

A comparison of algorithms for the vibroacoustic optimization of a beam: Gradient-based versus evolutionary

Cameron A. McCormick and Micah R. Shepherd

Citation: *Proc. Mtgs. Acoust.* **42**, 022005 (2020); doi: 10.1121/2.0001545

View online: <https://doi.org/10.1121/2.0001545>

View Table of Contents: <https://asa.scitation.org/toc/pma/42/1>

Published by the [Acoustical Society of America](#)

ARTICLES YOU MAY BE INTERESTED IN

[Machine learning in acoustics: Theory and applications](#)

The Journal of the Acoustical Society of America **146**, 3590 (2019); <https://doi.org/10.1121/1.5133944>

[A kiloelectron-volt ultrafast electron micro-diffraction apparatus using low emittance semiconductor photocathodes](#)

Structural Dynamics **9**, 024302 (2022); <https://doi.org/10.1063/4.0000138>

[The statistical estimates of sensitivity in the problems of intensimetric and correlation acoustic tomography](#)

Proceedings of Meetings on Acoustics **43**, 055001 (2021); <https://doi.org/10.1121/2.0001560>

[Ray method for complete system of acoustic equations with simple caustic correction](#)

Proceedings of Meetings on Acoustics **45**, 070006 (2021); <https://doi.org/10.1121/2.0001556>

[Summary of Acoustics-in-Focus Panel Discussion Special Session "Underwater Sound Modeling: Comprehensive Environmental Descriptions"](#)

Proceedings of Meetings on Acoustics **43**, 002001 (2021); <https://doi.org/10.1121/2.0001469>

[The calculation of acoustic field in a metamaterial and determination of its inner structure using scattering coefficients](#)

Proceedings of Meetings on Acoustics **42**, 065005 (2020); <https://doi.org/10.1121/2.0001514>



Why Publish in POMA?

Watch Now 

179th Meeting of the Acoustical Society of America*Acoustics Virtually Everywhere*

7-11 December 2020

Computational Acoustics: Paper 3aCAb1

A comparison of algorithms for the vibroacoustic optimization of a beam: Gradient-based versus evolutionary**Cameron A. McCormick and Micah R. Shepherd***Applied Research Laboratory, Graduate Program in Acoustics, The Pennsylvania State University, State College, PA, 16802; c.arthur.mccormick@gmail.com; mrs30@arl.psu.edu*

Two common classes of optimization algorithm used in structural and/or multidisciplinary optimization are gradient-based algorithms and evolutionary algorithms. In the case of smooth, unimodal objective spaces, gradient-based algorithms are generally faster, requiring fewer iterations to reach a solution. Evolutionary algorithms, on the other hand, are more robust against objective spaces that are nonlinear, discontinuous, and/or multimodal. This talk will present an optimization of the thickness distribution of a cantilever beam, inspired by a similar study carried out by Berggren et al. ["Sound vibration damping optimization with application to the design of speakerphone casings," 10th World Congress on Structural and Multidisciplinary Optimization, 2013]. The objective is to minimize the vibration response within a certain region of the beam at discrete frequencies, with constraints on total mass and static compliance. The objective space is expected to be nonlinear and potentially multimodal. A transfer matrix method is used to evaluate the objective function and constraints, and optimal solutions are found using both a gradient-based algorithm and an evolutionary algorithm. Qualitative and quantitative results will be presented in comparing the optimized distribution to that of Berggren et al. and in discussing the benefits and limitations of the two algorithms for vibroacoustic optimization.

1. INTRODUCTION

When performing structural-acoustic optimization, it is important 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), which fall into the latter category, only require the values of the objective function and therefore can be used as a black box. However, when used within a robust structural optimization framework, they can be very effective for difficult search spaces.^{1,2}

The objectives of the work presented in this research are twofold. The first is to describe a general 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 a specific EA (referred to as Borg³) 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 for structural-acoustic optimization problems.

2. PROBLEM DESCRIPTION

The chosen test problem is taken with minimal modification from Berggren et al.,⁴ 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. The tilde indicates that the variable is a complex quantity. 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 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 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:

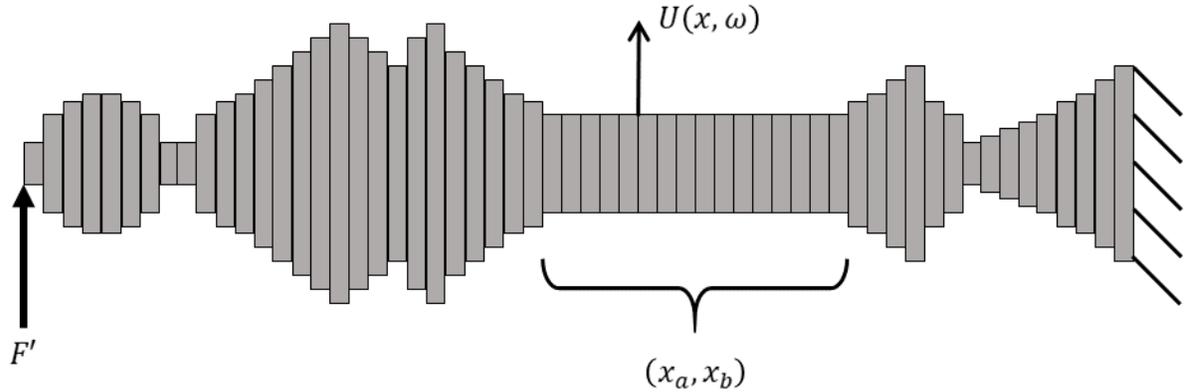


Figure 1: Graphical depiction of the current optimization problem. The beam geometry is represented by 100 beam segments, where each segment has constant thickness. The left-hand boundary condition is an imposed shear, while the right-hand boundary condition is zero displacement and zero rotation, i.e., fixed. The microphone region, (x_a, x_b) , has constant thickness h_{mic} throughout.

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

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, given by h^- and h^+ respectively. $J(\mathbf{h})$ 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} . Additionally, $f(x, \omega)$ is the external loading as a function of position and frequency while F' is a scaled shear force within the beam. 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. The beam vibration was determined using the transfer matrix method.⁵

3. 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. 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 1. As in Berggren,⁴ three different optimizations were carried out, one for each of three sets of frequencies. The first case is broadband optimization, $\omega \in \mathcal{W}_{\text{BB}}$ which is a frequency range from 300 Hz to 3400 Hz with 50 equally-spaced frequencies within this range. Two other cases include 300 Hz to 800 Hz and 2300 Hz to 2800 Hz. From each of these frequency ranges, 50 equally-spaced evaluation frequencies were again used 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. 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`. It should be noted 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} .

