentiated only by the dimensions of the host structure. The first variable, v_1 , is the ratio of the ABH radius to one half the edge length of the plate. The second variable, v_2 , is equivalent to the ratio of the bending wavenumber at the minimum thickness to the wavenumber at the maximum thickness, assuming the plate material is homogeneous. The third variable, v_3 , is the taper power, m , while the fourth variable, v_4 , is such that when $v_4 = 1$, the gradient at the center of the ABH is zero.

Given the nondimensionalized design variables, we now define two dimensionless objectives related to the performance of the ABH vibration absorber. Namely,

$$J_1(\mathbf{h}) = 4\frac{R_{\text{ABH}}^2}{a^2} = v_1^2 \quad (5.3)$$

and

$$J_2(\mathbf{h}) = \left(\frac{1}{(ka)_b - (ka)_a} \int_{(ka)_a}^{(ka)_b} \Pi(ka) \, d(ka) \right)^{-1}, \quad (5.4)$$

where

$$\Pi(ka) = -\frac{2}{\omega \text{Im}\{u_z(\mathbf{x}_f, ka)\}} \int_{\Omega_d} \langle \dot{D}(\mathbf{x}, ka) \rangle \, d\Omega_d \quad (5.5)$$

is called here the ‘power dissipation ratio’, and is the ratio of power dissipated in the damping layer (defined by the domain Ω_d) to the power input to the system at the drive point, \mathbf{x}_f . $\langle \dot{D} \rangle$ is the rate of dissipated energy averaged over one period and u_z is the (complex) displacement in the z -direction. The details of how Equation (5.5) is derived will be given in Section 5.3. Note that the ka used here and throughout the paper refers to the bending wavenumber in the host structure. To be properly scale-independent, J_2 should also be normalized by the volume of Ω_d . However, because the dimension of this volume is captured by J_1 , this normalization is omitted. J_1 is simply the square of v_1 , which is related to the volume of Ω_d , while J_2 measures directly the effectiveness of the vibration absorber to dissipate energy. Note the reciprocal in the definition of J_2 , so that maximum power is dissipated when J_2 is minimized.

With the two objectives so defined, the problem can be written formally as

$$\begin{aligned}
& \min_{\mathbf{h} \in \mathcal{H}} \mathbf{J} = [J_1, J_2]^T \\
& \text{subject to} \\
& \mathcal{H} = \left\{ \mathbf{h} \in \mathbb{R}^4 : v_1^- \leq v_1 \leq v_1^+, v_2^- \leq v_2 \leq v_2^+, v_3^- \leq v_3 \leq v_3^+, v_4^- \leq v_4 \leq v_4^+ \right\} \\
& \quad - (ka)^4 \mathbf{u} - \frac{3a^4}{h_1^2} \frac{\lambda + 2\mu}{\mu(\lambda + \mu)} \operatorname{div}(\boldsymbol{\sigma}(\mathbf{u})) = \mathbf{f}, \quad \forall \mathbf{x} \in \Omega \\
& \quad \left. \begin{aligned} & \mathbf{u}(\pm 0.5, y, z) = \mathbf{u}(x, \pm 0.5, z) = \mathbf{0} \\ & \mathbf{f} = \frac{3a^4}{h_1^2} \frac{\lambda + 2\mu}{\mu(\lambda + \mu)} \begin{bmatrix} 0 & 0 & \delta(\mathbf{x} - \mathbf{x}_f) \end{bmatrix}^T \end{aligned} \right\} \forall ka \in [(ka)_a, (ka)_b]
\end{aligned}$$

In words, the above says that the goal of the optimization is to simultaneously minimize $J_1(\mathbf{h})$ and $J_2(\mathbf{h})$, where \mathbf{h} is restricted to exist in the set of valid designs, \mathcal{H} . \mathcal{H} is defined by the permissible lower and upper bounds of each design variable. The dynamics of the problem must obey time-harmonic linear elastodynamics, written here in terms of ka , the Lamé parameters, λ and μ , and the (complex) stress, $\boldsymbol{\sigma}$. The edges of the plate will be fixed, and the plate is driven in the z -direction by a point force at some appropriate location, \mathbf{x}_f .

5.3 Theory

5.3.1 Elastodynamic theory

While it will not be derived here, the equation of motion for linear elastodynamics is given as

$$\rho \ddot{\mathbf{u}} - \operatorname{div}(\boldsymbol{\sigma}) = \mathbf{f} \quad (5.6)$$

or, equivalently,

$$\rho \ddot{u}_i - \sum_j \sigma_{ij,j} = f_i, \quad (5.7)$$

where \mathbf{u} is the three-dimensional vector of displacement components, $\ddot{\mathbf{u}}$ is its second derivative with respect to time, $\boldsymbol{\sigma}$ is the second-order stress tensor, $\operatorname{div}(\cdot)$ is the divergence operator, and \mathbf{f} is the three-dimensional vector of external forces *per unit volume*. The ‘,’ notation in Equation (5.7) is that commonly encountered in continuum mechanics, e.g.,

$$u_{1,1} = \frac{\partial u_x}{\partial x},$$

$$\sigma_{13,2} = \frac{\partial \sigma_{xz}}{\partial y},$$

and

$$f_{2,13} = \frac{\partial^2 f_y}{\partial z \partial x}.$$

The stress tensor is defined as the double contraction of the fourth-order elasticity tensor, \mathbf{C} , and the second-order strain tensor, $\boldsymbol{\epsilon}$,

$$\boldsymbol{\sigma} = \mathbf{C} : \boldsymbol{\epsilon} \quad (5.8)$$

or

$$\sigma_{ij} = \sum_k \sum_l C_{ijkl} \epsilon_{kl}. \quad (5.9)$$

The strain tensor itself is defined as the symmetric gradient of the displacement,

$$\boldsymbol{\epsilon} = \frac{1}{2} (\text{grad}(\mathbf{u}) + \text{grad}(\mathbf{u})^T) \quad (5.10)$$

or

$$\epsilon_{ij} = \frac{1}{2} (u_{i,j} + u_{j,i}), \quad (5.11)$$

where $\text{grad}(\cdot)$ is the gradient operator. In the case of an isotropic material, $C_{ijkl} = \lambda \delta_{ij} \delta_{kl} + \mu (\delta_{ik} \delta_{jl} + \delta_{il} \delta_{jk})$, where λ and μ are the first and second Lamé parameters, respectively. Under such conditions, the stress tensor can be written as

$$\boldsymbol{\sigma} = \lambda \text{div}(\mathbf{u}) \mathbf{I} + \mu (\text{grad}(\mathbf{u}) + \text{grad}(\mathbf{u})^T), \quad (5.12)$$

where \mathbf{I} is the second-order identity tensor, functionally equivalent to the identity matrix. Because the stress and strain tensors are defined in terms of \mathbf{u} , they will be denoted hereafter as operations on \mathbf{u} ; that is, $\boldsymbol{\sigma}(\mathbf{u})$ and $\boldsymbol{\epsilon}(\mathbf{u})$, respectively.

For the steady-state solution of Equation (5.6), a complex time-harmonic solution is assumed such that

$$\mathbf{u} \rightarrow (\mathbf{v} + i\mathbf{w}) e^{i\omega t}, \quad (5.13)$$

$$\mathbf{f} \rightarrow (\mathbf{f} + i\mathbf{0}) e^{i\omega t}, \quad (5.14)$$

$$\lambda \rightarrow \lambda(1 + i\eta), \quad (5.15)$$

$$\mu \rightarrow \mu(1 + i\eta), \quad (5.16)$$

where η is the material loss factor. The following boundary conditions will also be imposed:

$$\mathbf{u} = \mathbf{0}, \quad \forall \mathbf{x} \in \Gamma_1, \text{ and} \quad (5.17)$$

$$\boldsymbol{\sigma}(\mathbf{u})\mathbf{n} = \mathbf{0}, \quad \forall \mathbf{x} \in \Gamma_2, \quad (5.18)$$

for the outward unit normal vector \mathbf{n} . Γ_1 is the boundary surface of the domain on which displacement is zero, while Γ_2 is the boundary surface on which surface pressures (tractions) are zero. The union of Γ_1 and Γ_2 forms the entire boundary of the domain, Ω . More formally, $\partial\Omega = \Gamma_1 \cup \Gamma_2 \equiv \Gamma$. These boundary conditions, together with the complex time-harmonic assumption, give the following system of equations:

$$-\omega^2 \rho \mathbf{v} - \operatorname{div}(\boldsymbol{\sigma}(\mathbf{v}) - \eta \boldsymbol{\sigma}(\mathbf{w})) = \mathbf{f} \quad \forall \mathbf{x} \in \Omega, \quad (5.19)$$

$$-\omega^2 \rho \mathbf{w} - \operatorname{div}(\eta \boldsymbol{\sigma}(\mathbf{v}) + \boldsymbol{\sigma}(\mathbf{w})) = \mathbf{0} \quad \forall \mathbf{x} \in \Omega, \quad (5.20)$$

$$\mathbf{v} = \mathbf{0} \quad \forall \mathbf{x} \in \Gamma_1, \quad (5.21)$$

$$\mathbf{w} = \mathbf{0} \quad \forall \mathbf{x} \in \Gamma_1, \quad (5.22)$$

$$(\boldsymbol{\sigma}(\mathbf{v}) - \eta \boldsymbol{\sigma}(\mathbf{w}))\mathbf{n} = \mathbf{0} \quad \forall \mathbf{x} \in \Gamma_2, \quad (5.23)$$

$$(\eta \boldsymbol{\sigma}(\mathbf{v}) + \boldsymbol{\sigma}(\mathbf{w}))\mathbf{n} = \mathbf{0} \quad \forall \mathbf{x} \in \Gamma_2. \quad (5.24)$$

This system of equations is the strong form, from which the weak and Galerkin forms will be derived for use in the finite element method.

5.3.2 FEM formulation and preconditioning

By expanding the formalism of linear algebra to include matrices with operator elements, Equations (5.19) and (5.20) can be written as

$$\left(\begin{bmatrix} -\omega^2 \rho & 0 \\ 0 & -\omega^2 \rho \end{bmatrix} + \begin{bmatrix} -\operatorname{div}(\boldsymbol{\sigma}(\cdot)) & \operatorname{div}(\eta \boldsymbol{\sigma}(\cdot)) \\ -\operatorname{div}(\eta \boldsymbol{\sigma}(\cdot)) & -\operatorname{div}(\boldsymbol{\sigma}(\cdot)) \end{bmatrix} \right) \begin{bmatrix} \mathbf{v} \\ \mathbf{w} \end{bmatrix} = \begin{bmatrix} \mathbf{f} \\ \mathbf{0} \end{bmatrix}. \quad (5.25)$$

This makes derivation of the weak form straightforward, so that it can be written as

$$\int_{\Omega} \begin{bmatrix} \boldsymbol{\phi}_i \\ \boldsymbol{\psi}_i \end{bmatrix}^T \begin{bmatrix} -\omega^2 \rho & 0 \\ 0 & -\omega^2 \rho \end{bmatrix} \begin{bmatrix} \mathbf{v} \\ \mathbf{w} \end{bmatrix} d\Omega$$

$$+ \int_{\Omega} \begin{bmatrix} \phi_i \\ \psi_i \end{bmatrix}^T \begin{bmatrix} -\operatorname{div}(\boldsymbol{\sigma}(\cdot)) & \operatorname{div}(\eta\boldsymbol{\sigma}(\cdot)) \\ -\operatorname{div}(\eta\boldsymbol{\sigma}(\cdot)) & -\operatorname{div}(\boldsymbol{\sigma}(\cdot)) \end{bmatrix} \begin{bmatrix} \mathbf{v} \\ \mathbf{w} \end{bmatrix} d\Omega = \int_{\Omega} \begin{bmatrix} \phi_i \\ \psi_i \end{bmatrix}^T \begin{bmatrix} \mathbf{f} \\ \mathbf{0} \end{bmatrix} d\Omega, \quad (5.26)$$

where $\phi_i(\mathbf{x})$ is a test function for the possible values of $\mathbf{v}(\mathbf{x})$ and $\psi_i(\mathbf{x})$ is the corresponding test function for $\mathbf{w}(\mathbf{x})$. The reason for the index of i in ϕ_i and ψ_i will be explained below when deriving the Galerkin form. For now, expanding Equation 5.26 gives the equivalent equation

$$\begin{aligned} & \int_{\Omega} \phi_i \cdot (-\omega^2 \rho \mathbf{v}) d\Omega + \int_{\Omega} \psi_i \cdot (-\omega^2 \rho \mathbf{w}) d\Omega \\ & - \int_{\Omega} \phi_i \cdot \operatorname{div}(\boldsymbol{\sigma}(\mathbf{v})) d\Omega + \int_{\Omega} \phi_i \cdot \operatorname{div}(\eta\boldsymbol{\sigma}(\mathbf{w})) d\Omega \\ & - \int_{\Omega} \psi_i \cdot \operatorname{div}(\eta\boldsymbol{\sigma}(\mathbf{v})) d\Omega - \int_{\Omega} \psi_i \cdot \operatorname{div}(\boldsymbol{\sigma}(\mathbf{w})) d\Omega = \int_{\Omega} \phi_i \cdot \mathbf{f} d\Omega. \end{aligned} \quad (5.27)$$

Without derivation, the following identity is given for some first-order tensor[†] field $\mathbf{b}(\mathbf{x})$ and second-order tensor field $\mathbf{A}(\mathbf{x})$:

$$\operatorname{div}(\mathbf{A}\mathbf{b}) = \mathbf{b} \cdot \operatorname{div}(\mathbf{A}^T) + \operatorname{tr}(\mathbf{A}\operatorname{grad}(\mathbf{b})), \quad (5.28)$$

where $\operatorname{tr}(\mathbf{A}) = \sum_i A_{ii}$. Rearranging and integration produces the relevant Green's identity

$$\begin{aligned} \mathbf{b} \cdot \operatorname{div}(\mathbf{A}) &= \operatorname{div}(\mathbf{A}^T \mathbf{b}) - \operatorname{tr}(\mathbf{A}^T \operatorname{grad}(\mathbf{b})) \\ \implies \mathbf{b} \cdot \operatorname{div}(\mathbf{A}) &= \operatorname{div}(\mathbf{A}^T \mathbf{b}) - \operatorname{tr}(\operatorname{grad}(\mathbf{b})^T \mathbf{A}) \\ \implies \mathbf{b} \cdot \operatorname{div}(\mathbf{A}) &= \operatorname{div}(\mathbf{A}^T \mathbf{b}) - \operatorname{grad}(\mathbf{b}) : \mathbf{A} \\ \implies \int_{\Omega} \mathbf{b} \cdot \operatorname{div}(\mathbf{A}) d\Omega &= \int_{\Omega} \operatorname{div}(\mathbf{A}^T \mathbf{b}) d\Omega - \int_{\Omega} \operatorname{grad}(\mathbf{b}) : \mathbf{A} d\Omega \\ \implies \int_{\Omega} \mathbf{b} \cdot \operatorname{div}(\mathbf{A}) d\Omega &= \int_{\Gamma} \mathbf{A}^T \mathbf{b} \cdot \mathbf{n} d\Gamma - \int_{\Omega} \operatorname{grad}(\mathbf{b}) : \mathbf{A} d\Omega \\ \implies \int_{\Omega} \mathbf{b} \cdot \operatorname{div}(\mathbf{A}) d\Omega &= \int_{\Gamma} \mathbf{b} \cdot \mathbf{A}\mathbf{n} d\Gamma - \int_{\Omega} \operatorname{grad}(\mathbf{b}) : \mathbf{A} d\Omega, \end{aligned} \quad (5.29)$$

where $\mathbf{B} : \mathbf{A} = \sum_i \sum_j B_{ij} A_{ij}$ is the double contraction of two second-order tensors. The weak form can then be written

$$\begin{aligned} & -\omega^2 \int_{\Omega} \rho(\phi_i \cdot \mathbf{v}) d\Omega - \omega^2 \int_{\Omega} \rho(\psi_i \cdot \mathbf{w}) d\Omega \\ & + \int_{\Omega} \operatorname{grad}(\phi_i) : \boldsymbol{\sigma}(\mathbf{v}) d\Omega - \int_{\Omega} \eta(\operatorname{grad}(\phi_i) : \boldsymbol{\sigma}(\mathbf{w})) d\Omega \end{aligned}$$

[†]For the purposes of this chapter, a first-order tensor is a vector.

$$\begin{aligned}
& + \int_{\Omega} \eta(\text{grad}(\boldsymbol{\psi}_i) : \boldsymbol{\sigma}(\mathbf{v})) \, d\Omega + \int_{\Omega} \text{grad}(\boldsymbol{\psi}_i) : \boldsymbol{\sigma}(\mathbf{w}) \, d\Omega \\
& = \\
& \int_{\Omega} \boldsymbol{\phi}_i \cdot \mathbf{f} \, d\Omega + \int_{\Gamma} \boldsymbol{\phi}_i \cdot \boldsymbol{\sigma}(\mathbf{v})\mathbf{n} \, d\Gamma - \int_{\Gamma} \eta(\boldsymbol{\phi}_i \cdot \boldsymbol{\sigma}(\mathbf{w})\mathbf{n}) \, d\Gamma \\
& + \int_{\Gamma} \eta(\boldsymbol{\psi}_i \cdot \boldsymbol{\sigma}(\mathbf{v})\mathbf{n}) \, d\Gamma + \int_{\Gamma} \boldsymbol{\psi}_i \cdot \boldsymbol{\sigma}(\mathbf{w})\mathbf{n} \, d\Gamma. \tag{5.30}
\end{aligned}$$

Finally, recall the fact that $\boldsymbol{\phi}_i = \boldsymbol{\psi}_i = \mathbf{0}$, $\forall \mathbf{x} \in \Gamma_1$ and $(\boldsymbol{\sigma}(\mathbf{v}) - \eta\boldsymbol{\sigma}(\mathbf{w}))\mathbf{n} = (\eta\boldsymbol{\sigma}(\mathbf{v}) + \boldsymbol{\sigma}(\mathbf{w}))\mathbf{n} = \mathbf{0}$, $\forall \mathbf{x} \in \Gamma_2$. This leads to the final weak form

$$\begin{aligned}
& -\omega^2 \int_{\Omega} \rho(\boldsymbol{\phi}_i \cdot \mathbf{v}) \, d\Omega - \omega^2 \int_{\Omega} \rho(\boldsymbol{\psi}_i \cdot \mathbf{w}) \, d\Omega \\
& + \int_{\Omega} \text{grad}(\boldsymbol{\phi}_i) : \boldsymbol{\sigma}(\mathbf{v}) \, d\Omega - \int_{\Omega} \eta(\text{grad}(\boldsymbol{\phi}_i) : \boldsymbol{\sigma}(\mathbf{w})) \, d\Omega \\
& + \int_{\Omega} \eta(\text{grad}(\boldsymbol{\psi}_i) : \boldsymbol{\sigma}(\mathbf{v})) \, d\Omega + \int_{\Omega} \text{grad}(\boldsymbol{\psi}_i) : \boldsymbol{\sigma}(\mathbf{w}) \, d\Omega = \int_{\Omega} \boldsymbol{\phi}_i \cdot \mathbf{f} \, d\Omega. \tag{5.31}
\end{aligned}$$

In addition, however, note that Equation (5.12) implies

$$\begin{aligned}
\text{grad}(\boldsymbol{\phi}_i) : \boldsymbol{\sigma}(\mathbf{v}) & = \text{grad}(\boldsymbol{\phi}_i) : (\lambda \text{div}(\mathbf{v})\mathbf{I}) + \text{grad}(\boldsymbol{\phi}_i) : (2\mu\boldsymbol{\epsilon}(\mathbf{v})) \\
& = \lambda \text{div}(\boldsymbol{\phi}_i)\text{div}(\mathbf{v}) + 2\mu(\boldsymbol{\epsilon}(\boldsymbol{\phi}_i) : \boldsymbol{\epsilon}(\mathbf{v})) \tag{5.32}
\end{aligned}$$

and similarly for the other terms. This form is most convenient, because `deal.II` has optimized built-in functions for the divergence, gradient, symmetric gradient, and the double contraction of two such symmetric gradients.

For the finite element method (FEM), it is necessary to have a system of equations of the form $\mathbf{G}\mathbf{d} = \mathbf{b}$. The usual assumption is that $\mathbf{v} = \sum d_j \boldsymbol{\phi}_j$ and $\mathbf{w} = \sum d_j \boldsymbol{\psi}_j$. This is the reason that the test functions were denoted with the index i from the start: they are actually the finite basis functions that will be used to interpolate the solution. In actuality, test functions are usually assumed to be of the form $\boldsymbol{\phi} = \sum \boldsymbol{\phi}_i$ and $\boldsymbol{\psi} = \sum \boldsymbol{\psi}_i$, which is known as the Galerkin method. The derivation of the weak form is somewhat easier, however, if it is considered for a single basis function from the outset.

Next, a mass matrix, \mathbf{M} , stiffness matrix, \mathbf{K} , damping matrix, \mathbf{C} , and body force vector, \mathbf{b} , can be defined as follows:

$$\begin{aligned}
M_{ij} & = \int_{\Omega} \rho(\boldsymbol{\phi}_i \cdot \boldsymbol{\phi}_j + \boldsymbol{\psi}_i \cdot \boldsymbol{\psi}_j) \, d\Omega, \tag{5.33} \\
K_{ij} & = \int_{\Omega} \lambda(\text{div}(\boldsymbol{\phi}_i)\text{div}(\boldsymbol{\phi}_j) + \text{div}(\boldsymbol{\psi}_i)\text{div}(\boldsymbol{\psi}_j))
\end{aligned}$$

$$+ 2\mu \left(\boldsymbol{\epsilon}(\boldsymbol{\phi}_i) : \boldsymbol{\epsilon}(\boldsymbol{\phi}_j) + \boldsymbol{\epsilon}(\boldsymbol{\psi}_i) : \boldsymbol{\epsilon}(\boldsymbol{\psi}_j) \right) d\Omega, \quad (5.34)$$

and

$$C_{ij} = \int_{\Omega} \eta \left[\lambda \left(\operatorname{div}(\boldsymbol{\psi}_i) \operatorname{div}(\boldsymbol{\phi}_j) - \operatorname{div}(\boldsymbol{\phi}_i) \operatorname{div}(\boldsymbol{\psi}_j) \right) + 2\mu \left(\boldsymbol{\epsilon}(\boldsymbol{\psi}_i) : \boldsymbol{\epsilon}(\boldsymbol{\phi}_j) - \boldsymbol{\epsilon}(\boldsymbol{\phi}_i) : \boldsymbol{\epsilon}(\boldsymbol{\psi}_j) \right) \right] d\Omega \quad (5.35)$$

$$b_i = \int_{\Omega} \boldsymbol{\phi}_i \cdot \mathbf{f} d\Omega \quad (5.36)$$

Such that

$$\mathbf{G} \mathbf{d} = \left(-\omega^2 \mathbf{M} + \mathbf{K} + \mathbf{C} \right) \mathbf{d} = \mathbf{b}. \quad (5.37)$$

This is certainly the most natural way to think of the problem in the dynamical sense. However, note that for n_{dof} degrees of freedom, $\mathbf{M}, \mathbf{K}, \mathbf{C} \in \mathbb{R}^{n_{\text{dof}} \times n_{\text{dof}}}$ and their sum is (in general) neither symmetric nor definite. It is therefore useful to choose $\boldsymbol{\phi}_i$ and $\boldsymbol{\psi}_i$ in such a way that the solution is separable into real and imaginary parts. One particularly useful way is to choose $\boldsymbol{\phi}_i = \mathbf{0}$ for $i \in [1, \frac{1}{2}n_{\text{dof}}]$, $\boldsymbol{\psi}_i = \mathbf{0}$ for $i \in [\frac{1}{2}n_{\text{dof}}+1, n_{\text{dof}}]$, and $\boldsymbol{\phi}_i = \boldsymbol{\psi}_{(i+\frac{1}{2}n_{\text{dof}})}$ for $i \in [1, \frac{1}{2}n_{\text{dof}}]$. This allows the definition of blocks $\mathbf{Q}, \mathbf{U}, \mathbf{V} \in \mathbb{R}^{\frac{1}{2}n_{\text{dof}} \times \frac{1}{2}n_{\text{dof}}}$

$$Q_{ij} = \int_{\Omega} -\omega^2 \rho(\boldsymbol{\phi}_i \cdot \boldsymbol{\phi}_j) + \lambda \operatorname{div}(\boldsymbol{\phi}_i) \operatorname{div}(\boldsymbol{\phi}_j) + 2\mu \left(\boldsymbol{\epsilon}(\boldsymbol{\phi}_i) : \boldsymbol{\epsilon}(\boldsymbol{\phi}_j) \right) d\Omega, \quad (5.38)$$

$$U_{ij} = \int_{\Omega} -\eta \left[\lambda \operatorname{div}(\boldsymbol{\phi}_i) \operatorname{div}(\boldsymbol{\psi}_j) + 2\mu \left(\boldsymbol{\epsilon}(\boldsymbol{\phi}_i) : \boldsymbol{\epsilon}(\boldsymbol{\psi}_j) \right) \right] d\Omega, \quad (5.39)$$

and

$$V_{ij} = \int_{\Omega} -\omega^2 \rho(\boldsymbol{\psi}_i \cdot \boldsymbol{\psi}_j) + \lambda \operatorname{div}(\boldsymbol{\psi}_i) \operatorname{div}(\boldsymbol{\psi}_j) + 2\mu \left(\boldsymbol{\epsilon}(\boldsymbol{\psi}_i) : \boldsymbol{\epsilon}(\boldsymbol{\psi}_j) \right) d\Omega, \quad (5.40)$$

such that

$$\mathbf{G} \mathbf{d} = \begin{bmatrix} \mathbf{Q} & \mathbf{U} \\ -\mathbf{U}^T & \mathbf{V} \end{bmatrix} \begin{bmatrix} \boldsymbol{\alpha} \\ \boldsymbol{\gamma} \end{bmatrix} = \begin{bmatrix} \boldsymbol{\beta} \\ \mathbf{0} \end{bmatrix} = \mathbf{b}. \quad (5.41)$$

This block form of the system produces the same solution as the original system in Equation (5.37), but has some additional properties that are beneficial. In particular, given the relationship between $\boldsymbol{\phi}_i$ and $\boldsymbol{\psi}_i$, it is easy to show that $\mathbf{Q} = \mathbf{V}$ and $\mathbf{U}^T = \mathbf{U}$. Moreover, \mathbf{U} is positive definite, although \mathbf{Q} is (in general) not. For this reason, it is

advantageous to solve the modified system

$$\mathbf{G}'\mathbf{d} = \begin{bmatrix} \mathbf{U} & -\mathbf{V} \\ \mathbf{V} & \mathbf{U} \end{bmatrix} \begin{bmatrix} \boldsymbol{\alpha} \\ \boldsymbol{\gamma} \end{bmatrix} = \begin{bmatrix} \mathbf{0} \\ \boldsymbol{\beta} \end{bmatrix} = \mathbf{b}'. \quad (5.42)$$

Again, Equation (5.42) produces exactly the same solution for $\boldsymbol{\alpha}$ and $\boldsymbol{\gamma}$ as Equation (5.41). As a matter of fact, this modified system can be derived directly by multiplying Equation (5.6) by \mathbf{i} and carrying out the previous derivation. The matrix \mathbf{V} captures the undamped dynamic mass and stiffness portion of the weak form, and is symmetric but not generally definite, while \mathbf{U} captures the hysteretic material damping part of the weak form and is symmetric positive-definite (SPD). Having an SPD matrix on the diagonal permits some advantageous tactics in terms of solving the block system. In particular, it can be shown that

$$\begin{bmatrix} \mathbf{U} & -\mathbf{V} \\ \mathbf{V} & \mathbf{U} \end{bmatrix} = \begin{bmatrix} \mathbf{I} & \mathbf{0} \\ \mathbf{V}\mathbf{U}^{-1} & \mathbf{I} \end{bmatrix} \begin{bmatrix} \mathbf{U} & \mathbf{0} \\ \mathbf{0} & \mathbf{S} \end{bmatrix} \begin{bmatrix} \mathbf{I} & -\mathbf{U}^{-1}\mathbf{V} \\ \mathbf{0} & \mathbf{I} \end{bmatrix}, \quad (5.43)$$

where $\mathbf{S} = \mathbf{U} + \mathbf{V}\mathbf{U}^{-1}\mathbf{V}$ is the Schur complement. The inverse of Equation (5.43) is then easily written as

$$\begin{bmatrix} \mathbf{U} & -\mathbf{V} \\ \mathbf{V} & \mathbf{U} \end{bmatrix}^{-1} = \begin{bmatrix} \mathbf{I} & \mathbf{U}^{-1}\mathbf{V} \\ \mathbf{0} & \mathbf{I} \end{bmatrix} \begin{bmatrix} \mathbf{U}^{-1} & \mathbf{0} \\ \mathbf{0} & \mathbf{S}^{-1} \end{bmatrix} \begin{bmatrix} \mathbf{I} & \mathbf{0} \\ -\mathbf{V}\mathbf{U}^{-1} & \mathbf{I} \end{bmatrix}. \quad (5.44)$$

Note that \mathbf{S} is generally dense, even if \mathbf{U} and \mathbf{V} are sparse. Practically speaking, this means that the inverse of \mathbf{S} is expensive to calculate. Moreover, the inverse is nonlinearly dependent upon frequency, which means it must be recalculated for each analysis frequency. However, \mathbf{U} is frequency-independent and SPD. Bearing this in mind, the following precondition matrix was chosen to speed up solution of Equation (5.42):

$$\mathbf{P}^{-1} = \begin{bmatrix} \mathbf{I} & \mathbf{U}^{-1}\mathbf{V} \\ \mathbf{0} & \mathbf{I} \end{bmatrix} \begin{bmatrix} \mathbf{U}^{-1} & \mathbf{0} \\ \mathbf{0} & \mathbf{U}^{-1} \end{bmatrix} = \begin{bmatrix} \mathbf{U}^{-1} & \mathbf{U}^{-1}\mathbf{V}\mathbf{U}^{-1} \\ \mathbf{0} & \mathbf{U}^{-1} \end{bmatrix}. \quad (5.45)$$

With this definition of \mathbf{P}^{-1} , the preconditioned system is

$$\mathbf{G}'\mathbf{P}^{-1}\mathbf{P}\mathbf{d} = \mathbf{G}'\mathbf{P}^{-1}\mathbf{y} = \begin{bmatrix} \mathbf{U} & -\mathbf{V} \\ \mathbf{V} & \mathbf{U} \end{bmatrix} \begin{bmatrix} \mathbf{U}^{-1} & \mathbf{U}^{-1}\mathbf{V}\mathbf{U}^{-1} \\ \mathbf{0} & \mathbf{U}^{-1} \end{bmatrix} \begin{bmatrix} \mathbf{y}_\alpha \\ \mathbf{y}_\gamma \end{bmatrix}$$

$$= \begin{bmatrix} \mathbf{I} & \mathbf{0} \\ \mathbf{V}\mathbf{U}^{-1} & \mathbf{I} + (\mathbf{V}\mathbf{U}^{-1})^2 \end{bmatrix} \begin{bmatrix} \mathbf{y}_\alpha \\ \mathbf{y}_\gamma \end{bmatrix}. \quad (5.46)$$

In practice, the particular values in \mathbf{y} are not of interest. Rather, the preconditioned system is solved for \mathbf{y} as an intermediate step, and then \mathbf{d} is determined from \mathbf{y} using $\mathbf{d} = \mathbf{P}^{-1}\mathbf{y}$. The performance of this preconditioner will depend on the spectral radius of $\mathbf{V}\mathbf{U}^{-1}$, often denoted $\rho(\mathbf{V}\mathbf{U}^{-1})$. The smaller $\rho(\mathbf{V}\mathbf{U}^{-1})$ is, the closer the preconditioned system is to the identity matrix and the more effective the preconditioner is. The spectral radius is estimated to be on the order of $\frac{1}{\eta} \left| \left(\frac{\omega}{\omega_0} \right)^2 - 1 \right|$, where ω_0 is the first modal frequency of the system *in the absence of damping*. This comes from the estimate that $\mathbf{C} \approx i\eta\mathbf{K}$ in many practical cases. Thus, the preconditioner works best for low frequencies and large values of η .

The significance of \mathbf{U} being frequency-independent is that its inverse needs to be calculated only once, immediately after the matrices are assembled. The inverse can then be saved and used for all future calculations. The significance of it being SPD is that if the inverse is expensive to calculate directly, there are efficient methods to solve for it iteratively for each instance of vector multiplication with \mathbf{U}^{-1} . The prototypical solver is the conjugate gradient (CG) method, which is remarkably fast but only operates on SPD matrices. This preconditioner therefore allows the user some amount of control to choose a direct inner method or an iterative inner method depending on the complexity of \mathbf{U} .

5.3.3 Dissipated power formulation

As a preliminary to the following power formulation, it is important to know when using the complex notation $e^{i\omega t}$ carries comparable physical meaning to using $\cos(\omega t)$ and when it doesn't. For example, it is common to represent a harmonic force and velocity as

$$f(t) = F \cos(\omega t) \quad (5.47)$$

and

$$v(t) = V \cos(\omega t + \theta), \quad (5.48)$$

where θ is the phase difference between the force and the resultant velocity. Alternatively, these quantities can be expressed in complex notation as

$$\tilde{f}(t) = \tilde{F}e^{i\omega t} \quad (5.49)$$

and

$$\tilde{v}(t) = \tilde{V}e^{i\omega t}, \quad (5.50)$$

where the tilde is used to represent a complex quantity, and in most cases $\tilde{F} = F$ and $\tilde{V} = Ve^{i\theta}$. The complex representation is convenient in many applications, but it is important to keep in mind that only the real part represents a physical, measurable quantity. Indeed, using the definitions above, it is clear that

$$f(t) = \text{Re}\{\tilde{f}(t)\} = \frac{1}{2}(\tilde{f} + \tilde{f}^*) \quad (5.51)$$

and

$$v(t) = \text{Re}\{\tilde{v}(t)\} = \frac{1}{2}(\tilde{v} + \tilde{v}^*). \quad (5.52)$$

So for a real instantaneous power, $q(t) = f(t)v(t)$, it is fine to write the following equivalent form using complex notation: $q(t) = \frac{1}{4}(\tilde{f} + \tilde{f}^*)(\tilde{v} + \tilde{v}^*)$. It is then straightforward to show that the average (real) power over one period can be expressed as

$$\langle q(t) \rangle = \frac{\omega}{2\pi} \int_0^{\frac{2\pi}{\omega}} q(t) dt = \frac{1}{2}|\tilde{F}||\tilde{V}|\cos(\theta). \quad (5.53)$$

While mathematically valid, an instantaneous complex power, $\tilde{q}(t) = \tilde{f}(t)\tilde{v}(t)$, does not carry the same physical meaning. In particular,

$$\langle \tilde{q}(t) \rangle = \frac{\omega}{2\pi} \int_0^{\frac{2\pi}{\omega}} \tilde{q}(t) dt = \frac{\omega}{2\pi} \int_0^{\frac{2\pi}{\omega}} \tilde{F}\tilde{V}e^{i2\omega t} dt = 0. \quad (5.54)$$

Therefore, for the following power formulation, only real power will be considered.