4. 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.⁴ are presented, both for reference and to compare against the results of this current work.

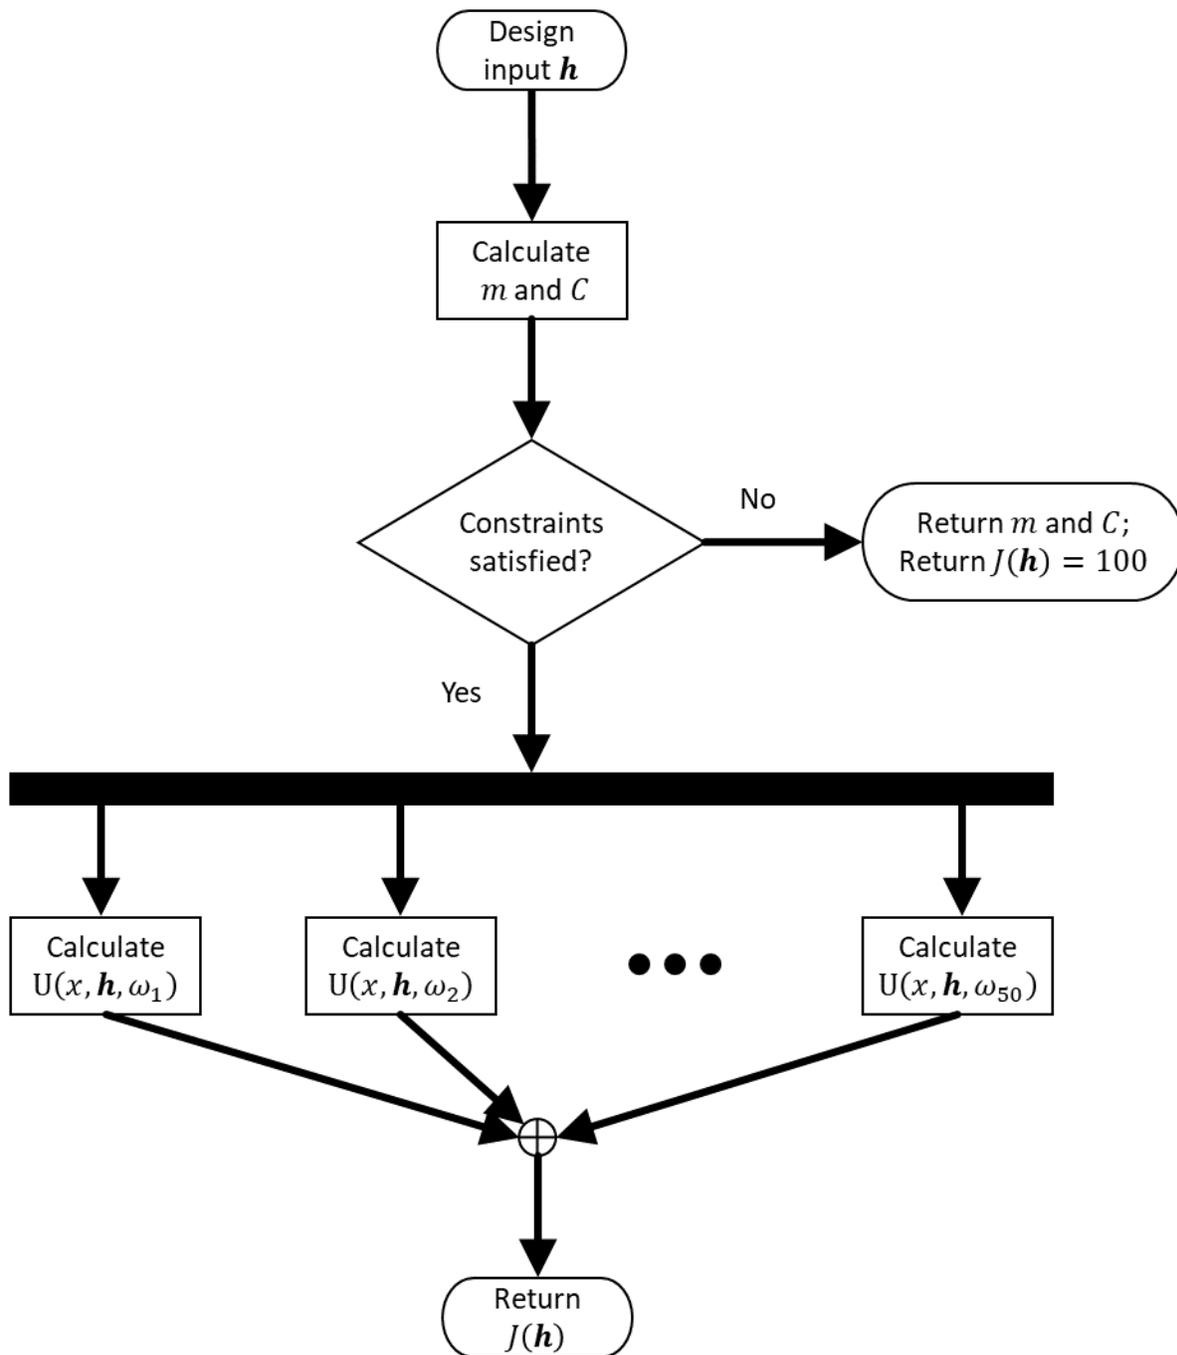


Figure 2: Flowchart representation of the design evaluation function called by both optimization algorithms.

Table 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
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.

A. $\omega \in \mathcal{W}_{\text{BB}}$

Figures 3 and 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 from⁴ is reproduced in Figure 5 for reference. To facilitate direct comparison, the optimization results are formatted in the same way as in⁴ for each of the three frequency sets. Taking Figure 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. 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 3 and 5, the problem does not include damping and so the displacement at resonance should theoretically be infinite. Therefore, the peak heights in the figures reproduced from⁴ are artificially low, possibly due to undersampling. Looking again at Figure 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 3 and 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 5. The authors' results are more similar to the results of `fmincon`, with the notable exception that the results of Berggren et al. show a thickening of the beam near the clamped end, while the `fmincon` favored distributing the mass to other locations.