The (real) power flow in a continuum will now be derived, following the logic of [117]. As already mentioned, it is common for a harmonic displacement to be represented as a complex exponential, which means that the representation of stress and strain will be similarly affected. In particular, if $\mathbf{u} \rightarrow \tilde{\mathbf{u}} = \tilde{\mathbf{U}}e^{i\omega t}$ then $\boldsymbol{\epsilon} \rightarrow \tilde{\boldsymbol{\epsilon}} = \tilde{\mathbf{E}}e^{i\omega t}$ and $\boldsymbol{\sigma} \rightarrow \tilde{\boldsymbol{\sigma}} = \tilde{\boldsymbol{\Sigma}}e^{i\omega t}$. In such a case, it is also common to represent the hysteretic material damping as a modification of the stiffness tensor. Namely, $\mathbf{C} \rightarrow \tilde{\mathbf{C}} = \mathbf{C}(1 + i\eta)$, where η is the material loss factor. However, as demonstrated by Equation (5.54), it is only physically valid to deal with real quantities when calculating power. For that reason, note that

$$\sigma_{ij} = \text{Re}\{\tilde{\sigma}_{ij}\} = \sum_{k=1}^3 \sum_{l=1}^3 C_{ijkl} \text{Re}\{\tilde{\epsilon}_{kl}\} - \eta C_{ijkl} \text{Im}\{\tilde{\epsilon}_{kl}\}$$

$$\begin{aligned}
&= \sum_{k=1}^3 \sum_{l=1}^3 C_{ijkl} \epsilon_{kl} + \frac{\eta}{\omega} C_{ijkl} \dot{\epsilon}_{kl} \\
&= \sigma_{ij}^e + \sigma_{ij}^d.
\end{aligned} \tag{5.55}$$

The first term, σ_{ij}^e , represents the elastic stress from storage of energy, while the second term, σ_{ij}^d , represents the dissipative stress from viscous losses.

To derive the power in a viscoelastic continuum, note that the kinetic energy density (i.e., kinetic energy per unit volume) at time t can be expressed as

$$T = \int_0^t \frac{\partial}{\partial t} \left(\frac{1}{2} \rho \dot{\mathbf{u}} \cdot \dot{\mathbf{u}} \right) d\tau = \int_0^t \rho \dot{\mathbf{u}} \cdot \ddot{\mathbf{u}} d\tau. \tag{5.56}$$

As a reminder, $\mathbf{u} = \text{Re}\{\tilde{\mathbf{u}}\}$ for the current analysis. Similarly to kinetic energy density, the potential energy density and the dissipated energy density can be expressed as

$$V = \int_0^t \boldsymbol{\sigma}^e : \dot{\boldsymbol{\epsilon}} d\tau = \int_0^t \sum_{i,j,k,l} C_{ijkl} \epsilon_{kl} \dot{\epsilon}_{ij} d\tau \tag{5.57}$$

and

$$D = \int_0^t \boldsymbol{\sigma}^d : \dot{\boldsymbol{\epsilon}} d\tau = \int_0^t \sum_{i,j,k,l} \frac{\eta}{\omega} C_{ijkl} \dot{\epsilon}_{kl} \dot{\epsilon}_{ij} d\tau, \tag{5.58}$$

respectively. The colon notation used in Equations (5.57) and (5.58) represents a double contraction of two second-order tensors, i.e., $\mathbf{A} : \mathbf{B} = \sum_i \sum_j A_{ij} B_{ij}$. When expressed as above, it is clear from inspection that the time rates of change of these energies are

$$\dot{T} = \rho \dot{\mathbf{u}} \cdot \ddot{\mathbf{u}}, \tag{5.59}$$

$$\dot{V} = \boldsymbol{\sigma}^e : \dot{\boldsymbol{\epsilon}}, \tag{5.60}$$

and

$$\dot{D} = \boldsymbol{\sigma}^d : \dot{\boldsymbol{\epsilon}}. \tag{5.61}$$

The power input to the system (per unit volume) is $\Pi_I = \mathbf{f} \cdot \dot{\mathbf{u}}$, while the power absorbed by the system (per unit volume) is the combined total of the three above rates of energy change, $\Pi_A = \dot{T} + \dot{V} + \dot{D}$. In the case of a unit point force in the z -direction at \mathbf{x}_f , the time-averaged input power can be calculated directly from the complex displacement amplitude as

$$\langle \Pi_I \rangle = \frac{\omega}{2\pi} \int_0^{\frac{2\pi}{\omega}} \mathbf{f} \cdot \dot{\mathbf{u}} dt = -\frac{1}{2} \omega \text{Im} \left\{ \tilde{U}_z(\mathbf{x}_f) \right\}. \tag{5.62}$$

Without going through the full derivation, note that the time-averaged rate of change of

Table 5.1. Table of parameters used in the optimization.

Material Parameters	
Plate Young's modulus, E	70 GPa
Plate density, ρ	$2,700 \text{ kg}\cdot\text{m}^{-3}$
Plate Poisson's ratio, ν	0.33
Plate loss factor, η	0.004
Damping Parameters	
Damping Young's modulus, E_d	0.7 GPa
Damping density, ρ_d	$1,800 \text{ kg}\cdot\text{m}^{-3}$
Damping Poisson's ratio, ν_d	0.33
Damping loss factor, η_d	0.4
Geometric Parameters	
Plate width, a	1 m
Plate thickness, h_1	1 cm
Damping thickness, h_d	1 cm
Runtime Parameters	
Variable 1 bounds, $[v_1^-, v_1^+]$	$[0.001, 1]$
Variable 2 bounds, $[v_2^-, v_2^+]$	$[1, \sqrt{10}]$
Variable 3 bounds, $[v_3^-, v_3^+]$	$[2, 8]$
Variable 4 bounds, $[v_4^-, v_4^+]$	$[-10^6, 1]$
Dimensionless frequency range, $[(ka)_a, (ka)_b]$	$[0.1, 10]$

kinetic energy density and strain energy density are zero when averaged over one period (it is sufficient to consider the relative phases of \mathbf{u} , $\dot{\mathbf{u}}$, and $\ddot{\mathbf{u}}$). This means that the time-averaged absorbed power per unit volume is

$$\langle \Pi_A \rangle = \langle \dot{D} \rangle = \frac{\omega}{2\pi} \int_0^{\frac{2\pi}{\omega}} \dot{D} \, dt = \frac{1}{2} \eta \omega \boldsymbol{\Sigma}^0 : \mathbf{E}^0 = \sum_{i,j,k,l} \frac{1}{2} \eta \omega C_{ijkl} E_{kl}^0 E_{ij}^0, \quad (5.63)$$

where E_{ij}^0 is the amplitude of the (i, j) component of the harmonic strain tensor, $\tilde{\boldsymbol{\epsilon}}$. That is, $E_{ij}^0 = |\tilde{E}_{ij}|$. Integrating $\langle \dot{D} \rangle$ over some domain of interest then gives the average power dissipated within that domain. This is the source of Equation (5.5) used in the definition of J_2 .

5.4 Procedure

The specific parameters used in the optimization are given in Table 5.1. Note that although the problem is formulated in dimensionless quantities, in practice it was calculated in the frequency domain according to the solution strategy outlined in the previous section. For each analysis frequency, the system in Equation (5.42) was solved using the GMRES

algorithm together with the preconditioner given by Equation (5.45). The GMRES algorithm was chosen because both the original block system and the preconditioned system are not symmetric and generally not definite.

The remainder of this section outlines the development and some key features of a bespoke FEM solver written in C++ using the deal.II library [59]. Moving from two to three spatial dimensions increases the problem complexity exponentially. This means exponentially increased computational effort and, potentially, exponentially increased solution time. Because of this, it is worth considering the tools used to solve the physical dynamics of the problem, so as to avoid slowing the optimization process to the point that the time required to reach an optimum becomes intractable.

The commercial software NASTRAN was previously used in Chapter 4 as the finite element method (FEM) solver in the case of two spatial dimensions, and could also be used for the full three-dimensional problem. However, NASTRAN—as well as other commercial FEM software—requires an external interface, separate from the compiled code. Oftentimes, the external interface is graphical to aid the user in formulation of the FEM problem. However, for multiple repeated calculations with many different designs, use of a graphical interface is impractical. NASTRAN does allow one to submit data files that can be used for batch processing, but if the geometry changes dynamically then it is still up to the user to update the data files. Combined with the fact that Borg has up to this point been used through MATLAB, the result is a substantial amount of input and output communication. Having a single executable therefore suggests the potential to significantly speed up solution, as well as simplify use.

5.4.1 Overview of deal.II

To write FEM code requires a handful of key mathematical structures, as well as the numerical tools to represent and operate upon these structures. Although a comprehensive overview of the finite element method is outside the scope of this dissertation, a list of these structures might include

1. A partitioning of the domain into a mesh of cells
2. A means of enumerating cells & nodes
3. Definition of basis functions on a reference cell (the ‘element’)
4. A mapping between the mesh and the reference cell

5. A means of evaluating integrals
6. A means of evaluating function values, gradients, etc.
7. Linear algebra objects (vectors & matrices)
8. A linear solver

Although each of these structures can be implemented in code rather straightforwardly, to write an FEM program that is both accurate and efficient is a non-trivial task, and even expertly coded programs can quickly become very large and complex. It is expedient, then, to use the work of others as building blocks. To that end, the deal.II library was selected to aid in the development of an FEM program for the current work.

deal.II[‡] is a C++ program library targeted at the computational solution of partial differential equations using adaptive finite elements [59]. Note that Borg itself is written in C, and so can be integrated directly with deal.II code. deal.II uses modern programming techniques to offer data structures and algorithms that behave as close to the base mathematical structures as possible. The primary aim of deal.II is to enable rapid development of finite element codes using adaptive meshes, among other tools. It aims to be a program library that takes care of the details of grid handling and refinement, handling of degrees of freedom, input of meshes and output of results, etc. Likewise, support for several space dimensions at once is at the core of the code base, such that programs can be written independent of the space dimension. In reference to the list above, deal.II provides the following classes:

1. `Triangulation<dim> tria;`
2. `DoFHandler<dim> dof_handler(tria);`
3. `FiniteElement<dim> fe(fe_data);`
4. `MappingFE<dim> map(fe);`
5. `Quadrature<dim> quad(points, weights);`
6. `FEValues<dim> fe_values(map, fe, quad);`
7. `Vector<number>` and `SparseMatrix<number>` and many others

[‡]The name ‘deal.II’ stems from its origin as a successor to the Differential Equations Analysis Library of the same authors. Obviously, it has since grown into something else entirely.

8. SparseDirectUMFPACK and many others

Note that all of the classes listed above are internal to deal.II. Their details can be found in the online documentation [118]. While the standalone capabilities of deal.II are sufficient for many applications, the library is also designed to interface with many high-performance computing (HPC) libraries, including (Sca)LAPACK [119, 120], ARPACK [121], PETSc [122], Trilinos [123], and p4est [124]. The latter two in particular were used in the current work for solution of linear equations and to aid in parallelization.

5.4.2 Parallelization

To further improve the finite element solver, the code was parallelized in several ways. Firstly, deal.II wrapper classes were used to utilize vectors and matrices from the Trilinos software library [123]. These vector and matrix classes are designed to be distributed across multiple processors using MPI. As such, large data structures can be split, vectors and matrices can be assembled in parallel, and vector-vector and matrix-vector operations can also be performed in parallel. Secondly, the domain mesh was distributed across processors using the p4est library [124]. This level of parallelization allows large meshes to be split, but also allows cell-level calculations to be performed in parallel. Third, the matrix \mathbf{U}^{-1} in Equation (5.43) was calculated using the parallel direct solver, SuperLU_DIST [125]. Although \mathbf{U}^{-1} could be solved for iteratively at each application, it was decided to solve for it once at the beginning because the problem size is small enough that a direct solution is not unreasonably time-consuming.

Finally, for the outer computations and for exchanging data with the Borg algorithm, a manager/worker scheme was built using MPI. That is, the single manager node exchanges design variables and objective values with the Borg algorithm, as well as instructing the worker nodes on which calculations to carry out. A flowchart representation of the relationships is given in Figure 5.2. The full frequency range was subdivided according to the expected number of modes, using the formula of Xie et al. [126], so that there was theoretically only one mode within each subdivision. A subset of all processors was assigned to each subdivision, so that multiple frequencies could be evaluated simultaneously. It should be noted that the efficiency of this layer of parallelization will depend on the problem size. In particular, there is a trade-off between the number of simultaneous frequency calculations and the amount of data handled by each processor. Although this outer level of parallelization was designed for the particular optimization problem at hand, the architecture could be modified for other use cases

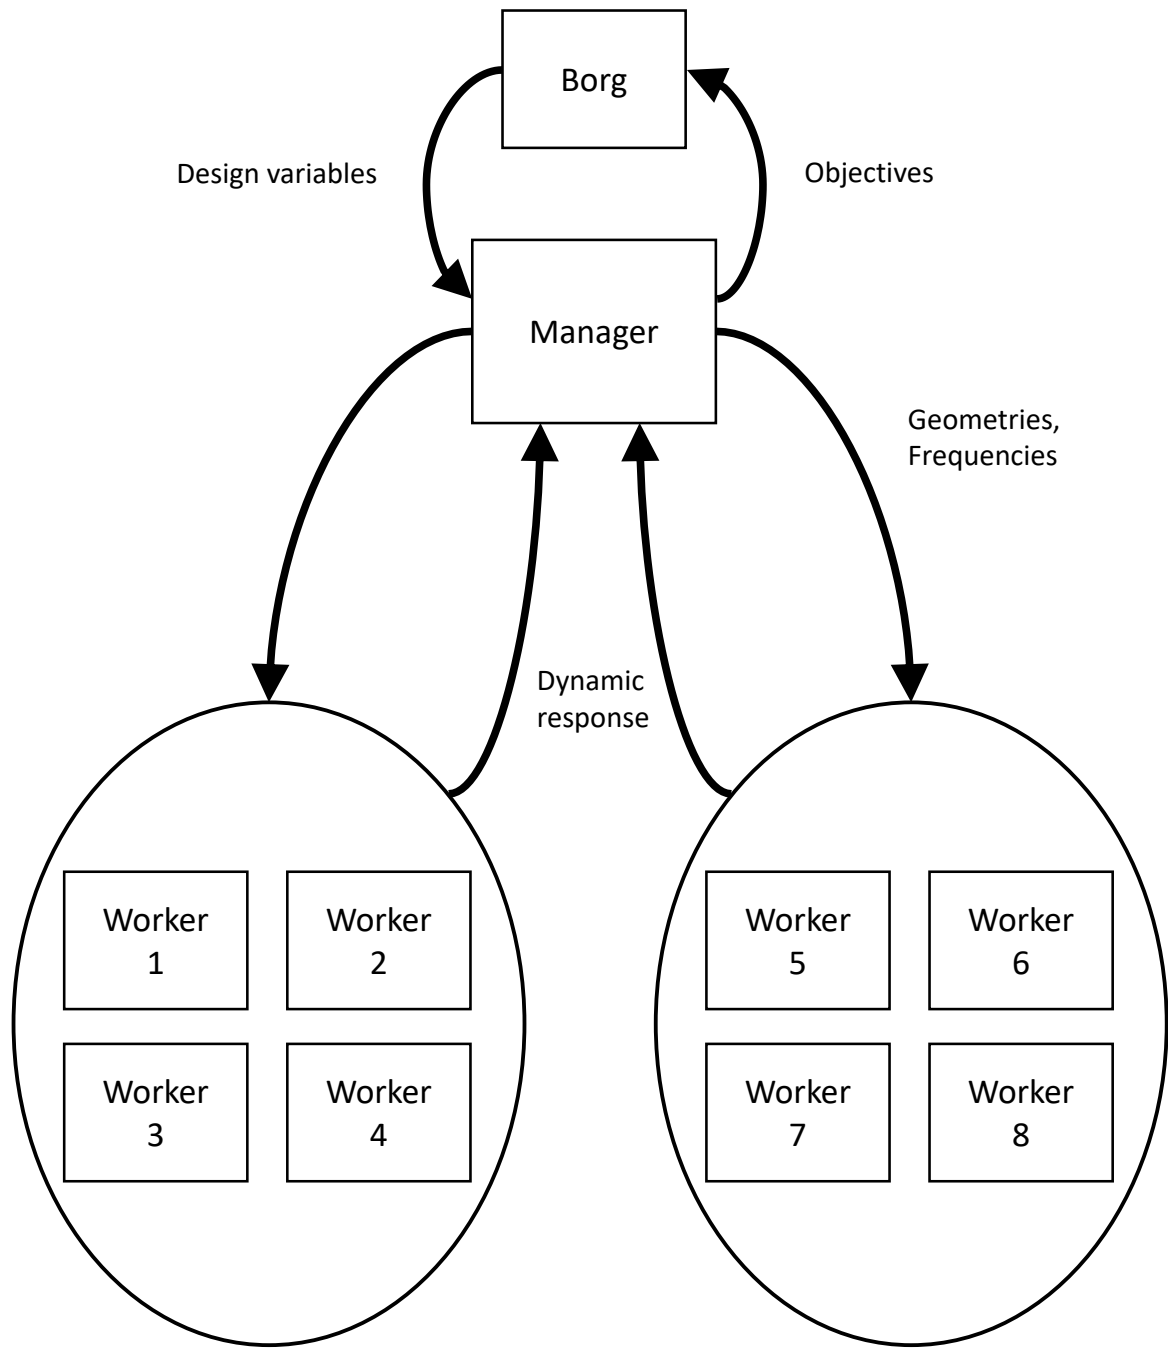


Figure 5.2. Flowchart representation of the multi-level parallelization framework used. The manager process interfaces with Borg to interpret the design variables and to calculate the objectives. Note that workers 1–4 share the same analysis frequencies but each own a different portion of the degrees of freedom. The same is true for workers 5–8. In this way, the construction of matrices and vectors is also parallelized.




Figure 5.3. An example of drive point mobility demonstrating artificial stiffness due to locking phenomena. The solid curve represents the mobility as calculated by NASTRAN with shear and volumetric locking control. The dashed line represents the result when no measures are taken to relieve locking. Clearly, the dashed line shows peaks that are shifted up in frequency, indicating increased stiffness.

that might benefit from parallelization.

5.4.3 Control of shear locking

In finite element analysis of thin-walled structures, so-called ‘locking’ phenomena are a well-known issue for solid elements, such as are used here [127]. At a high-level understanding, locking phenomena occur because the linear or quadratic basis functions cannot adequately represent the full stress field within a given element, and so the element becomes artificially stiff. An example of drive point mobility that demonstrates this element stiffening is shown in Figure 5.3 A simple fix is to use reduced integration, in which fewer integration points are used when evaluating the integrals that include stress and strain. For linear elements, the stress and strain are only evaluated at the center of

the element. However, this often results in zero-energy modes or so-called ‘hourglassing’, in which the strain at the center is zero but the deformation at the nodes is nonzero [127]. Physically, this means that the element has zero stiffness and so the solution results in a structure that is artificially compliant. It should be noted that there are methods to control hourglassing so that reduced integration can be and is used in practice.