B. $\omega \in \mathcal{W}_{\text{LF}}$

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 6, 7, and 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.

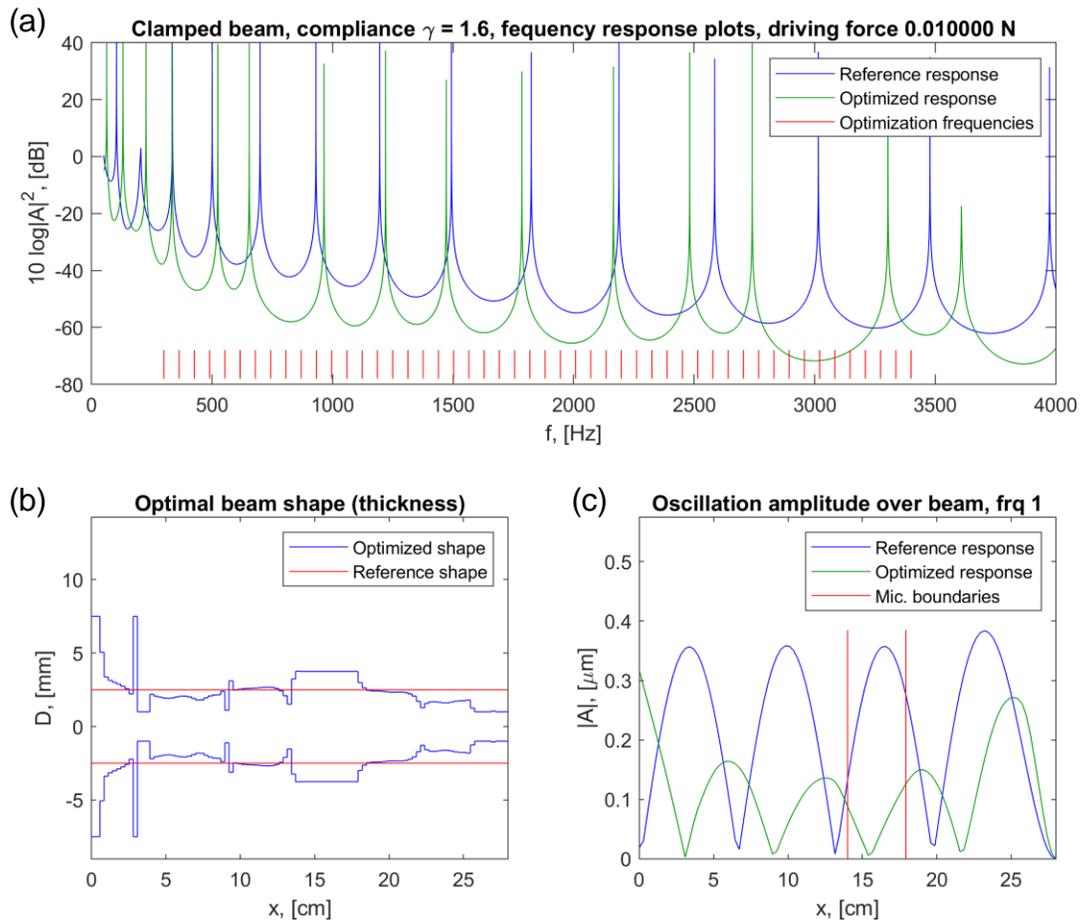


Figure 3: Optimization results for the broadband case, $\omega \in \mathcal{W}_{\text{BB}}$, using `fmincon`. The top plot (a) shows the spectrum of $J_\omega(\omega, h)$ in decibels for the optimal beam design (green) and a reference beam with a uniform thickness of 5 mm (blue). The evaluation frequencies 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. Note that Berggren et al. use the variables D and A in place of h and U , respectively.