A more rigorous and robust method is to define an auxiliary variable, $p = -\frac{1}{3}\text{tr}(\boldsymbol{\sigma}) = -(\lambda + \frac{2\mu}{3})\text{div}(\mathbf{u})$. In using this definition, p is precisely the hydrostatic pressure for a compressible elastic material, where the minus sign is because a positive pressure is usually interpreted as enacting compression, which is a negative strain. With this auxiliary pressure variable, the elastodynamic problem becomes a system of equations in both \mathbf{u} and p and the finite element method is said to be ‘mixed’. At first, this 33% increase in the number of degrees of freedom may seem like a reasonable trade-off for eliminating all locking phenomena. However, the smallest numerically stable element type requires the basis functions to be quadratic in displacement and linear in pressure [128, 129]. Compared to a linear displacement element, this means a 270% increase in the number of degrees of freedom in three dimensions. As such, an alternative means of relieving locking is necessary.

Rather than using reduced integration or mixed methods, the approach used here is sometimes called a selective substitution method. This form of selective substitution is most likely the method used in NASTRAN to alleviate shear locking [130]. The method consists of modifying the off-diagonal terms of the strain tensor at the quadrature points and replacing them with a Jacobian-weighted average. To illustrate, consider first the two-dimensional problem, with a reference cell and Gaussian quadrature points shown in Figure 5.4. The points within some physical cell in the mesh are usually mapped to and from the reference cell using some mapping, $M : (\xi, \eta) \mapsto (x, y)$. This allows the formulation of integrals in terms of a single reference cell, rather than modifying the process for each particular cell. The Jacobian of the mapping from the reference cell to the physical cell is $j(\xi, \eta) = \det(\mathbf{J}(\xi, \eta))$, where $\mathbf{J}(\xi, \eta)$ is the Jacobian matrix of the linear map M , and is related to the change in size between the physical cell and the reference cell.

The selective substitution method replaces the off-diagonal term in the strain, ϵ_{xy} , with a Jacobian-weighted average,

$$\epsilon_{xy} \rightarrow \frac{\sum_q \epsilon_{xy}(\xi_q, \eta_q) j(\xi_q, \eta_q)}{\sum_q j(\xi_q, \eta_q)}, \quad (5.64)$$




Figure 5.4. A 2D reference cell. The points represent the quadrature points for 2×2 Gaussian quadrature. ξ and η (not to be confused with loss factor) are the coordinates in the reference domain.

for each quadrature point on the reference cell, (ξ_q, η_q) . Looking again at Figure 5.4, quadrature points 1–4 are all substituted for this averaged value, so that $\epsilon_{xy}(\xi_1, \eta_1) = \epsilon_{xy}(\xi_2, \eta_2) = \epsilon_{xy}(\xi_3, \eta_3) = \epsilon_{xy}(\xi_4, \eta_4)$. In the case of three dimensions, the situation is more complicated. Referring to Figure 5.5, note that the quadrature points form six planes instead of just one. Additionally, there are three off-diagonal components instead of only one in the two-dimensional case. As such, the substituted value will depend not only on which off-diagonal component is being calculated, but will also depend on which plane the quadrature point lies in. The full list of substitutions is

$$\epsilon_{xy}(\boldsymbol{\xi}_q) = \frac{\sum_{q=1,2,3,4} \epsilon_{xy}(\boldsymbol{\xi}_q) j(\boldsymbol{\xi}_q)}{\sum_{q=1,2,3,4} j(\boldsymbol{\xi}_q)}, \quad \text{for } q = 1, 2, 3, 4; \quad (5.65)$$

$$\epsilon_{xy}(\boldsymbol{\xi}_q) = \frac{\sum_{q=5,6,7,8} \epsilon_{xy}(\boldsymbol{\xi}_q) j(\boldsymbol{\xi}_q)}{\sum_{q=5,6,7,8} j(\boldsymbol{\xi}_q)}, \quad \text{for } q = 5, 6, 7, 8; \quad (5.66)$$




Figure 5.5. A 3D reference cell. The points represent the quadrature points for $2 \times 2 \times 2$ Gaussian quadrature. ξ , η , and ζ are the coordinates in the reference domain. Although the axes are shown off to the corner for ease of readability, the origin is usually taken to be the center of the cell.

$$\epsilon_{xz}(\boldsymbol{\xi}_q) = \frac{\sum_{q=1,2,5,6} \epsilon_{xy}(\boldsymbol{\xi}_q) j(\boldsymbol{\xi}_q)}{\sum_{q=1,2,5,6} j(\boldsymbol{\xi}_q)}, \quad \text{for } q = 1, 2, 5, 6; \quad (5.67)$$

$$\epsilon_{xz}(\boldsymbol{\xi}_q) = \frac{\sum_{q=3,4,7,8} \epsilon_{xy}(\boldsymbol{\xi}_q) j(\boldsymbol{\xi}_q)}{\sum_{q=3,4,7,8} j(\boldsymbol{\xi}_q)}, \quad \text{for } q = 3, 4, 7, 8; \quad (5.68)$$

$$\epsilon_{yz}(\boldsymbol{\xi}_q) = \frac{\sum_{q=1,3,5,7} \epsilon_{xy}(\boldsymbol{\xi}_q) j(\boldsymbol{\xi}_q)}{\sum_{q=1,3,5,7} j(\boldsymbol{\xi}_q)}, \quad \text{for } q = 1, 3, 5, 7; \quad (5.69)$$

and

$$\epsilon_{yz}(\boldsymbol{\xi}_q) = \frac{\sum_{q=2,4,6,8} \epsilon_{xy}(\boldsymbol{\xi}_q) j(\boldsymbol{\xi}_q)}{\sum_{q=2,4,6,8} j(\boldsymbol{\xi}_q)}, \quad \text{for } q = 2, 4, 6, 8. \quad (5.70)$$

Comparing to Figure 5.5, it is clear that for shear in the x - y direction, the quadrature




Figure 5.6. Same mobility as in Figure 5.3, but now the dashed line includes selective substitution to relieve shear locking.

points 1–4 lie in the same plane, while the points 5–8 form in their own plane. The case is similar for shear in the x - z and y - z directions.

Figure 5.6 shows the same mobility calculation as Figure 5.3 but after applying selective substitution to the shear stresses and strains. Clearly the main locking phenomenon is alleviated and the results very nearly match those of a commercial finite element solver. Note that a possible extension would be to relieve any volumetric locking using, e.g., the B-bar method from [61]. This might be important in the case of materials that are nearly incompressible. However, Figure 5.6 demonstrates that shear locking accounts for the majority of the artificial stiffness for the materials considered in this chapter. As such, no steps were taken to avoid what little volumetric locking may be present.




Figure 5.7. Approximate Pareto front for the current problem. Also indicated is the ‘knee’ of the front, at which the trade-off between the two objectives is approximately equal. Although the set of designs found by Borg does not include any with J_1 greater than 0.4, the scale of the abscissa is such that the figure shows all possible values for J_1 . Conversely, the values for J_2 are potentially infinite, and so it has been limited.

5.5 Results and discussion

The optimization was run for a total of 3,000 function evaluations using the Borg algorithm, to produce the approximate Pareto front displayed in Figure 5.7. Note that the abscissa of Figure 5.7 has been scaled to show the full range of possible values for the first objective, J_1 . Although the approximate Pareto front is not entirely complete, it shows a clear trend in trade-off between the two objectives. Also marked in Figure 5.7 is the approximate location of the ‘knee’ in the front, the point at which the trade-off in the two objectives is equal. This acts as an interesting reference point when looking at the variable trends shown in Figure 5.8.

As expected, there is a clear relationship between the first variable, v_1 , and the shape of the Pareto front because $J_1 = v_1^2$. Looking at the other variables, however, there does

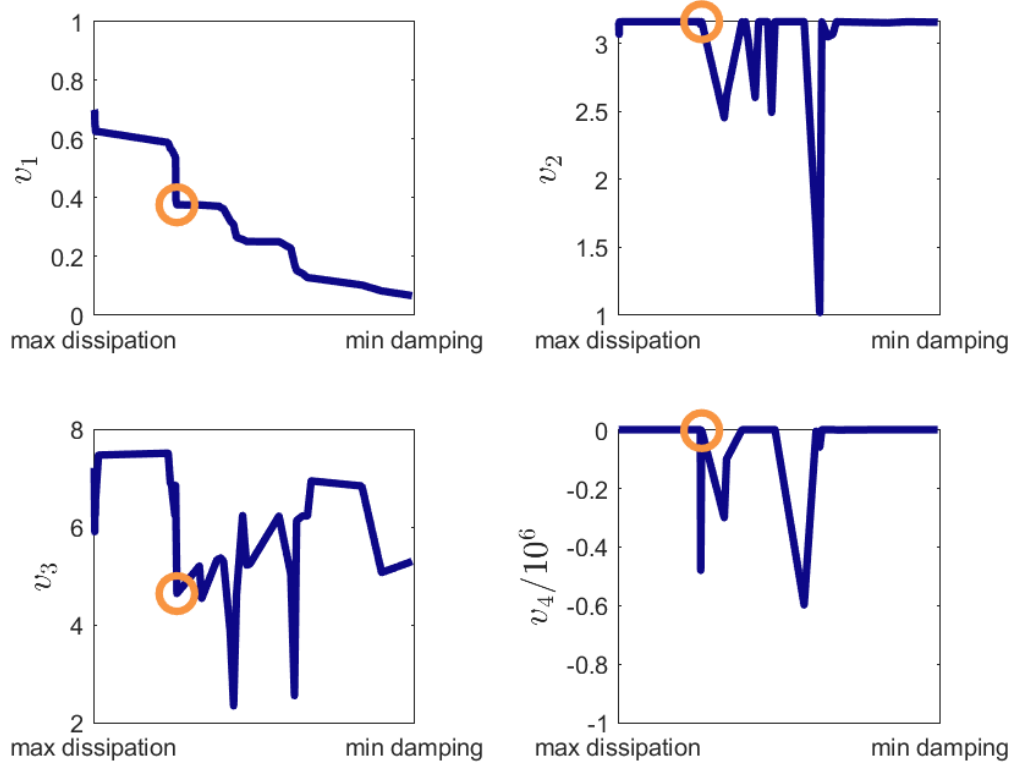


Figure 5.8. Variable trends for the Pareto front in Figure 5.7. The demarcation ‘max dissipation’ indicates the end at which J_2 is minimized, while the ‘min damping’ indicates the end at which J_1 is minimized. The location of the knee is also shown as a circular marker.

not appear to be an obvious trend, with the exception that there appears to be two regions separated by the knee. As the size of the ABH decreases from the knee point, the precise values of v_2 , v_3 , and v_4 fluctuate significantly. On the other hand, as the size of the ABH increases, these variables hold relatively stable, with a clear preference for $v_2 = \sqrt{10}$ and $v_4 = 1$. Although not quite as stable, the value of v_3 is also restricted to between 6 and 8 in this region. This would suggest that there exists a threshold beyond which (for the current damping design) the performance of the ABH vibration absorber is less sensitive to certain variations in the ABH taper profile.

To gain insight into the ka dependence of the dissipated power measure, the ratio of dissipated power to input power was calculated for the design that minimized J_2 . This spectrum is shown in Figure 5.9. By way of comparison, the ratio of dissipated power to input power was also calculated for a uniform plate with the equivalent damping. That is, the damping layer size, shape, and position are equivalent to that of the design that




Figure 5.9. Ratio of power dissipated by the damping to power input to the system, as a function of ka . The spectrum for the ABH design that optimizes J_2 is shown as a solid line. The spectrum for a uniform plate with the equivalent added damping layer is shown as a dashed line. Note that the first modal frequency of the unmodified plate occurs at around $ka = 6$.

minimizes J_2 ; however, the thickness of the plate is uniformly equal to h_1 across its entire span.

In both cases, there is relatively little energy dissipated by the damping layer below a ka of 5. For reference, this corresponds to a frequency of about 50 Hz for the present geometric and material parameters. The ratio of power dissipated increases dramatically in both cases around $ka = 6$, which corresponds to the first modal frequency of the plate. The key difference between the ABH design and the uniform plate is that above $ka = 6$ the ABH plate continues to couple well with the damping layer and provide consistent effective power dissipation. Conversely, the uniform plate sometimes does not couple well to the damping layer.

For comparison, the value of J_2 for the ABH plate represented in Figure 5.9 is 2.13, while the value of J_2 for the uniform plate with equivalent damping is 3.25. Considering these values and looking back to the Pareto front in Figure 5.7, it is clear that even the

uniform plate performs significantly better than many designs in the Pareto optimal set. However, the important point to consider is that for the same damping configuration, the plate with an ABH vibration absorber is around 50% more effective on average at dissipating energy through the damping layer.

5.6 Conclusions

In this chapter, a novel dissipated power ratio was proposed and derived as a measure of ABH vibration absorber performance. The ratio is that of the power dissipated in an attached damping layer to the power input to the system by some excitation, averaged over a range of ka and normalized by the quantity of added damping. This measure is significant in that it is largely independent of scale. Because it is evaluated as a function of ka , structures of different sizes can be directly compared. In the case of ABH structures, this comparison is also facilitated by a set of nondimensional design variables, $\mathbf{h} = \left[2\frac{R_{\text{ABH}}}{a} \quad \sqrt{\frac{h_1}{h_0}} \quad m \quad \frac{\gamma}{h_0} \right]^T = [v_1 \quad v_2 \quad v_3 \quad v_4]^T$. Note that to be truly scale-independent, it would be necessary to account for the effective loss factor of the composite structure in some way, possibly using an analytical model such as that of Oberst [5] or Ross, Kerwin, and Ungar [28]. Normalization by the input power also means that results of two tests can be compared even when the excitation characteristics are somewhat different.

It is important to mention that very recently Huang et al. used a nearly identical power dissipation measure for the simultaneous shape optimization of a two-dimensional ABH and topological optimization of its viscoelastic damping [131]. However, the current measure is distinct in two ways. Firstly, the current power dissipation ratio is a steady-state quantity, using the power flow averaged over one cycle, whereas Huang et al. integrate the dissipated power in the time domain over a period of time comparable to the transit time across the plate. Secondly, the current measure is normalized by the input power, whereas Huang et al. only considered the absolute dissipated energy. This normalization ensures that the measure is independent of the excitation magnitude.

To test the application of the novel power dissipation ratio, a shape optimization was carried out on an ABH vibration absorber. The considered geometry was an ABH embedded in a square plate with an attached free damping layer that covers the entire bottom of the ABH. Using the above nondimensional parameters, the objective was to simultaneously maximize the power dissipation ratio and minimize the volume of applied damping (as such, the power dissipation ratio was not spatially normalized in this case).

The resulting Pareto optimal set showed a clear trade-off between the two objectives that was strongly dependent on the size of the ABH. Also of note was the fact that for the particular damping configuration, designs smaller than a certain size appear to be less sensitive to ABH design. The ABH feature is superior to a uniform plate with the equivalent damping configuration as it allows the damping to dissipate a greater proportion of the input power over a broader range of frequencies.

Although the current study was meant to be a proof of concept, it behooves one to expand the study to additional damping configurations, as well as to different geometric configurations. Plates with differing aspect ratios and damping with different material properties would be of particular interest, as the current power dissipation ratio has no way to account for these variations.

Finally, in addition to the novel dissipated power ratio, a novel finite element method (FEM) solver was developed to analyze the dynamics of the optimization problem. Using a modified block structure, a preconditioner was derived that performs especially well at low frequencies and for structures with large amounts of hysteretic material damping. This particular preconditioner-accelerated solution strategy has not been found in the literature, and may have applications to many elastodynamic problems.

Chapter 6 |

Conclusions

6.1 Summary of findings

This dissertation has investigated the structural optimization of acoustic black hole (ABH) vibration absorbers using a rigorous and robust optimization framework and the Borg multi-objective evolutionary algorithm (MOEA). Such a thorough and systematic series of optimization studies has not been carried out before in the literature. Moreover, at each step of the way the physics modelling techniques have been scrupulously documented, so that they may be easily and accurately applied by others to their own needs or to replicating the findings of this dissertation. Errors in the physical modelling can significantly alter results in structural optimization, so expounding upon the methods so rigorously is not insignificant. Altogether, the work reported in this dissertation has determined several important things:

- In Chapter 2, an optimization framework was established for solving vibroacoustic problems using the multi-objective optimization algorithm (MOEA) called Borg. It was shown that Borg, as an EA, is robust when compared to a standard gradient-based algorithm, `fmincon`, applied to the same structural optimization problem. In particular, the problem was to determine the thickness profile of a cantilever beam that minimizes transmission of vibration energy to a specific region. The problem was taken directly from a conference paper presented at the 10th World Congress on Structural and Multidisciplinary Optimization [2] so as to compare results. Borg reliably converged upon more optimal designs within the same overall number of function evaluations, although on average it took the gradient-based `fmincon` many fewer function evaluations to converge upon a single locally optimal design. Borg's ability to better traverse the search space indicated that EAs are

preferable to gradient-based algorithms for the kinds of vibroacoustic optimization problems considered in this dissertation. This is in no small part due to the fact that acoustic optimization problems—and especially those that include frequency-based objective functions—tend to have search spaces that are noisy, highly nonlinear, and therefore difficult for an optimization algorithm to search effectively. The fact that an EA like Borg might take longer to converge upon a solution* is offset by the fact that the design variables of interest—namely, the ABH shape and position—can be boiled down into only a handful of key parameters. Furthermore, Borg naturally extends to the optimization of multiple objectives in a way that most standard gradient-based algorithms do not. An EA like Borg was therefore determined to be the preferred optimization algorithm to use for the optimization problems considered in this dissertation.