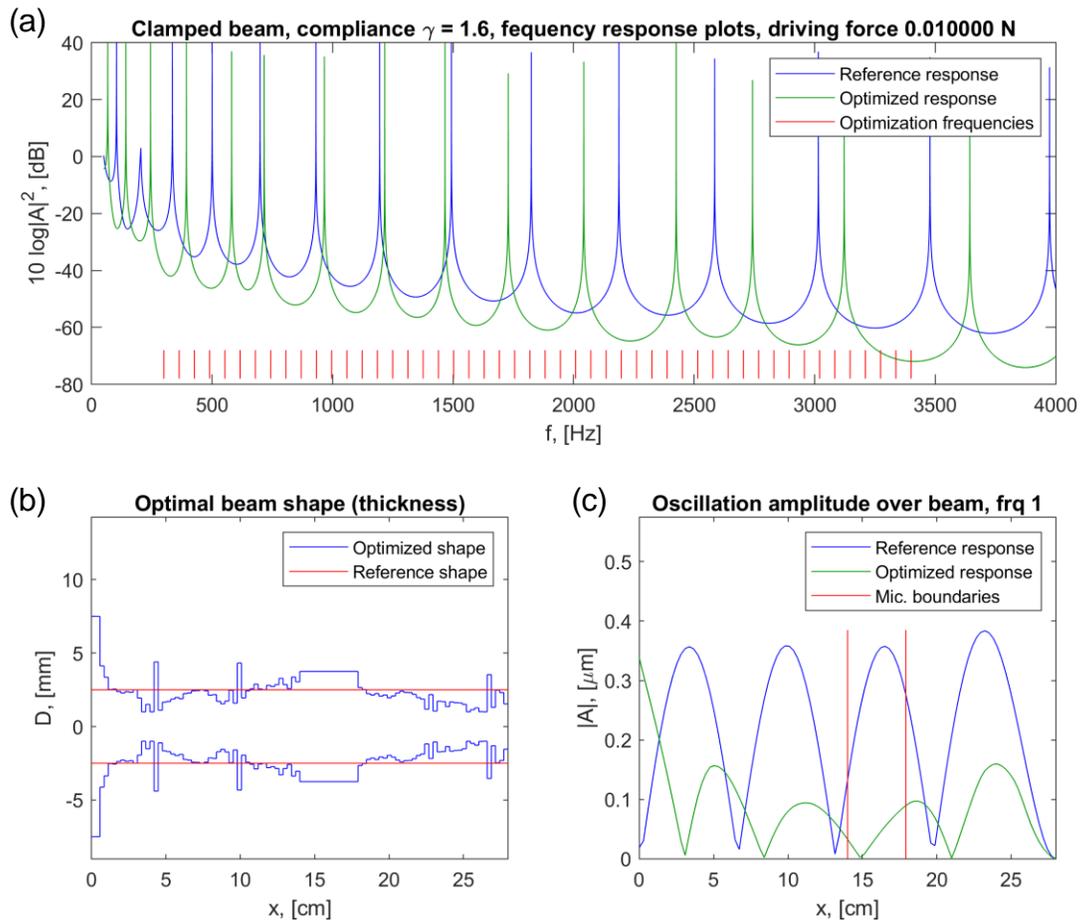


Figure 4: Optimization results for the broadband case, $\omega \in \mathcal{W}_{\text{BB}}$, using Borg. The top plot (a) shows the spectrum of $J_\omega(\omega, h)$ in decibels for the optimal beam design (green) and a reference beam with a uniform thickness of 5 mm (blue). The evaluation frequencies 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. Note that Berggren et al. use the variables D and A in place of h and U , respectively.

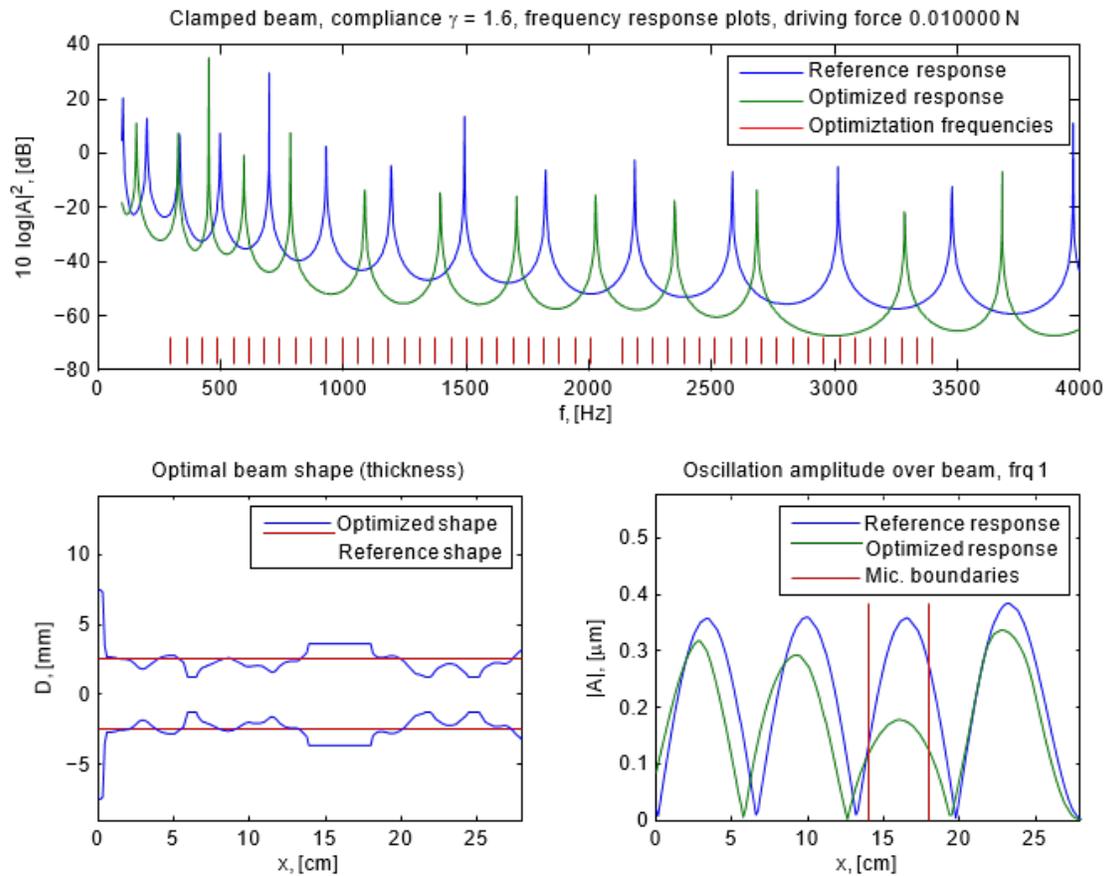


Figure 5: Figure 3 from Berggren et al.,⁴ which is the condition equivalent to Figures 3 and 4 in the present work. Note that Berggren et al. use the variables D and A in place of h and U , respectively.

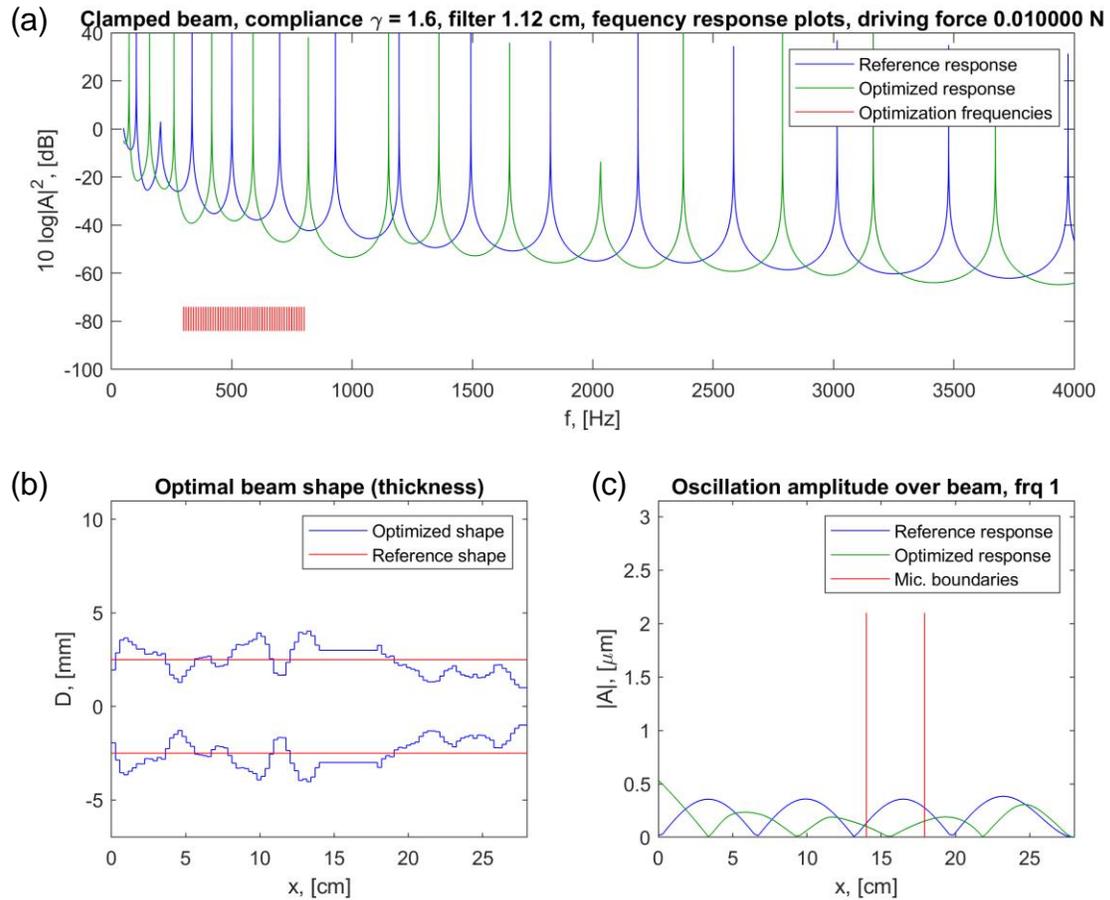


Figure 6: Optimization results for the low-frequency case, $\omega \in \mathcal{W}_{\text{LF}}$, using `fmincon`. The top plot (a) shows the spectrum of $J_\omega(\omega, h)$ in decibels for the optimal beam design (green) and a reference beam with a uniform thickness of 5 mm (blue). The evaluation frequencies 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. Note that Berggren et al. use the variables D and A in place of h and U , respectively.

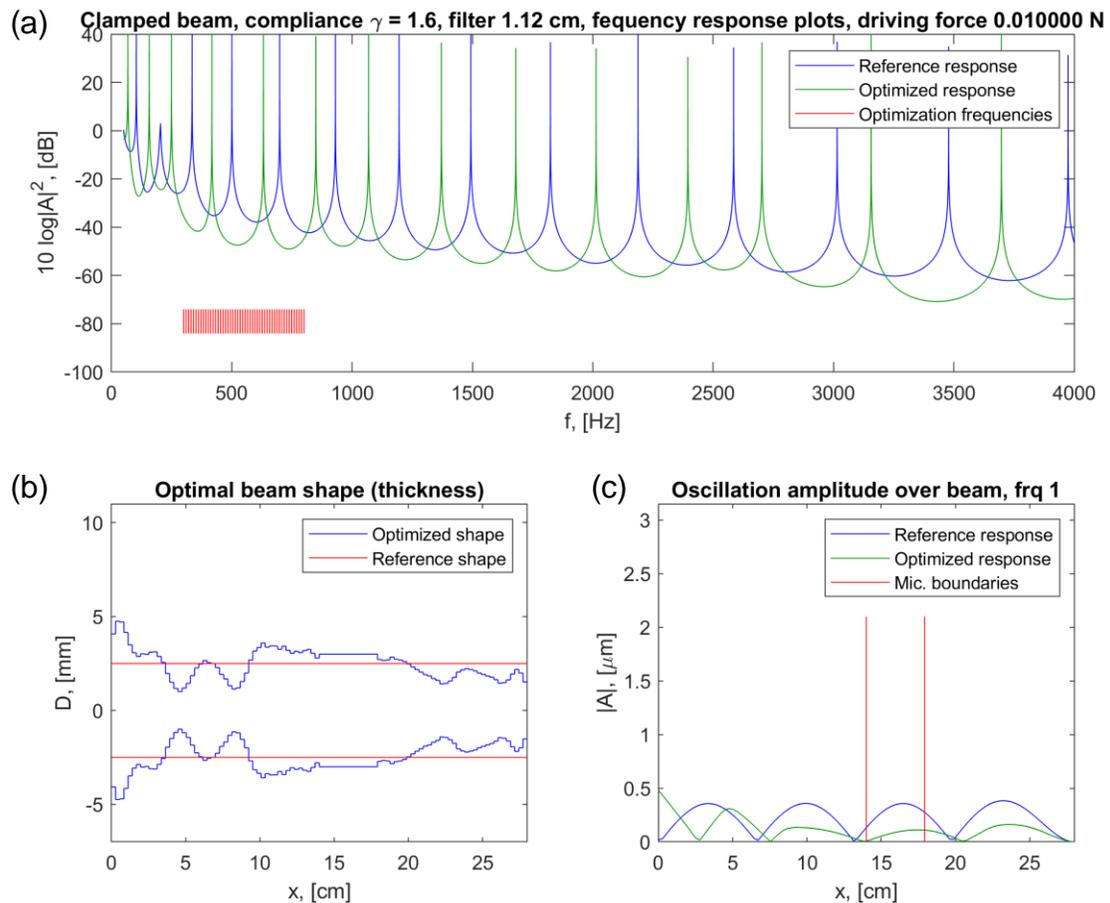


Figure 7: Optimization results for the low-frequency case, $\omega \in \mathcal{W}_{LF}$, using Borg. The top plot (a) shows the spectrum of $J_\omega(\omega, h)$ in decibels for the optimal beam design (green) and a reference beam with a uniform thickness of 5 mm (blue). The evaluation frequencies are 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. Note that Berggren et al. use the variables D and A in place of h and U , respectively.