- Chapter 3 studied the optimization of an ABH termination at the free end of a cantilever beam. The single objective was to minimize the spatially-averaged squared vibration response across the entire beam. A significant amount of damping was included in the model, but was itself not a design variable. Unlike the benchmark problem used in Chapter 2, the response was evaluated across a continuous frequency band, rather than at discrete frequencies. The results in Chapter 2, which used discrete analysis frequencies, showed that in at least one case the frequency resolution was much too coarse to effectively evaluate the effect of the beam’s thickness profile on its dynamic response. With this in mind, the beam’s response was instead integrated using a logarithmically-spaced set of analysis frequencies with sufficient resolution to capture peaks in the dynamic response. To speed up analysis and improve analysis quality, a novel, sparse-data formulation of the transfer matrix method (TMM) was developed based on the numerically stable Riccati TMM. The results of the optimization mostly aligned with previous analysis based on reflection coefficients [7, 26]. However, whereas analytical results predict a higher taper power to be optimal, the results in this study found a relatively low taper power to be preferable. The reason for this once again comes down to the problem formulation. While not totally unique, an ABH feature is especially adept at increasing the modal density of a structure. For this reason, the inclusion of an ABH might increase the overall response within a certain frequency range

*As mentioned in Chapter 1, Borg is designed such that it never actually converges. However, in most practical applications there will be a point at which Borg will continue to restart without any significant improvement in the solution.

by increasing the number of resonances, even if the average response at resonance is significantly less than that for the unmodified structure. The optimization was performed again using a frequency range of the same bandwidth but higher in frequency and these additional results confirmed that the increased modal density in a finite structure with an ABH leads to an inherent trade-off in design criteria. While this effect had been mentioned previously in experimental investigations [10, 33], analytical analysis prior to this study had largely considered the base structure to be effectively infinite. As a result, at the time the study in Chapter 3 had been carried out, it was the first known rigorous demonstration that a more extreme taper profile can lead to sub-optimal ABH design.

- Chapter 4 sought to clarify and codify the effects that variation in ABH taper implementation can have on vibration reduction performance. In particular, three styles from the literature—denoted ‘standard symmetric’, ‘standard non-symmetric’, and ‘double-leaf’—were studied using identical optimization frameworks applied to the same optimization problem. That problem was to minimize the spatially-averaged squared velocity response of a beam fixed at either end, and jointly minimize the overall mass of the beam. That is, the problem was a multi-objective optimization problem. A plane strain finite element method was used instead of the TMM because a one-dimensional model does not properly account for the variation in taper implementation. Heeding the findings of Chapter 3, the frequency range in Chapter 4 was chosen carefully by sampling the search space uniformly at random and identifying a region in which the bounds were not likely to straddle a mode. In other words, a slight change in design variables would not significantly change the objective function value. Also, in addition to the taper profile parameters used in Chapter 3, the position of the ABH feature along the beam and the amount of applied damping material were included as design variables. The results showed that the choice of ABH style does have a significant effect on vibration reduction performance. In particular, the standard non-symmetric style performs more effectively than the other two styles. However, it should be noted that the double-leaf style performed comparably to the standard symmetric style and has additional benefits related to stability under static load. Perhaps most interesting, the trade-off between the two objectives—velocity response and overall mass—was shown to be dominated by the amount of added damping. That is, when a significant quantity of added damping material is applied, the performance of an ABH vibration absorber is largely determined by the properties of the damping material. This is significant

because it demonstrates a) there is flexibility in the design of an ABH feature when damping is added, and b) the design of the damping is especially important when applied to an ABH feature. The first point emphasizes that if an ABH taper profile needs to be made sub-optimal to accommodate for other design considerations, the performance of the vibration absorber will not suffer significantly. The second point emphasizes that consideration of the damping application—as well as the style of ABH—should be taken seriously when implementing an ABH vibration absorber. Following the study presented in Chapter 4, other authors have explored the topological optimization of damping layers added to ABH features [34, 36]. They found that for low frequencies below the critical frequency of the base plate, the historical strategy of focusing damping at the thinnest portion of the ABH—as is done in this dissertation—is in fact optimal for vibration reduction. However, for frequencies above the plate’s critical frequency, it was found that for the same quantity of damping material, it is preferable to distribute it evenly across the ABH. Unfortunately, the authors did not provide a comparison of vibration reduction for the distributed damping versus damping concentrated in the thinnest portion. However, their analysis showed that a further reduction in radiated sound power of almost 4 dB could be achieved by distributing the damping material instead of concentrating it. The key takeaway from this and the results of Chapter 4 is that damping design should be a key factor in optimal ABH design and implementation, just as Chapter 4 showed that the choice of ABH taper style should be a key factor in implementation.

- In Chapter 5 a novel performance measure was proposed and tested. The measure, termed the ‘power dissipation ratio’ is the frequency- and spatially-averaged ratio of the power dissipated by an ABH vibration absorber to the power input to the system. Its definition was designed *ab initio* to be independent of the scale of the ABH vibration absorber and host structure, so that one- and two-dimensional ABH structures of differing sizes and materials can be directly compared so long as their dimensions are congruent. After deriving the power dissipation ratio, the measure was used in a test problem. The problem consisted of optimizing the shape of a two-dimensional ABH vibration absorber in a thin plate with fixed boundary conditions. The goal of the optimization was to simultaneously maximize the power dissipation ratio and minimize the side of the ABH. (As such, the power dissipation ratio was not spatially-normalized.) In order to facilitate the optimization, a novel preconditioner-accelerated solution scheme was derived for the

finite element method (FEM) using hexahedral displacement elements to solve the time-harmonic linear elastodynamic equations. The preconditioner relies on a block formulation of the problem and exploits the inherent structure of the system when hysteretic material damping is included. Analysis shows that the preconditioner is especially effective for large damping and relatively low frequencies. It therefore has potential for use in many different applications. Using this preconditioner-accelerated solution scheme, the optimization was carried out for 3,000 function evaluations to produce an approximate Pareto front. This Pareto front showed a strong trade-off between the ABH size and the power dissipation ratio. It may also be possible that performance is somewhat less sensitive to the particular design of the ABH in the case of the damping configuration used. The ABH design that maximized the power dissipation ratio was compared against a uniform plate with the same damping configuration. The results showed that the ABH is more than 50% more effective at dissipating power compared to the uniform plate. By considering the power dissipation ratio as a function of frequency, the ABH is able to couple with the damping layer more effectively across a broader range of frequencies. The results of this test problem show that the power dissipation ratio is useful as a measure in the design of ABH vibration absorbers.

6.2 Suggestions for future work

As with all research, there are many ways in which the research presented in this dissertation could be built upon or its ideas furthered. Some suggestions include

- Only flat beams and plates were considered in this dissertation. However, it might be interesting to perform similar optimization studies on curved panels, which have practical uses in marine and aerospace applications. It would be especially enlightening to consider base structures of varying radii of curvature and the effect this has on optimal ABH design.
- A formal comparison like that in Chapter 4 could be done for two-dimensional ABHs using the same three styles.
- Non-circular perimeters could be considered for two-dimensional ABH geometries, such as ellipses or even a generalized periodic function of angle. Similarly, ABH taper parameters could vary as a function of angle.

- Similarly, only a power-law ABH profile of the form $h(x) = \epsilon(x + \beta)^m + \gamma$ has been considered in this dissertation. Another profile that exhibits the ABH effect is of the form $h(x) = \epsilon \sin^m(\alpha x + \beta) + \gamma$. This optional taper profile could be considered in future research. Other taper profiles that do not produce zero reflection but which may be useful for vibration reduction are given in [39].
- Damping layers with non-uniform thickness could be included in the optimization problem. The thickness could vary proportionally to the thickness of the beam or plate, or with some power-law relation.
- Despite deriving a useful set of nondimensional design variables in Chapter 5, it is still necessary to have four parameters in order to uniquely describe a given ABH profile. It would be advantageous to instead have a single design parameter that can uniquely describe or otherwise significantly categorize all possible ABH profiles. One potential route might be to use the total accumulated phase from the minimum thickness point to the maximum thickness point, although this quantity by itself can describe multiple possible ABH designs.
- Multiple ABH features could be optimized simultaneously in a given structure. This could be a grid of two-dimensional ABHs in a plate, in which case the grid structure might be a design variable, or the ABHs could be allowed to move independently, in which case their positions would be design variables.
- In this dissertation, the ABH effect has been produced by varying the thickness of the base structure. It is also theoretically possible to produce this effect by varying the material properties. This may be preferable in some circumstances and the optimization studies in this dissertation could be easily modified to investigate how this different approach may produce similar or different results.
- A point excitation has been used throughout with the intent of exciting all relevant modes of the structure. However, some authors have used a more realistic excitation model, such as a diffuse field model [132] or a spatially correlated model like turbulent boundary flow [133]. Utilizing such excitation models in the optimization problem formulation might produce different and interesting results.
- Finally, no fluid loading effects were considered in this dissertation. While these are usually negligible for in-air applications, fluid loading becomes significant in

under-water applications. The inclusion of these effects in the optimization process might prove especially enlightening.

Appendix A |

Mesh convergence study

A.1 Introduction

As with many numerical methods, a fundamental aspect of the finite element method (FEM) is defining the problem domain in terms of a number of discrete, connected, elements. In a practical sense, this means dividing some geometry into a number of basic shapes, thus defining the ‘mesh’ or ‘grid’ from which the FEM formulation will define variables and constraints. Ultimately, it is not only the FEM formulation, but also the mesh that determines the numerical data structures that are ultimately used to solve the problem. It is not surprising, then, that in many real-world problems of interest, creating a quality mesh is a problem as complex as the original physical problem [134].

One aspect in the problem of mesh generation is the minimum spatial resolution required to obtain a useful result. If one makes the elements too large, then it is likely that the FEM solution will not have the desired accuracy or will not capture relevant physical characteristics and thereby produce unphysical results. On the other hand, the smaller the size of the elements, the more degrees of freedom (DoFs) and so the longer the calculations will take. At the extreme limit, it may be that there are so many DoFs that the computing resources become exhausted and a numerical solution cannot be obtained at all.

A number of different measures can be derived for the acceptable resolution of a mesh. In the field of acoustics, a popular rule of thumb is that there should be somewhere between six and twelve elements per wavelength [135]. How this rule of thumb came to be is not entirely clear, but Marburg posits that the idea of using a fixed number of elements per wavelength stems from the sampling theorem of Shannon and Nyquist [136]. The theory, which states that the maximum detectable frequency is that which corresponds to two points per period (in time) or per wavelength (in space), is ubiquitous in acoustics

and other fields involving signal processing. In fact, an earlier paper by Margburg indicated that the eigenvalue distribution of a one-dimensional FEM mesh can easily be related to the Nyquist frequency [137]. Namely, for both linear and quadratic elements, the largest eigenfrequency is slightly larger than the frequency which corresponds to two points per wavelength. It is not inconceivable, then, that engineers would build a rule of thumb off of this strict lower limit. The additional factor of three to six may have come from a desire for increased accuracy near this limit.

In 2017, Langer et al. performed an extensive study into the effect of element size on the accuracy of the FEM in thin-walled domains [138]. This class of problem includes beams, plates, and shells, and is of great importance in acoustics. The authors analyzed different element types as well as the number and quality of elements. Their recommendation was that the domain be divided into 20 elements per wavelength for quadratic elements, with a maximum aspect ratio of 1:10, and up to 500 elements per wavelength for linear elements, with a maximum aspect ratio of 1:4. Considering that most software implementations use 8-node linear hexahedral elements and 20-node quadratic elements, this recommendation equates to a tenfold increase in the number of DoFs when using linear elements versus quadratic elements.

It is important to point out a few subtleties from the results of Langer et al. Firstly, the test geometry from which the authors derived their recommendations was a simple beam. The mesh was thus a regular (structured) orthogonal three-dimensional grid—that is, all elements were the same shape and size, and element edges were either perpendicular or parallel at intersection points. The authors did include a test case on a more complex geometry with an irregular (unstructured) grid, but their analysis was limited to the number of DoFs and not directly related to the number of elements per wavelength.

Secondly, the linear elements used in the study used full integration. That is, there were no measures taken to relieve volumetric or shear locking phenomena, which are especially relevant in elements with moderate aspect ratios, as in the case of thin-walled domains. These locking phenomena will result in elements that are artificially stiff, which is precisely what was observed by the authors: the eigenfrequencies predicted with linear elements were significantly higher than the true eigenfrequencies.

The purpose of the work described in this appendix is, in small part, to evaluate the disparate recommendations regarding the number of elements per wavelength. More so, however, its purpose is to establish a bespoke rule of thumb by performing a mesh convergence study on an acoustic black hole (ABH) test case. Additionally, the relative accuracy of linear elements will be compared to that of quadratic elements. The use of

Table A.1. Material parameters and dimensions of the test case. Uniform 1 is the uniform portion to the left of the ABH taper, and Uniform 2 is the uniform portion to the right of it. For the ABH taper profile, $\varepsilon = 3.93 \text{ m}^{-1}$ and $h_0 = 0.508 \text{ mm}$.

	Uniform 1	Taper	Damping 1	Uniform 2	Damping 2
Young's modulus, E (GPa)	70	70	0.1	70	0.1
Density, ρ ($\text{kg}\cdot\text{m}^{-3}$)	2700	2700	1000	2700	1000
Loss factor, η	0	0	0.9	0	1.8
Thickness, h (mm)	6.8	$\varepsilon x^2 + h_0$	2	0.508	2
Length, L (cm)	20.5	4	4	34.3	34.3



Figure A.1. Graphical depiction (not to scale) of the test geometry whose material parameters and dimensions are given in Table A.1. Green denotes the beam material, yellow denotes the Damping 1 material, and orange denotes the Damping 2 material.

linear elements is preferable from a computational standpoint, since the resultant linear systems are as small as possible for a given number of elements.

A.2 Test case description

The material parameters and dimensions of the test case are given in Table A.1 and a graphical depiction is shown in Figure A.1. The lengths of the two uniform segments were chosen so as to ensure no pollution of the received signal at the measurement points. That is, the lengths of Uniform 1 and Uniform 2 are such that any reflections from the end of Uniform 2 or between the start of the taper and the start of Uniform 1 will not reach the measurement points until the incident wave has fully passed the measurement points. The remainder of the dimensions, as well as the material parameters, were taken from a similar study by Feurtado and Conlon [25].

The geometry was first meshed with the particular number of elements per wavelength, using either linear or quadratic hexahedral elements. The wavelength used in the taper was that used for the anechoic termination, Uniform 2. Next, if any of the elements had aspect ratios greater than 1:5, those elements were split into two elements, unless doing so would make the aspect ratio of another element greater than 1:5. Finally, if at any point there was a transition from fewer elements through the thickness to more—or vice versa—then transition regions were meshed using pyramidal elements. An example mesh, showing these transition regions, is given in Figure A.2.




Figure A.2. Example mesh, showing the transition regions using pyramidal elements. This happens to be the mesh used for twelve elements per wavelength with linear elements.

The excitation was an enforced displacement at the far left end of the beam. The excitation profile was a 10 kHz pulse, shaped with a cosine filter to be 0.5 ms in duration, with a peak amplitude of $5 \mu\text{m}$. This excitation signal is shown in Figure A.3. For analysis of the beam displacement response, time-domain FEM calculations were carried out using NASTRAN with a time step of $0.6 \mu\text{s}$. The vertical displacement of the top surface of the beam was recorded at two measurement points: the beginning of the ABH taper and the end of the ABH taper. An example of the raw time series at the two points is given in Figure A.3.

From the recorded time series, cross-correlation functions were computed between the excitation point and the measurement points. In each case, the lag corresponding to the maximum in the cross-correlation was taken as the travel time between the excitation and the response. Each response signal was then filtered using the same cosine filter used for the excitation profile. An envelope was calculated for the excitation and the two filtered measurements using the Hilbert transform, and the peak value of these envelopes were used as measures of the pulses' amplitudes. An example of the filtered signals and envelopes is given in Figure A.4. The ratio of the taper end amplitude to the taper start amplitude was calculated for all combinations of mesh resolution and element type in order to determine convergence.




Figure A.3. Excitation displacement pulse (black), along with an example of the response at the start of the ABH taper (blue) and at the end of the taper (yellow).

A.3 Results

Figure A.5 shows the ratio of the taper end amplitude to the taper start amplitude, using linear and quadratic elements at mesh resolutions ranging from six to 21 elements per wavelength. As is to be expected, the calculated ratio starts at some value and then shifts as the mesh resolution increases, until the change is minimal and the models converge upon a value. In Figure A.5, the curve appears to flatten out at around twelve

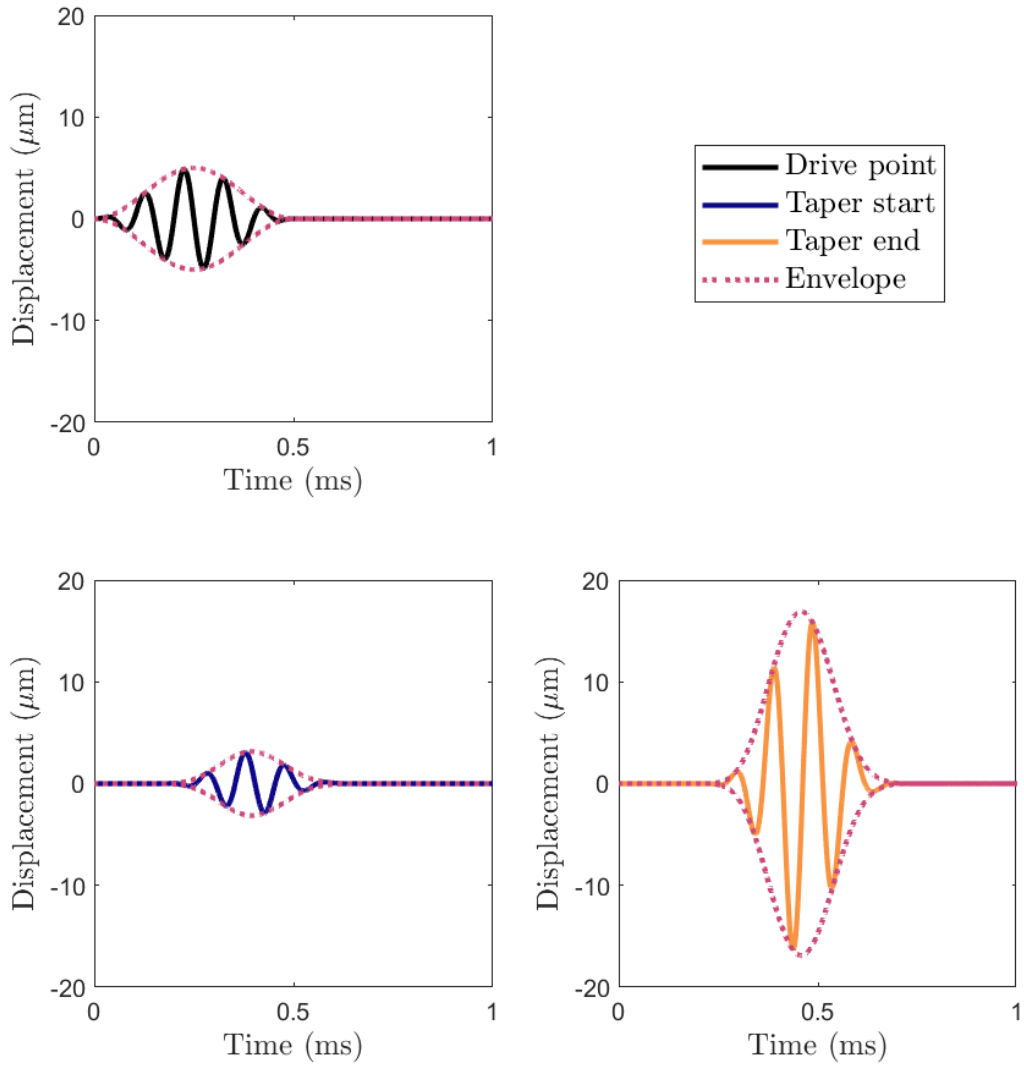


Figure A.4. Same time series as shown in Figure A.3, except that the signals have been filtered with a cosine filter of 0.5 ms width. Also shown are the envelopes calculated using the Hilbert transform (dotted lines).

or fifteen elements per wavelength. That is, there appears not to be much additional accuracy gained from using much more than twelve elements per wavelength. It should be pointed out that, although results for both linear and quadratic elements show similar ratios at similar mesh resolutions, the ratio values obtained using quadratic elements are more steady. That is, there is less variation between using six elements per wavelength or using 21 elements per wavelength in the case of quadratic elements.