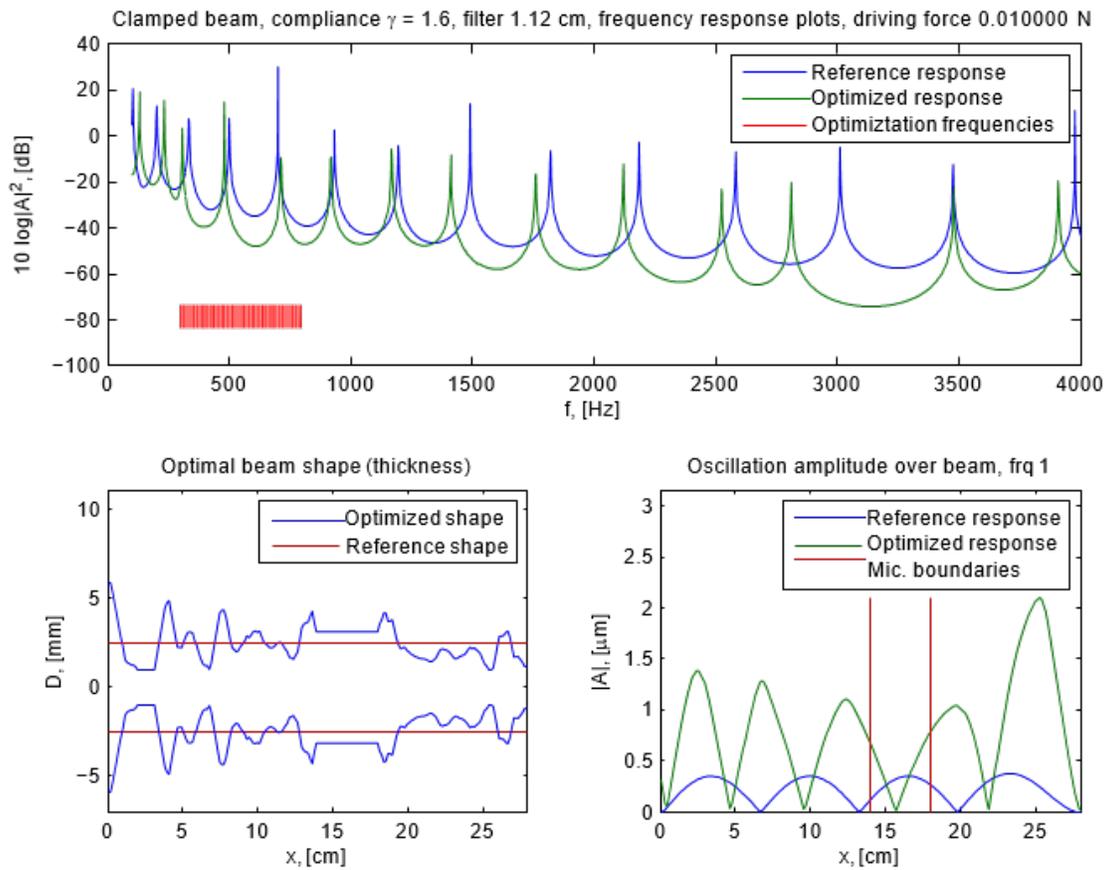


Figure 8: Figure 5 from Berggren *et al.*,⁴ which is the condition equivalent to Figures 6 and 7 in the present work. Note that Berggren *et al.* use the variables D and A in place of h and U , respectively.

Table 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

C. $\omega \in \mathcal{W}_{\text{HF}}$

Finally, the optimization results for the high-frequency case, $\omega \in \mathcal{W}_{\text{HF}}$, are shown in Figures 9 and 10 for fmincon and Borg, respectively, with the corresponding plot from⁴ reproduced in Figure 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 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 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. As a result, one would expect two new resonance peaks flanking the one original resonance frequency. Indeed, inspection of the spectra in Figure 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 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 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.

D. COMPARING FMINCON AND BORG

A summary of the optimization results for all three frequency sets is given in Table 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

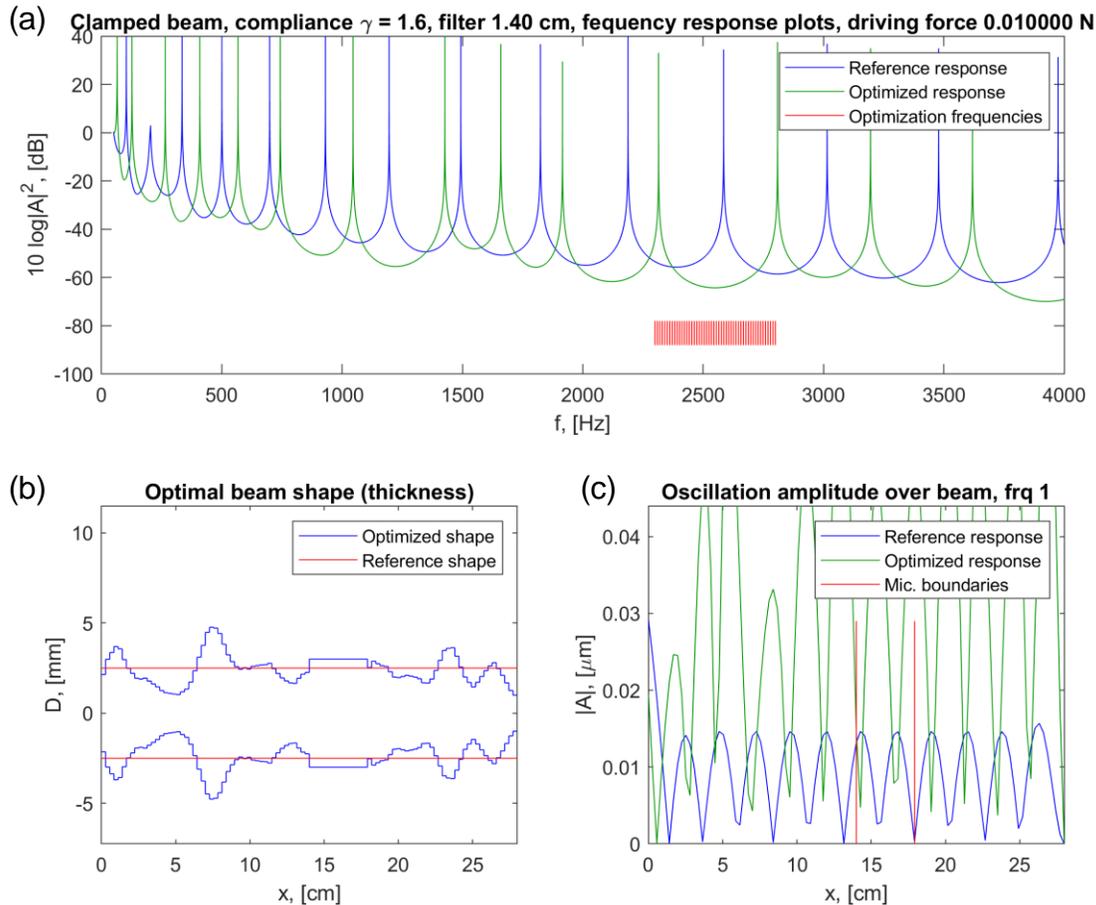


Figure 9: Optimization results for the high-frequency case, $\omega \in \mathcal{W}_{\text{HF}}$, using `fmincon`. The top plot (a) shows the spectrum of $J_\omega(\omega, h)$ in decibels for the optimal beam design (green) and a reference beam with a uniform thickness of 5 mm (blue). The evaluation 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. Note that Berggren et al. use the variables D and A in place of h and U , respectively.