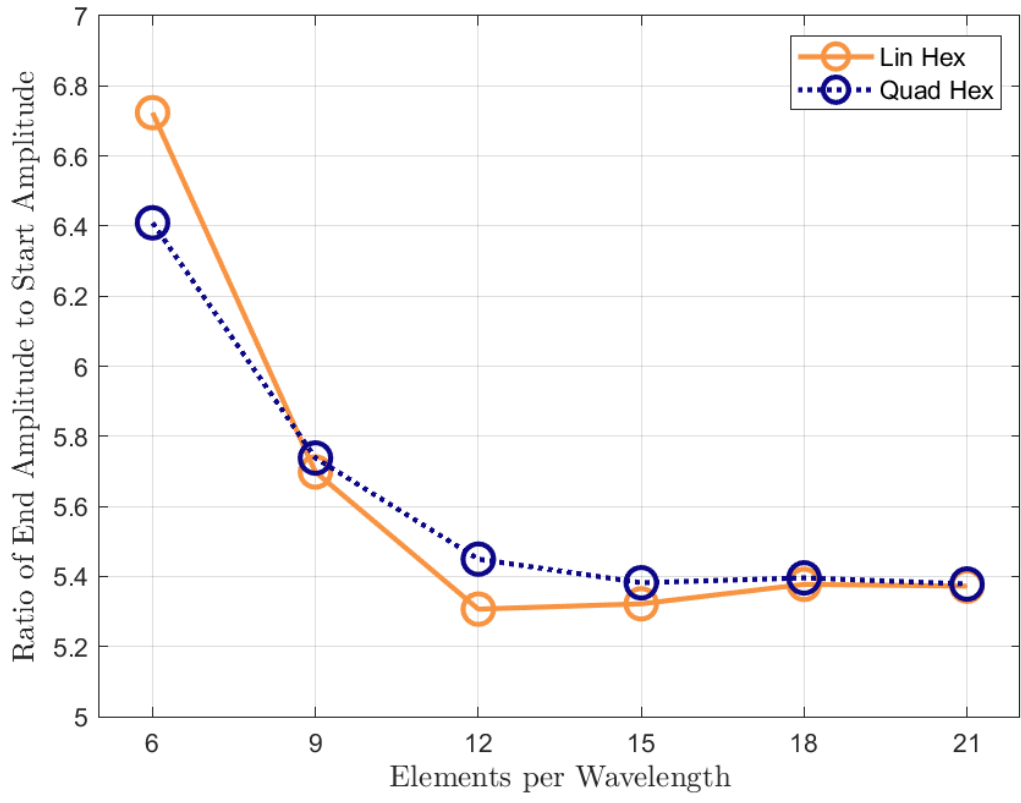


Figure A.5. Ratio of taper end amplitude to taper start amplitude for various combinations of mesh resolution and element type. Note that the theoretical ratio in the absence of any loss mechanisms would be approximately 7, according to WKB theory.

Figure A.6 shows the relative error compared to the results using 21 quadratic elements per wavelength. Figure A.6 appears to show a reduction in error of approximately tenfold for every six additional elements per wavelength. That is, the error for six elements per wavelength is approximately ten times greater than that for twelve elements per wavelength. It is interesting that this trend is present for both linear and quadratic elements, although it should be noted that the trend for quadratic elements may be slightly better, but the maximum number of elements per wavelength is not sufficient to show this difference.

A.4 Conclusions

A fundamental aspect of the finite element method (FEM) is defining the problem domain in terms of elements. However, it is often up to the practitioner to decide the minimum




Figure A.6. Relative L^2 error in the amplitude ratios shown in Figure A.5, using the results of 21 quadratic elements per wavelength as ‘true’. Note the logarithmic ordinate scale. Because the error between 21 quadratic elements per wavelength and itself is zero, it is not shown.

spatial resolution that will produce a useful result. In the field of acoustics, a popular rule of thumb is that there should be somewhere between six and twelve elements per wavelength. The work of Langer et al. has indicated that it may be necessary to use as many as 500 elements per wavelength to obtain accurate results.

A convergence study was carried out to establish a reasonable mesh resolution in the context of acoustic black hole (ABH) tapers. Results from using both linear and quadratic elements indicate that beyond a certain resolution, linear elements are able to obtain accuracy comparable to that of quadratic elements. An important point to highlight is that the linear elements used by NASTRAN incorporate a reduced integration, which alleviates shear and volumetric locking. This is a major difference from the work of Langer et al., and may be the key to why the linear elements performed relatively well against the quadratic elements. The use of linear elements is preferable from a computational standpoint, since the resultant linear systems are as small as possible for

a given number of elements.

Based off of the above convergence results and error analysis, twelve elements per wavelength appears to be a reasonable choice to achieve a balance between a high accuracy and a low number of degrees of freedom. It is interesting that this number corresponds to the upper edge of the popular rule of thumb. At this resolution, there was about a 1% relative error and it is predicted that to reduce the error further would require approximately 50% more degrees of freedom. Therefore, twelve elements per wavelength has been used as the desired mesh resolution when using the FEM in this dissertation.

Appendix B |

FEM meshing calculations

Two profiles that theoretically produce the ABH effect are

$$h(r) = \varepsilon(r + \beta)^m + \gamma \quad (\text{B.1a})$$

$$h(r) = \varepsilon \sin^m(\alpha r + \beta) + \gamma \quad (\text{B.1b})$$

In theory, zero reflection is achieved when $\beta = \gamma = 0$. In the following analysis, however, it is assumed that $\beta \geq 0$ and $-\infty < \gamma \leq h_0$. In profile (B.1a), given a desired minimum thickness, h_0 , maximum thickness, h_1 , length, R , taper power, m , and a prescribed γ , the parameters can be calculated as

$$\beta = \frac{R(h_0 - \gamma)^{\frac{1}{m}}}{(h_1 - \gamma)^{\frac{1}{m}} - (h_0 - \gamma)^{\frac{1}{m}}} \quad (\text{B.2a})$$

$$\varepsilon = \frac{h_1 - \gamma}{(R + \beta)^m} \quad (\text{B.2b})$$

In this way, γ determines the slope of the profile at the minimum thickness. In particular, $\lim_{\gamma \rightarrow h_0} h'(0) = 0$ and $\lim_{\gamma \rightarrow -\infty} h'(0) = \frac{h_1 - h_0}{R}$.

For profile (B.1b), we define a new variable, z , such that $\alpha R + \beta = 2\text{atan}(z)$.* In such a case, the parameters are

$$\varepsilon = \frac{h_1 - \gamma}{2^m} \frac{(1 + z^2)^m}{z^m} \quad (\text{B.3a})$$

*The reason for this is that such a definition facilitates calculation of the extrema of $h'(r)$. In particular, the maximum occurs when $\alpha r + \beta = 2\text{atan}\left(\sqrt{\frac{m-2\sqrt{m+1}}{m-1}}\right)$ and the minimum occurs when $\alpha r + \beta = \frac{\pi}{2}$. Thus, z is allowed to vary from $\sqrt{\frac{m-2\sqrt{m+1}}{m-1}}$ to 1.

$$\beta = \text{asin}\left(\left(\frac{h_0 - \gamma}{\varepsilon}\right)^{\frac{1}{m}}\right) \quad (\text{B.3b})$$

$$\alpha = \frac{2\text{atan}(z) - \beta}{R} \quad (\text{B.3c})$$

The slopes at either end of the taper are determined by the unique combination of γ and z , though in general γ has a greater effect on the slope at the minimum thickness and z has a greater effect on the slope at the maximum thickness.

Now, suppose one wishes to have a certain number of elements per wavelength to describe the taper in the radial direction. The length of the taper, in terms of wavelengths, can be calculated by integrating the wavenumber, k , along the profile. In particular, the number of wavelengths, N_r , is calculated as

$$N_r = \int_0^R \frac{k(r)}{2\pi} dr \quad (\text{B.4})$$

The wavenumber for bending waves in a thin isotropic plate is given as

$$k(r) = \left(\frac{12\omega^2\rho(1-\nu^2)}{E(h(r))^2}\right)^{\frac{1}{4}} \quad (\text{B.5})$$

The integrand in Equation (B.4) is thus

$$\frac{k}{2\pi} = \frac{1}{2\pi} \left(\frac{12\omega^2\rho(1-\nu^2)}{E}\right)^{\frac{1}{4}} h^{-\frac{1}{2}}(r) = Ch^{-\frac{1}{2}}(r) \quad (\text{B.6})$$

For n elements per wavelength along the taper, the number of elements is[†] $M_r = nN_r$. Furthermore, for M_r elements spanning the length of the taper, Equation (B.4) can be written equivalently

$$N_r = \sum_{i=1}^{M_r} \int_{r_{i-1}}^{r_i} \frac{k}{2\pi} dr = \sum_{i=1}^{M_r} \int_{r_{i-1}}^{r_i} Ch^{-\frac{1}{2}}(r) dr \quad (\text{B.7})$$

where the r_i are the positions of the nodes, with $r_0 = 0$ and $r_{M_r} = R$. Given the

[†]If one wishes to impose that M_r is an integer, then in practice the number of elements is $M_r = \text{ceil}(nN_r)$.

relationship between M_r and N_r , it is particularly convenient to assume

$$\int_{r_{i-1}}^{r_i} Ch^{-\frac{1}{2}}(r) dr = \frac{N_r}{M_r} \quad (\text{B.8})$$

for all i . In such a case, given that $r_0 = 0$, the other node positions may be calculated *a priori* as

$$\int_0^{r_i} h^{-\frac{1}{2}}(r) dr = \frac{iN_r}{CM_r} \quad (\text{B.9})$$

In the angular direction, we might assume that the shape of the ABH can be generalized as an ellipse. In such a case, the outer radius, R , is related to the angle by

$$R(\theta) = \frac{ab}{\sqrt{a^2\sin^2(\theta) + b^2\cos^2(\theta)}} \quad (\text{B.10})$$

where the width of the ellipse is $2a$ and its height is $2b$.

The differential length of such an ellipse can be calculated as

$$\begin{aligned} ds &= \sqrt{R^2(\theta) + \left(\frac{dR}{d\theta}\right)^2} d\theta \\ &= R(\theta) \sqrt{1 + \frac{(b^2 - a^2)^2}{4(ab)^4} R^4(\theta) \sin^2(2\theta)} d\theta \\ &= s'(\theta) d\theta \end{aligned} \quad (\text{B.11})$$

Similar to Equation (B.4), the arc length of one quarter of the ellipse, in terms of wavelengths, can be calculated as

$$N_\theta = \int \frac{k}{2\pi} ds = \int_0^{\frac{\pi}{2}} Ch_1^{-\frac{1}{2}} s'(\theta) d\theta \quad (\text{B.12})$$

Notice that the thickness along the outer radius is $h(\theta) = h_1$, independent of angle. If $M_\theta = nN_\theta$ is the number of elements in the angular direction, then an expression similar to Equation (B.9) can be derived for the node positions along the outer radius

$$\int_0^{\theta_i} s'(\theta) d\theta = \frac{iN_\theta h_1^{\frac{1}{2}}}{CM_\theta} \quad (\text{B.13})$$

where $\theta_0 = 0$ and $\theta_{M_\theta} = \frac{\pi}{2}$.

Now that we have expressions that can be used to calculate r_i and θ_i , we take the

analysis one step further. In the limit that M_r and M_θ go to infinity, Equations B.9 and B.13 define continuous mappings to the domain $[0, 1]$. In particular, we define $\xi = \frac{i}{M_\theta}$ and $\eta = \frac{i}{M_r}$ to give the mappings $\theta \mapsto \xi$ and $r \mapsto \eta$

$$\int_0^\theta s'(t) dt = \xi \int_0^{\frac{\pi}{2}} s'(t) dt \quad (\text{B.14})$$

$$\int_0^r h^{-\frac{1}{2}}(t) dt = \eta \int_0^R h^{-\frac{1}{2}}(t) dt \quad (\text{B.15})$$

in which we have used the definitions for N_r and N_θ in Equations B.4 and B.12, respectively, to eliminate the factors C and $h_1^{\frac{1}{2}}$. To determine the interior nodes, we use transfinite interpolation (TFI)

$$\begin{aligned} \theta(\xi, \eta) &= (1 - \eta)\theta_b(\xi) + \eta\theta_t(\xi) + (1 - \xi)\theta_l(\eta) + \xi\theta_r(\eta) \\ &\quad - \left[\xi\eta\theta_t(1) + \xi(1 - \eta)\theta_b(1) + \eta(1 - \xi)\theta_t(0) + (1 - \xi)(1 - \eta)\theta_b(0) \right] \end{aligned} \quad (\text{B.16})$$

$$\begin{aligned} r(\xi, \eta) &= (1 - \eta)r_b(\xi) + \eta r_t(\xi) + (1 - \xi)r_l(\eta) + \xi r_r(\eta) \\ &\quad - \left[\xi\eta r_t(1) + \xi(1 - \eta)r_b(1) + \eta(1 - \xi)r_t(0) + (1 - \xi)(1 - \eta)r_b(0) \right] \end{aligned} \quad (\text{B.17})$$

with

$$\theta_b(\xi) = \theta(\xi), \quad r_b(\xi) = 0 \quad (\text{B.18})$$

$$\theta_t(\xi) = \theta(\xi), \quad r_t(\xi) = R(\theta(\xi)) \quad (\text{B.19})$$

$$\theta_l(\eta) = 0, \quad r_l(\eta) = r^a(\eta) \quad (\text{B.20})$$

$$\theta_r(\eta) = \frac{\pi}{2}, \quad r_r(\eta) = r^b(\eta) \quad (\text{B.21})$$

where $\theta(\xi)$ is the inverse mapping of Equation (B.14), $r^b(\eta)$ is the inverse mapping of Equation (B.15) for the y -axis thickness profile $h^b(r)$, and $r^a(\eta)$ is the inverse mapping of Equation (B.15) for the x -axis thickness profile $h^a(r)$. Note that these mappings also allow us to define the interior thickness profile through TFI. Namely,

$$\begin{aligned} h(\xi, \eta) &= (1 - \eta)h_b(\xi) + \eta h_t(\xi) + (1 - \xi)h_l(\eta) + \xi h_r(\eta) \\ &\quad - \left[\xi\eta h_t(1) + \xi(1 - \eta)h_b(1) + \eta(1 - \xi)h_t(0) + (1 - \xi)(1 - \eta)h_b(0) \right] \end{aligned} \quad (\text{B.22})$$

with

$$h_b(\xi) = h_0 \quad (\text{B.23})$$

$$h_t(\xi) = h_1 \quad (\text{B.24})$$

$$h_l(\eta) = h^a(r^a(\eta)) \quad (\text{B.25})$$

$$h_r(\eta) = h^b(r^b(\eta)) \quad (\text{B.26})$$

In this way, one can define a mapping from the unit cube $[0, 1]^3$ to the desired profile and vice versa, namely $(\xi, \eta, \zeta) \mapsto (\theta, r, h)$. Unsurprisingly, ζ is defined as $\zeta = \frac{i}{M_h}$ and leads to

$$\theta(\xi, \eta, \zeta) = \theta(\xi, \eta) \quad (\text{B.27})$$

$$r(\xi, \eta, \zeta) = r(\xi, \eta) \quad (\text{B.28})$$

$$h(\xi, \eta, \zeta) = \zeta h(\xi, \eta) \quad (\text{B.29})$$

where $\theta(\xi, \eta)$, $r(\xi, \eta)$, and $h(\xi, \eta)$ are defined in Equations B.16, B.17, and B.22, respectively.