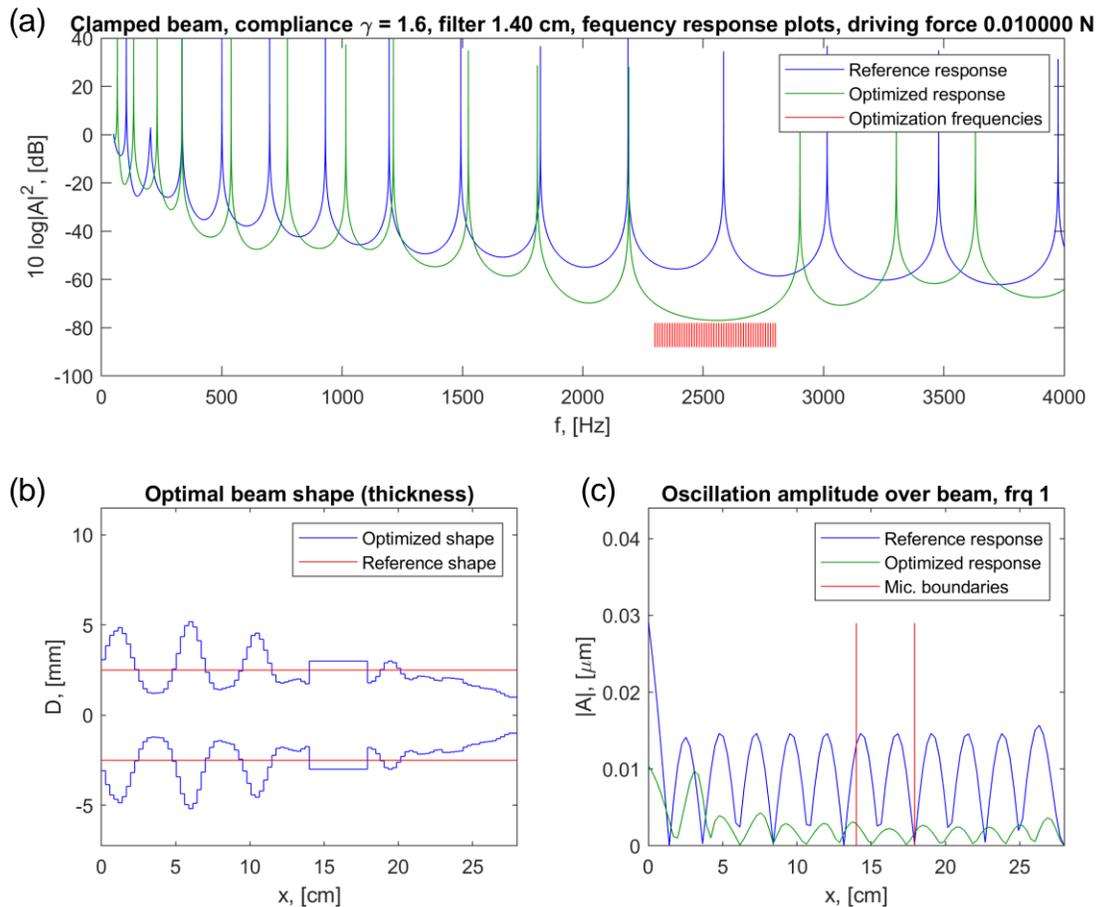


Figure 10: Optimization results for the high-frequency case, $\omega \in \mathcal{W}_{\text{HF}}$, using Borg. The top plot (a) shows the spectrum of $J_\omega(\omega, h)$ in decibels for the optimal beam design (green) and a reference beam with a uniform thickness of 5 mm (blue). The evaluation frequencies 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. Note that Berggren et al. use the variables D and A in place of h and U , respectively.