For the transition region from the ABH mesh to the plate mesh, the TFI given by Equations B.16 and B.17 are still used, but now using the parametric curves

$$\theta_b(\xi) = \theta(\xi) \quad (\text{B.30})$$

$$r_b(\xi) = R(\theta(\xi)) \quad (\text{B.31})$$

$$\theta_t(\xi) = \begin{cases} \arctan\left(\frac{b_{\text{out}} + a_{\text{out}}}{a_{\text{out}}} \xi\right) & \text{for } 0 \leq \xi \leq \frac{b_{\text{out}}}{b_{\text{out}} + a_{\text{out}}} \\ \arctan\left(\frac{b_{\text{out}}}{(b_{\text{out}} + a_{\text{out}})(1 - \xi)}\right) & \text{for } \frac{b_{\text{out}}}{b_{\text{out}} + a_{\text{out}}} < \xi \leq 1 \end{cases} \quad (\text{B.32})$$

$$r_t(\xi) = \begin{cases} c\sqrt{a_{\text{out}}^2 + (b_{\text{out}} + a_{\text{out}})^2 \xi^2} & \text{for } 0 \leq \xi \leq \frac{b_{\text{out}}}{b_{\text{out}} + a_{\text{out}}} \\ c\sqrt{(b_{\text{out}} + a_{\text{out}})^2 (1 - \xi)^2 + b_{\text{out}}^2} & \text{for } \frac{b_{\text{out}}}{b_{\text{out}} + a_{\text{out}}} < \xi \leq 1 \end{cases} \quad (\text{B.33})$$

$$\theta_l(\eta) = 0 \quad (\text{B.34})$$

$$r_l(\eta) = a_{\text{out}}((c - 1)\eta + 1) \quad (\text{B.35})$$

$$\theta_r(\eta) = \frac{\pi}{2} \quad (\text{B.36})$$

$$r_r(\eta) = b_{\text{out}}((c - 1)\eta + 1) \quad (\text{B.37})$$

where it is assumed that the outer perimeter of the transition region is a $2ca_{\text{out}}$ by $2cb_{\text{out}}$

rectangle, with $c > 1$.

To avoid singular elements, there needs to be a rectangular inner part of the mesh. The transition region to this inner rectangular part of the mesh can be defined similarly to the outer transition region using the parametric curves

$$\theta_b(\xi) = \begin{cases} \arctan\left(\frac{b_{\text{in}}+a_{\text{in}}}{a_{\text{in}}}\xi\right) & \text{for } 0 \leq \xi \leq \frac{b_{\text{in}}}{b_{\text{in}}+a_{\text{in}}} \\ \arctan\left(\frac{b_{\text{in}}}{(b_{\text{in}}+a_{\text{in}})(1-\xi)}\right) & \text{for } \frac{b_{\text{in}}}{b_{\text{in}}+a_{\text{in}}} < \xi \leq 1 \end{cases} \quad (\text{B.38})$$

$$r_b(\xi) = \begin{cases} \frac{1}{c}\sqrt{a_{\text{in}}^2 + (b_{\text{in}} + a_{\text{in}})^2\xi^2} & \text{for } 0 \leq \xi \leq \frac{b_{\text{in}}}{b_{\text{in}}+a_{\text{in}}} \\ \frac{1}{c}\sqrt{(b_{\text{in}} + a_{\text{in}})^2(1-\xi)^2 + b_{\text{in}}^2} & \text{for } \frac{b_{\text{in}}}{b_{\text{in}}+a_{\text{in}}} < \xi \leq 1 \end{cases} \quad (\text{B.39})$$

$$\theta_t(\xi) = \theta(\xi) \quad (\text{B.40})$$

$$r_t(\xi) = r\left(\xi, \frac{1}{M_r}\right) \quad (\text{B.41})$$

$$\theta_l(\eta) = 0 \quad (\text{B.42})$$

$$r_l(\eta) = \frac{a_{\text{in}}}{c}\left((c-1)\eta + 1\right) \quad (\text{B.43})$$

$$\theta_r(\eta) = \frac{\pi}{2} \quad (\text{B.44})$$

$$r_r(\eta) = \frac{b_{\text{in}}}{c}\left((c-1)\eta + 1\right) \quad (\text{B.45})$$

where $a_{\text{in}} = r\left(0, \frac{1}{M_r}\right) = r^a\left(\frac{1}{M_r}\right)$ and $b_{\text{in}} = r\left(\frac{\pi}{2}, \frac{1}{M_r}\right) = r^b\left(\frac{1}{M_r}\right)$. As above, $c > 1$.

Bibliography

- [1] C. A. Coello Coello, “A comprehensive survey of evolutionary-based multiobjective optimization techniques,” *Knowl. Inf. Syst.*, vol. 1, no. 3, pp. 269–308, 1999.
- [2] M. Berggren, U. Laci, F. Lindström, and E. Wadbro, “Sound vibration damping optimization with application to the design of speakerphone casings,” in *Proceedings of the 10th World Congress on Structural and Multidisciplinary Optimization*, p. 5569, 2013.
- [3] V. B. Georgiev, J. Cuenca, F. Gautier, L. Simon, and V. V. Krylov, “Damping of structural vibrations in beams and elliptical plates using the acoustic black hole effect,” *J. Sound Vib.*, vol. 330, no. 11, pp. 2497–2508, 2011.
- [4] L. Cremer and M. Heckl, *Structure-Borne Sound*. Berlin: Springer, 2nd ed., 1988. Eric E. Ungar, trans.
- [5] H. Oberst and K. Frankenfeld, “Über die Dämpfung der Biegeschwingungen dünner Bleche durch fest haftende Beläge,” *Acustica*, vol. 2, no. AB4, pp. 181–194, 1952.
- [6] M. A. Mironov, “Propagation of a flexural wave in a plate whose thickness decreases smoothly to zero in a finite interval,” *Akust. Zh.*, vol. 34, no. 3, pp. 546–547, 1988. [*Sov. Phys. Acoust.*, vol. 34, no. 3, pp. 318–319, 1988].
- [7] V. V. Krylov, “New type of vibration dampers utilising the effect of acoustic ‘black holes’,” *Acta Acust. united Acust.*, vol. 90, no. 5, pp. 830–837, 2004.
- [8] V. V. Krylov and F. J. B. S. Tilman, “Acoustic ‘black holes’ for flexural waves as effective vibration dampers,” *J. Sound Vib.*, vol. 274, no. 3–5, pp. 605–619, 2004.
- [9] V. V. Krylov and R. E. T. B. Winward, “Experimental investigation of the acoustic black hole effect for flexural waves in tapered plates,” *J. Sound Vib.*, vol. 300, no. 1–2, pp. 43–49, 2007.
- [10] E. P. Bowyer, D. J. O’Boy, V. V. Krylov, and F. Gautier, “Experimental investigation of damping flexural vibrations in plates containing tapered indentations of power-law profile,” *Appl. Acoust.*, vol. 74, no. 4, pp. 553–560, 2013.

- [11] E. P. Bowyer and V. V. Krylov, “Experimental study of sound radiation by plates containing circular indentations of power-law profile,” *Appl. Acoust.*, vol. 88, pp. 30–37, 2015.
- [12] S. C. Conlon, J. B. Fahnlone, and F. Semperlotti, “Numerical analysis of the vibroacoustic properties of plates with embedded grids of acoustic black holes,” *J. Acoust. Soc. Am.*, vol. 137, no. 1, pp. 447–457, 2015.
- [13] L. Tang and L. Cheng, “Enhanced acoustic black hole effect in beams with a modified thickness profile and extended platform,” *J. Sound Vib.*, vol. 391, pp. 116–126, 2017.
- [14] A. Climente, D. Torrent, and J. Sánchez-Dehesa, “Gradient index lenses for flexural waves based on thickness variations,” *Appl. Phys. Lett.*, vol. 105, no. 6, p. 064101, 2014.
- [15] W. Huang, H. Ji, J. Qiu, and L. Cheng, “Wave energy focalization in a plate with imperfect two-dimensional acoustic black hole indentation,” *J. Vib. Acoust.*, vol. 138, no. 6, p. 061004, 2016.
- [16] W. Huang, H. Ji, J. Qiu, and L. Cheng, “Analysis of ray trajectories of flexural waves propagating over generalized acoustic black hole indentations,” *J. Sound Vib.*, vol. 417, pp. 216–226, 2018.
- [17] E. P. Bower, D. J. O’Boy, V. V. Krylov, and J. L. Horner, “Effect of geometrical and material imperfections on damping flexural vibrations in plates with attached wedges of power law profile,” *Appl. Acoust.*, vol. 73, no. 5, pp. 514–523, 2012.
- [18] E. P. Bowyer, “Acoustic black hole manufacturing for practical applications and the effect of geometrical and material imperfection,” in *INTER-NOISE and NOISE-CON Congress and Conference Proceedings*, vol. 253, pp. 3924–3934, 2016.
- [19] V. Denis, A. Pelat, and F. Gautier, “Scattering effects induced by imperfections on an acoustic black hole placed at a structural waveguide termination,” *J. Sound Vib.*, vol. 362, pp. 56–71, 2016.
- [20] V. Denis, C. Touzé, and F. Gautier, “Effects of geometrical nonlinearities on the acoustic black hole effect,” in *INTER-NOISE and NOISE-CON Congress and Conference Proceedings*, vol. 253, pp. 4431–4441, 2016.
- [21] P. A. Feurtado and S. C. Conlon, “Wavenumber transform analysis for acoustic black hole design,” *J. Acoust. Soc. Am.*, vol. 140, no. 1, pp. 718–727, 2016.
- [22] S. C. Conlon and P. A. Feurtado, “Progressive phase trends in plates with embedded acoustic black holes,” *J. Acoust. Soc. Am.*, vol. 143, no. 2, pp. 921–930, 2018.

- [23] O. Unruh, C. Blech, and H. P. Monner, “Numerical and experimental study of sound power reduction performance of acoustic black holes in rectangular plates,” *SAE Int. J. Passeng. Cars - Mech. Syst.*, vol. 8, no. 3, pp. 956–963, 2015.
- [24] J.-G. Ih, M. Kim, I.-J. Lee, and J. S. Jensen, “Truncated acoustic black hole structure with the optimized tapering shape and damping coating,” in *INTER-NOISE and NOISE-CON Congress and Conference Proceedings*, vol. 253, pp. 6096–6103, 2016.
- [25] P. Feurtado and S. Conlon, “Investigation of boundary-taper reflection for acoustic black hole design,” *Noise Control Eng. J.*, vol. 63, no. 5, pp. 460–466, 2015.
- [26] M. R. Shepherd, P. A. Feurtado, and S. C. Conlon, “Multi-objective optimization of acoustic black hole vibration absorbers,” *J. Acoust. Soc. Am.*, vol. 140, no. 3, pp. EL227–EL230, 2016.
- [27] P. A. Feurtado and S. C. Conlon, “An experimental investigation of acoustic black hole dynamics at low, mid, and high frequencies,” *J. Vib. Acoust.*, vol. 138, no. 6, p. 061002, 2016.
- [28] D. Ross, E. E. Ungar, and E. M. Kerwin Jr., “Damping of plate vibrations by viscoelastic laminae,” in *Structural Damping* (J. E. Ruzicka, ed.), ch. 3, pp. 49–87, New York: ASME, 1959.
- [29] M. Dorn, C. Blech, and S. C. Langer, “Numerical studies on the optimal exploitation of the acoustic black hole effect in curved panels,” in *INTER-NOISE and NOISE-CON Congress and Conference Proceedings*, vol. 255, pp. 4194–4204, 2017.
- [30] X. Wang, X. Liu, J. Tai, T. He, and Y. Shan, “A novel method of reducing the acoustic emission wave reflected by boundary based on acoustic black hole,” *Ultrasonics*, vol. 94, pp. 292–304, 2019.
- [31] T. Kundu, “Acoustic source localization,” *Ultrasonics*, vol. 54, no. 1, pp. 25–38, 2014.
- [32] K. Hook, J. Cheer, and S. Daley, “A parametric study and modal analysis of an acoustic black hole on a beam,” *J. Acoust. Soc. Am.*, vol. 145, no. 6, pp. 3488–3498, 2019.
- [33] V. Denis, A. Pelat, F. Gautier, and B. Elie, “Modal Overlap Factor of a beam with an acoustic black hole termination,” *J. Sound Vib.*, vol. 333, no. 12, pp. 2475–2488, 2014.
- [34] M. Ouisse, D. Renault, P. Butaud, and E. Sadoulet-Reboul, “Damping control for improvement of acoustic black hole effect,” *J. Sound Vib.*, vol. 454, pp. 63–72, 2019.

- [35] S. Roth, V. Ghaffari Mejlaj, S. C. Langer, and T. Vietor, “Optimal adaptation of acoustic black holes by evolutionary optimization algorithms,” *Proc. Appl. Math. Mech.*, vol. 16, no. 1, pp. 625–626, 2016.
- [36] L. Ma and L. Cheng, “Topological optimization of damping layout for minimized sound radiation of an acoustic black hole plate,” *J. Sound Vib.*, vol. 458, pp. 349–364, 2019.
- [37] L. Ma, H.-W. Dong, and L. Cheng, “An alternative and optimized thickness profile of an acoustic black hole plate,” *J. Sound Vib.*, vol. 486, p. 115619, 2020.
- [38] J. Deng, O. Guasch, and L. Zheng, “Ring-shaped acoustic black holes for broadband vibration isolation in plates,” *J. Sound Vib.*, vol. 458, pp. 109–122, 2019.
- [39] A. Karlos, S. J. Elliott, and J. Cheer, “Higher-order WKB analysis of reflection from tapered elastic wedges,” *J. Sound Vib.*, vol. 449, pp. 368–388, 2019.
- [40] P. A. Feurtado, S. C. Conlon, and F. Semperlotti, “A normalized wave number variation parameter for acoustic black hole design,” *J. Acoust. Soc. Am.*, vol. 136, no. 2, pp. EL148–EL152, 2014.
- [41] C. A. Coello Coello, G. B. Lamont, and D. A. van Veldhuizen, *Evolutionary Algorithms for Solving Multi-Objective Problems*. Boston: Springer, 2nd ed., 2007.
- [42] J. H. Holland, *Adaptation in Natural and Artificial Systems*. Cambridge, MA: MIT Press, 1992.
- [43] R. Salomon, “Evolutionary algorithms and gradient search: Similarities and differences,” *IEEE Trans. Evolut. Comput.*, vol. 2, no. 2, pp. 45–55, 1998.
- [44] D. Hadka and P. Reed, “Borg: An auto-adaptive many-objective evolutionary computing framework,” *Evol. Comput.*, vol. 21, no. 2, pp. 231–259, 2013.
- [45] V. F. D. Pareto, *Manuale di Economia Politica: Con una Introduzione alla Scienza Sociale [Manual of Political Economy: With an Introduction to Social Science]*, vol. 13, pp. 164–165. Milan: Societa Editrice Libreria, 1906.
- [46] G. Rudolph, “Evolutionary search for minimal elements in partially ordered finite sets,” in *Evolutionary Programming VII* (V. W. Porto, N. Saravanan, D. Waagen, and A. E. Eiben, eds.), pp. 345–353, 1998.
- [47] M. Laumanns, L. Thiele, K. Deb, and E. Zitzler, “Combining convergence and diversity in evolutionary multiobjective optimization,” *Evol. Comput.*, vol. 10, no. 3, pp. 263–282, 2002.
- [48] J. A. Vrugt and B. A. Robinson, “Improved evolutionary optimization from genetically adaptive multimethod search,” in *Proceedings of the National Academy of Sciences*, vol. 104, pp. 708–711, 2007.

- [49] K. Deb and R. B. Agrawal, “Simulated binary crossover for continuous search space,” *Complex Syst.*, vol. 9, no. 2, pp. 115–148, 1995.
- [50] R. Storn and K. Price, “Differential Evolution - a simple and efficient heuristic for global optimization over continuous spaces,” *J. Global Optim.*, vol. 11, no. 4, pp. 341–359, 1997.
- [51] K. Deb, D. Joshi, and A. Anand, “Real-coded evolutionary algorithms with parent-centric recombination,” in *Proceedings of the 2002 Congress on Evolutionary Computation (CEC’02)*, vol. 1, pp. 61–66, 2002.
- [52] H. Kita, I. Ono, and S. Kobayashi, “Multi-parental extension of the unimodal normal distribution crossover for real-coded genetic algorithms,” in *Proceedings of the 1999 Congress on Evolutionary Computation (CEC’99)*, vol. 2, pp. 1581–1588, 1999.
- [53] S. Tsutsui, M. Yamamura, and T. Higuchi, “Multi-parent recombination with simplex crossover in real coded genetic algorithms,” in *Proceedings of the 1st Annual Conference on Genetic and Evolutionary Computation (GECCO’99)*, vol. 1, pp. 657–664, 1999.
- [54] M. J. Woodruff, P. M. Reed, and T. W. Simpson, “Many objective visual analytics: Rethinking the design of complex engineered systems,” *Struct. Multidiscip. O.*, vol. 48, no. 1, pp. 201–219, 2013.
- [55] R. Singh, P. M. Reed, and K. Keller, “Many-objective robust decision making for managing an ecosystem with a deeply uncertain threshold response,” *Ecol. Soc.*, vol. 20, no. 3, p. 12, 2015.
- [56] M. Giuliani, J. D. Quinn, J. D. Herman, A. Castelletti, and P. M. Reed, “Scalable multiobjective control for large-scale water resources systems under uncertainty,” *IEEE Trans. Control Syst. Tech.*, vol. 26, no. 4, pp. 1492–1499, 2018.
- [57] M. R. Shepherd, R. L. Campbell, and S. A. Hambric, “A parallel computing framework for performing structural-acoustic optimization with stochastic forcing,” *Struct. Multidiscip. O.*, vol. 61, no. 2, pp. 675–685, 2020.
- [58] R. S. Gupta, A. L. Hamilton, P. M. Reed, and G. W. Characklis, “Can modern multi-objective evolutionary algorithms discover high-dimensional financial risk portfolio tradeoffs for snow-dominated water-energy systems?,” *Adv. Water Resour.*, vol. 145, p. 103718, 2020.
- [59] W. Bangerth, T. Heister, L. Heltai, G. Kanschat, M. Kronbichler, M. Maier, and B. Turcksin, “The deal.II library, version 8.3,” *Arch. Numer. Soft.*, vol. 4, no. 100, pp. 1–11, 2016.