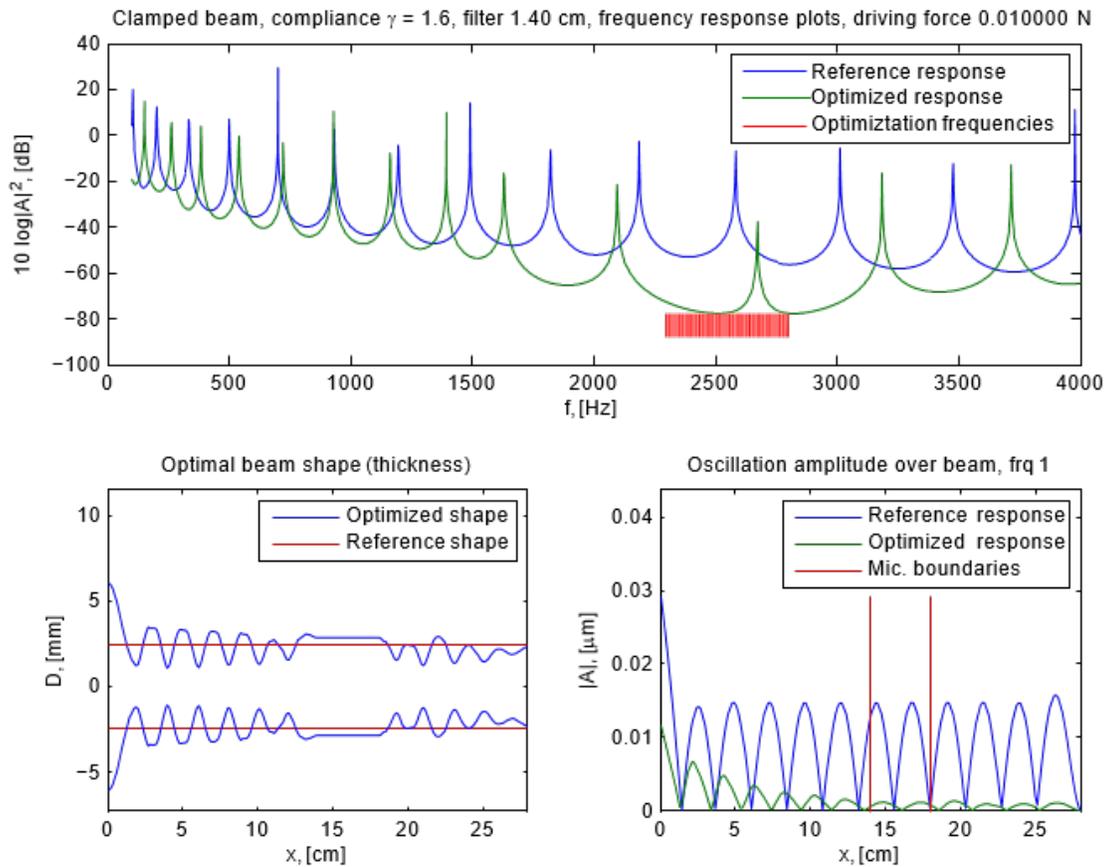


Figure 11: Figure 6 from Berggren et al.,⁴ which is the condition equivalent to Figures 9 and 10 in the present work. Note that Berggren et al. use the variables D and A in place of h and U , respectively.

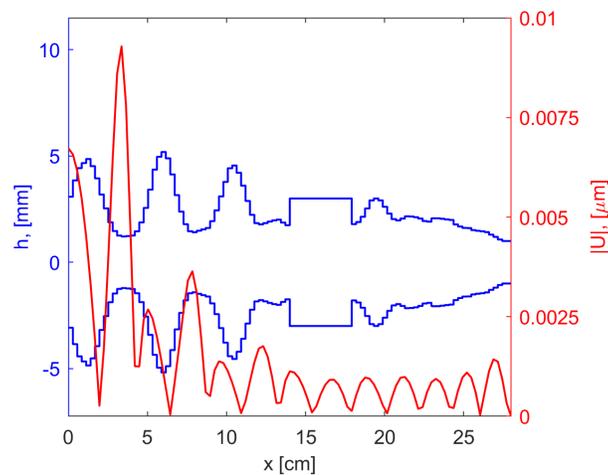


Figure 12: Optimal thickness distribution (blue) and displacement amplitude at $\omega = 2\pi \times 2557.32$ Hz (red) for the high-frequency case, $\omega \in \mathcal{W}_{\text{HF}}$, using Borg. Clearly, at 2557.32 Hz the vibration energy is concentrated at the left end away from the microphone region, (x_a, x_b) .

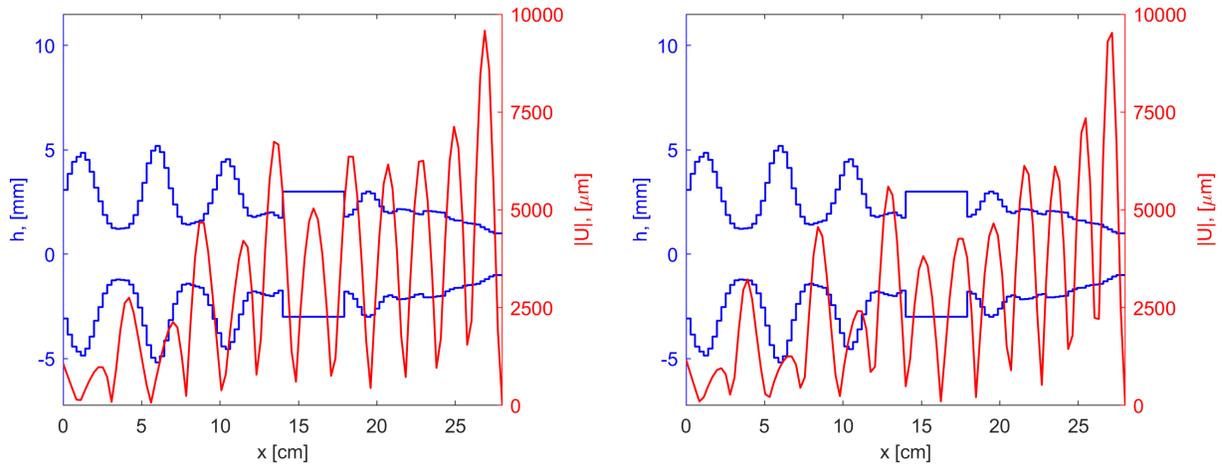


Figure 13: Optimal thickness distribution (blue) and displacement amplitude (red) for the high-frequency case, $\omega \in \mathcal{W}_{\text{HF}}$, using Borg. Analysis frequencies are $\omega = 2\pi \times 2191.23$ Hz (left) and $\omega = 2\pi \times 2902.02$ Hz (right). Whereas at 2557.32 Hz the vibration energy is concentrated at the drive point, here there is very little movement at the left end.

Table 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

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 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 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 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 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 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.

5. CONCLUSIONS

As mentioned in the introduction, the two goals of this work were to describe a general but 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.⁴ 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 method⁶ and steepest descent.⁷ 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 described here, mean that Borg can be trusted in future structural-acoustic optimization studies.

REFERENCES

- ¹ Cameron A. McCormick and Micah R. Shepherd. Design optimization and performance comparison of three styles of one-dimensional acoustic black hole vibration absorbers. *J. Sound Vib.*, 470, 2020.

- ² Micah R. Shepherd, Robert L. Campbell, and Stephen A. Hambric. A parallel computing framework for performing structural-acoustic optimization with stochastic forcing. *Struct. MultiDisp. Optimization*, 61(2):675–685, 2020.
- ³ David Hadka and Patrick Reed. Borg: An auto-adaptive many-objective evolutionary computing framework. *Evol. Comput.*, 21(2):231–259, 213.
- ⁴ Martin Berggren, Ugis Lacis, Fredric Lindström, and Eddie 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*, page 5569, 2013.
- ⁵ Cameron A. McCormick and Micah R. Shepherd. Optimization of an acoustic black hole vibration absorber at the end of a cantilever beam. *J. Acoust. Soc. Am.*, 145(4):EL591–EL597, 2019.
- ⁶ Jörg Fliege, Luis M. Graña Drummond, and Benar F. Svaiter. Newton’s method for multiobjective optimization. *SIAM J. Optim.*, 20(2):602–626, 2009.
- ⁷ Jörg Fliege and Benar F. Svaiter. Steepest descent methods for multicriteria optimization. *Math Meth. Oper. Res.*, 51:479–494, 2000.