- [60] C. A. McCormick and M. Shepherd, “A comparison of algorithms for the vibroacoustic optimization of a beam: Gradient-based versus evolutionary,” *J. Acoust. Soc. Am.*, vol. 148, no. 4, p. 2604, 2020.
- [61] T. J. R. Hughes, *The Finite Element Method*, ch. 4.5.2, pp. 232–237. Mineola, NY: Dover, 2000.
- [62] M. P. Bendsøe and O. Sigmund, *Topology Optimization: Theory, Methods, and Applications*, pp. 34–35. Berlin: Springer, 2nd ed., 2004.
- [63] J. Fliege, L. M. Graña Drummond, and B. F. Svaiter, “Newton’s method for multiobjective optimization,” *SIAM J. Optim.*, vol. 20, no. 2, pp. 602–626, 2009.
- [64] J. Fliege and B. F. Svaiter, “Steepest descent methods for multicriteria optimization,” *Math Meth. Oper. Res.*, vol. 51, pp. 479–494, 2000.
- [65] C. A. McCormick and M. R. Shepherd, “Optimization of an acoustic black hole vibration absorber at the end of a cantilever beam,” *J. Acoust. Soc. Am.*, vol. 145, no. 6, pp. EL593–EL597, 2019.
- [66] J. W. Dunkin, “Computation of modal solutions in layered, elastic media at high frequencies,” *Bull. Seismol. Soc. Am.*, vol. 55, no. 2, pp. 335–358, 1965.
- [67] P. Sautet and C. Joachim, “Electronic transmission coefficient for the single-impurity problem in the scattering-matrix approach,” *Phys. Rev. B*, vol. 38, no. 17, pp. 12238–12247, 1988.
- [68] C. C. Katsidis and D. I. Siapkas, “General transfer-matrix method for optical multilayer systems with coherent, partially coherent, and incoherent interference,” *Appl. Opt.*, vol. 41, no. 19, pp. 3978–3987, 2002.
- [69] B. Jonsson and S. T. Eng, “Solving the schrödinger equation in arbitrary quantum-well potential profiles using the transfer matrix method,” *IEEE J. Quantum Electron.*, vol. 26, no. 11, pp. 2025–2035, 1990.
- [70] D. Y. K. Ko and J. R. Sambles, “Scattering matrix method for propagation of radiation in stratified media: attenuated total reflection studies of liquid crystals,” *J. Opt. Soc. Am. A*, vol. 5, no. 11, pp. 1863–1866, 1988.
- [71] M. G. Moharam, D. A. Pommet, E. B. Grann, and T. K. Gaylord, “Stable implementation of the rigorous coupled-wave analysis for surface-relief gratings: enhanced transmittance matrix approach,” *J. Opt. Soc. Am. A*, vol. 12, no. 5, pp. 1077–1086, 1995.
- [72] L. Wang and S. I. Rokhlin, “Stable reformulation of transfer matrix method for wave propagation in layered anisotropic media,” *Ultrasonics*, vol. 39, no. 6, pp. 413–424, 2001.

- [73] G. C. Horner and W. D. Pilkey, “The Riccati transfer matrix method,” *J. Mech. Design*, vol. 100, no. 2, pp. 297–302, 1978.
- [74] G. B. Rybicki and P. D. Usher, “The generalized Riccati transformation as a simple alternative to invariant imbedding,” *Astrophys. J.*, vol. 146, no. 3, pp. 871–879, 1966.
- [75] S. Bittanti, “Count Riccati and the early days of the Riccati equation,” in *The Riccati Equation* (S. Bittanti, A. J. Laub, and J. C. Willems, eds.), ch. 1, pp. 1–10, Berlin: Springer, 1991.
- [76] C. A. McCormick and M. R. Shepherd, “Design optimization and performance comparison of three styles of one-dimensional acoustic black hole vibration absorbers,” *J. Sound Vib.*, vol. 470, p. 115164, 2020.
- [77] E. P. Bowyer and V. V. Krylov, “Experimental investigation of damping flexural vibrations in glass fibre composite plates containing one- and two-dimensional acoustic black holes,” *Compos. Struct.*, vol. 107, pp. 406 – 415, 2014.
- [78] T. Zhou, L. Tang, H. Ji, J. Qiu, and L. Cheng, “Dynamic and static properties of double-layered compound acoustic black hole structures,” *Int. J. Appl. Mech.*, vol. 09, no. 05, p. 1750074, 2017.
- [79] X. Li and Q. Ding, “Analysis on vibration energy concentration of the one-dimensional wedge-shaped acoustic black hole structure,” *J. Intel. Mat. Syst. Str.*, vol. 29, no. 10, pp. 2137–2148, 2018.
- [80] P. Zeng, L. Zheng, J. Deng, A. Elsabbagh, S. Xiang, T. Yan, and Y. Wu, “Flexural wave concentration in tapered cylindrical beams and wedge-like rectangular beams with power-law thickness,” *J. Sound Vib.*, vol. 452, pp. 82 – 96, 2019.
- [81] Y. Wang, J. Du, and L. Cheng, “Power flow and structural intensity analyses of acoustic black hole beams,” *Mech. Syst. Signal Pr.*, vol. 131, pp. 538 – 553, 2019.
- [82] V. Kralovic and V. V. Krylov, “Damping of flexural vibrations in tapered rods of power-law profile: Experimental studies,” in *Proceedings of the Institute of Acoustics*, vol. 29 (Part 5), pp. 66–73, Institute of Acoustics, 2007.
- [83] H. Ji, J. Luo, J. Qiu, and L. Cheng, “Investigations on flexural wave propagation and attenuation in a modified one-dimensional acoustic black hole using a laser excitation technique,” *Mech. Syst. Signal Pr.*, vol. 104, pp. 19 – 35, 2018.
- [84] X. Du, D. Huang, and J. Zhang, “Dynamic property investigation of sandwich acoustic black hole beam with clamped-free boundary condition,” *Shock Vib.*, vol. 2019, 2019. Article ID 6708138.

- [85] V. Denis, A. Pelat, C. Touzé, and F. Gautier, “Improvement of the acoustic black hole effect by using energy transfer due to geometric nonlinearity,” *Int. J. Nonlin. Mech.*, vol. 94, pp. 134 – 145, 2017.
- [86] H. Li, C. Touzé, A. Pelat, F. Gautier, and X. Kong, “A vibro-impact acoustic black hole for passive damping of flexural beam vibrations,” *J. Sound Vib.*, vol. 450, pp. 28 – 46, 2019.
- [87] L. Zhao, S. C. Conlon, and F. Semperlotti, “Broadband energy harvesting using acoustic black hole structural tailoring,” *Smart Mater. Struct.*, vol. 23, p. 065021, May 2014.
- [88] Y. Liang, H. Ji, J. Qiu, L. Cheng, Y. Wu, and C. Zhang, “Enhancement of piezo-electric energy harvesting using ABH structural tailoring,” in *2018 IEEE/ASME International Conference on Advanced Intelligent Mechatronics (AIM)*, pp. 1372–1377, July 2018.
- [89] C. Zhao and M. G. Prasad, “Studies on vibration energy harvesting using a cantilever beam with a modified acoustic black hole cavity,” *INTER-NOISE and NOISE-CON Congress and Conference Proceedings*, vol. 258, no. 4, pp. 3755–3761, 2018.
- [90] Y. Mi and H. Zheng, “Enhancement of acoustic black hole effect in beams using shunted piezo-electric patch,” *INTER-NOISE and NOISE-CON Congress and Conference Proceedings*, vol. 259, no. 9, pp. 242–248, 2019.
- [91] K. Hook, J. Cheer, and S. Daley, “Optimal feedforward control of a beam with an active acoustic black hole termination,” in *26th International Congress on Sound and Vibration (11/07/19)*, July 2019.
- [92] J. Deng and L. Zheng, “Exploration of admissible shape functions in semi-analytical model to characterize one-dimensional acoustic black hole beams,” *INTER-NOISE and NOISE-CON Congress and Conference Proceedings*, vol. 257, no. 1, pp. 162–172, 2018.
- [93] T. Zhou, J.-D. Chazot, E. Perrey-Debain, and L. Cheng, “Frequency responses of acoustic black hole wedges solved by the partition of unity finite element method,” *INTER-NOISE and NOISE-CON Congress and Conference Proceedings*, vol. 259, no. 6, pp. 3170–3179, 2019.
- [94] Y. Xiong, E. Smith, and S. C. Conlon, “An analytical model for investigating damping effects in beams with acoustic black holes,” *INTER-NOISE and NOISE-CON Congress and Conference Proceedings*, vol. 260, no. 1, pp. 614–621, 2019.
- [95] X. Li and Q. Ding, “Sound radiation of a beam with a wedge-shaped edge embedding acoustic black hole feature,” *J. Sound Vib.*, vol. 439, pp. 287 – 299, 2019.

- [96] N. Gao, Z. Wei, H. Hou, and A. O. Krushynska, “Design and experimental investigation of V-folded beams with acoustic black hole indentations,” *J. Acoust. Soc. Am.*, vol. 145, no. 1, pp. EL79–EL83, 2019.
- [97] N. Gao, Z. Wei, R. Zhang, and H. Hou, “Low-frequency elastic wave attenuation in a composite acoustic black hole beam,” *Appl. Acoust.*, vol. 154, pp. 68 – 76, 2019.
- [98] Y. Zhang, K. Chen, S. Zhou, and Z. Wei, “An ultralight phononic beam with a broad low-frequency band gap using the complex lattice of acoustic black holes,” *Appl. Phys. Express*, vol. 12, p. 077002, July 2019.
- [99] N. Gao, X. Guo, B. Cheng, Y. Zhang, Z. Wei, and H. Hou, “Elastic wave modulation in hollow metamaterial beam with acoustic black hole,” *IEEE Access*, vol. 7, pp. 124141–124146, 2019.
- [100] B. M. P. Chong, L. B. Tan, K. M. Lim, and H. P. Lee, “A review on acoustic black-holes (ABH) and the experimental and numerical study of ABH-featured 3D printed beams,” *Int. J. Appl. Mech.*, vol. 09, no. 06, p. 1750078, 2017.
- [101] S. Rothe, C. Blech, H. Watschke, T. Vietor, and S. C. Langer, “Numerical modelling of additively manufactured acoustic black holes,” in *Proceedings of the Inter-Noise 2018 47th International Congress and Exposition on Noise Control Engineering, Chicago, IL, USA*, pp. 26–29, 2018.
- [102] W. Huang, H. Zhang, D. J. Inman, J. Qiu, C. E. S. Cesnik, and H. Ji, “Low reflection effect by 3D printed functionally graded acoustic black holes,” *J. Sound Vib.*, vol. 450, pp. 96 – 108, 2019.
- [103] J. Y. Lee and W. Jeon, “Vibration damping using a spiral acoustic black hole,” *J. Acoust. Soc. Am.*, vol. 141, no. 3, pp. 1437–1445, 2017.
- [104] J. Y. Lee and W. Jeon, “Cut-off frequency for curvilinear acoustic black holes,” *INTER-NOISE and NOISE-CON Congress and Conference Proceedings*, vol. 259, no. 6, pp. 3697–3700, 2019.
- [105] S. Park, M. Kim, and W. Jeon, “Experimental validation of vibration damping using an Archimedean spiral acoustic black hole,” *J. Sound Vib.*, vol. 459, p. 114838, 2019.
- [106] C. A. McCormick and M. R. Shepherd, “Optimal design and position of an embedded one-dimensional acoustic black hole,” *INTER-NOISE and NOISE-CON Congress and Conference Proceedings*, vol. 258, no. 6, pp. 1345–1354, 2018.
- [107] V. Denis, F. Gautier, A. Pelat, and J. Poittevin, “Measurement and modelling of the reflection coefficient of an acoustic black hole termination,” *J. Sound Vib.*, vol. 349, pp. 67 – 79, 2015.

- [108] J. Leng, V. Romero, J.-P. Groby, A. Pelat, R. Pico, and F. Gautier, “The use of perfect absorption in the tunability of the resonant modes of an acoustic black hole,” *INTER-NOISE and NOISE-CON Congress and Conference Proceedings*, vol. 258, no. 3, pp. 4869–4879, 2018.
- [109] M. R. Shepherd, C. A. McCormick, S. C. Conlon, and P. A. Feurtado, “Modeling and optimization of acoustic black hole vibration absorbers,” *J. Acoust. Soc. Am.*, vol. 141, no. 5, pp. 4034–4034, 2017.
- [110] L. Tang and L. Cheng, “Broadband locally resonant band gaps in periodic beam structures with embedded acoustic black holes,” *J. Appl. Phys.*, vol. 121, no. 19, p. 194901, 2017.
- [111] L. Tang and L. Cheng, “Ultrawide band gaps in beams with double-leaf acoustic black hole indentations,” *J. Acoust. Soc. Am.*, vol. 142, no. 5, pp. 2802–2807, 2017.
- [112] Y. Xiong, E. Smith, and S. C. Conlon, “Vibroacoustic properties of plates with tuned acoustic black holes,” *INTER-NOISE and NOISE-CON Congress and Conference Proceedings*, vol. 258, no. 4, pp. 3637–3650, 2018.
- [113] D. J. O’Boy and V. V. Krylov, “Damping of flexural vibrations in circular plates with tapered central holes,” *J. Sound Vib.*, vol. 330, no. 10, pp. 2220 – 2236, 2011.
- [114] E. D. Williams and J. D. Maynard, “Numerical evaluation of the Rayleigh integral for planar radiators using the FFT,” *J. Acoust. Soc. Am.*, vol. 72, no. 6, pp. 2020–2030, 1982.
- [115] J. Deng, L. Zheng, P. Zeng, Y. Zuo, and O. Guasch, “Passive constrained viscoelastic layers to improve the efficiency of truncated acoustic black holes in beams,” *Mech. Syst. Signal Pr.*, vol. 118, pp. 461 – 476, 2019.
- [116] W. Huang, H. Ji, C. Tao, and J. Qiu, “Optimal design of the damping layer in plate with imperfect Acoustic Black Hole for wave energy dissipation,” in *INTER-NOISE and NOISE-CON Congress and Conference Proceedings*, vol. 259, pp. 7763–7771, 2019.
- [117] J. T. Xing and W. G. Price, “A power-flow analysis based on continuum dynamics,” *Proc. R. Soc. Lond. A*, vol. 455, no. 1982, pp. 401–436, 1999.
- [118] “The deal.II library documentation.” <https://www.dealii.org/8.3.0/doxygen/deal.II/index.html>. [Online; accessed October 18, 2021].
- [119] “LAPACK — Linear Algebra PACKage.” <https://www.netlib.org/lapack/>. [Online; accessed October 18, 2021].
- [120] “ScaLAPACK — Scalable Linear Algebra PACKage.” <https://www.netlib.org/scalapack/>. [Online; accessed October 18, 2021].

- [121] “ARPACK homepage.” <https://www.caam.rice.edu/software/ARPACK/>. [Online; accessed October 18, 2021].
- [122] “PETSc — Portable Extensible Toolkit for Scientific Computation.” <https://www.mcs.anl.gov/petsc/>. [Online; accessed October 18, 2021].
- [123] “Trilinos home page.” <https://trilinos.github.io/>. [Online; accessed October 18, 2021].
- [124] “p4est: Parallel AMR on forests of octrees.” <https://www.p4est.org/>. [Online; accessed October 18, 2021].
- [125] “SuperLU (Supernodal LU).” <https://portal.nersc.gov/project/sparse/superlu/>. [Online; accessed October 18, 2021].
- [126] G. Xie, D. J. Thompson, and C. J. C. Jones, “Mode count and modal density of structural systems: relationships with boundary conditions,” *J. Sound Vib.*, vol. 274, no. 3–5, pp. 621–651, 2004.
- [127] E. Q. Sun, “Shear locking and hourglassing in MSC Nastran, ABAQUS, and ANSYS,” in *MSC Software’s Virtual Product Development Conference*, 2006.
- [128] C. Taylor and P. Hood, “A numerical solution of the Navier-Stokes equations using the finite element technique,” *Comp. Fluids*, vol. 1, no. 1, pp. 73–100, 1973.
- [129] F. Brezzi, “On the existence, uniqueness and approximation of saddle-point problems arising from Lagrangian multipliers,” *RAIRO Numer. Anal.*, vol. 8, no. R2, pp. 129–151, 1974.
- [130] R. H. MacNeal, *Finite Elements: Their Design and Performance*, ch. 7.6.2, pp. 293–295. New York, Basel, Hong Kong: Marcel Dekker, 1994.
- [131] W. Huang, C. Tao, H. Ji, and J. Qiu, “Enhancement of wave energy dissipation in two-dimensional acoustic black hole by simultaneous optimization of profile and damping layer,” *J. Sound Vib.*, vol. 491, p. 115764, 2021.
- [132] P. A. Feurtado and S. C. Conlon, “Transmission loss of plates with embedded acoustic black holes,” *J. Acoust. Soc. Am.*, vol. 142, no. 3, pp. 1390–1398, 2017.
- [133] M. R. Shepherd, P. A. Feurtado, and S. C. Conlon, “Wavenumber interactions of turbulent boundary layer flow with structures exhibiting the acoustic black hole effect,” *Proc. Mtgs. Acoust.*, vol. 39, no. 1, p. 065002, 2019.
- [134] P. Knupp and S. Steinberg, *Fundamentals of Grid Generation*. CRC Press, 1993.
- [135] R. L. Kuhlemeyer and J. Lysmer, “Finite element method accuracy for wave propagation problems,” *J. Soil Mech. Found. Div.*, vol. 99, no. 5, pp. 421–427, 1973.

- [136] S. Marburg, “Discretization requirements: How many elements per wavelength are necessary?,” in *Computational Acoustics of Noise Propagation in Fluids - Finite and Boundary Element Methods* (S. Marburg and B. Nolte, eds.), pp. 309–332, Berlin, Heidelberg: Springer Berlin Heidelberg, 2008.
- [137] S. Marburg, “Normal modes in external acoustics. Part I: Investigation of the one-dimensional duct problem,” *Acta Acust. unite Acust.*, vol. 91, no. 6, pp. 1063–1078, 2005.
- [138] P. Langer, M. Mäder, C. Guist, M. Krause, and S. Marburg, “More than six elements per wavelength: The practical use of structural finite element models and their accuracy in comparison with experimental results,” *J. Comp. Acous.*, vol. 25, no. 04, p. 1750025, 2017.

Vita

Cameron A. McCormick

Born and raised in the San Francisco Bay Area, Cam received a B.S. in Physics in 2014 from Loyola University New Orleans, where he researched the chemical and crystallographic properties of cyano-anilinium and cyano-methyl-pyridinium salts. In 2015, he received an MSc in Sound and Vibration Studies from the ISVR in Southampton, England, where he developed an acoustic metamaterial comprised of Helmholtz resonators utilizing active noise control.