Continuation strategies to mitigate convergence to low-performing local optima in dynamic topology optimization

Tom De Weer[a,c], Vanessa Cool[b,c], Elke Deckers[a,c]

[a]*KU Leuven Campus Diepenbeek, Department of Mechanical Engineering, Wetenschapspark 27, Diepenbeek, 3590, Belgium*
[b] *KU Leuven, Department of Mechanical Engineering, Celestijnenlaan 300, Heverlee, 3001, Belgium*
[c]*Flanders Make@KU Leuven, Belgium*

Abstract

Solving dynamic topology optimization problems often yields low-performing local optima. Instead of converging towards a design that exploits dynamic mechanisms, a less interesting, mass-driven solution is often generated. This necessitates repeated and computationally expensive optimization reruns before a suitable optimum is found. In this work, an overview of three strategy classes that reduce the need for such reruns is presented: exclusion strategies, frequency shift methods and relaxation strategies. Novel variants for each strategy class are developed, implemented and compared via Monte Carlo sampling on a benchmark problem, namely the sound transmission loss optimization of a sandwich panel. Probabilities of achieving high-performing optima are estimated and all investigated strategies demonstrate quantifiable improvements and trade-offs. The study offers furthermore a quantitative comparison of the presented strategies, supporting researchers in making an informed choice when addressing convergence to poor local optima in dynamic topology optimization.

Keywords: dynamic topology optimization, local optima, continuation strategies

1. Introduction

Mechanical structures in numerous engineering applications are continuously subjected to dynamic loading. This complicates their design, since the dynamic load must be incorporated in the design process. However, adapting the design to shift its eigenfrequencies away from the loading frequencies can lead to an unwanted loss of stiffness, which poses a difficult trade-off. To strike a balance in this design challenge, the seminal introduction of topology optimization (TO) [1] in 1988 was quickly followed by its application to dynamic problems by Díaaz and Kikuchi [2] in 1992. Since then, the subfield of dynamic TO has seen research towards e.g, numerical sensitivity issues and mode switching [3], large scale computations and model order reduction [4, 5, 6] and applications such as noise-reducing sandwich structures [7], resonant structures [8] and Micro Electro Mechanical Systems (MEMS) [9].

A particular issue in dynamic TO is its many local optima. Although classical compliance minimization is already known for its strong nonconvexity, dynamic TO exacerbates this issue due the eigenfrequencies of the structure. Resonance and antiresonance frequencies create a design space with sharp peaks and valleys between which dynamic compliance optimizers get stuck [10]. If eigenvalue optimization formulations are used, then a choice remains as to whether they are tuned up or down [11]. A particular issue in dynamic TO is the existence of a class of numerically attractive but low-performing optima. Especially at low frequencies, TO routines tend to simply shift all frequencies upwards to produce stiff, mass-driven designs that do not exploit any compliant mechanisms or resonance behavior. To deal with the issue of local optima in (dynamic)

Email address: `tom.deweer at kuleuven.be` (Tom De Weer)

topology optimization problems, many techniques already exist in the literature. Yet, a clear overview and quantifiable comparison is lacking. To deal with the former gap, this work proposes to divide all strategies into four categories, which are now discussed.

The first, most general category consists of *global optimization methods*. Such methods are often applicable to generic optimization problems, of which design optimization is but a subcategory. A standard way is to start optimization from a large amount of initial guesses generated randomly by, e.g., Latin hypercube sampling, and select the best local optimum. One example that improves upon this sampling concept is recent work by Herrman et al. [12], who use a neural network material discretization to improve that chance at a better optimum for acoustic TO problems. Another instance of this category is by Assimi and Jamali [13], who show that a combination of Nelder-Mead and genetic programming can lead to lighter optimal truss layouts when considering a large amount of static and dynamic loads. Giannini et al. [14] use genetic algorithms for the sizing optimization of metamaterial plates. Zhang et al. [15] apply the tunneling method to compliance minimization, moving from one optimum to a better one by ignoring all designs that perform worse. Another example is from Liu et al. [16], who optimize the band gaps of bi-material beams and plates by employing a gradient-free Kriging-based methodology. Finally, recent work by Dalklint et al. [17] shows a way to find out how far a topology optimized solution is from the global optimum. This provides a (numerically expensive) way to gauge performance bounds and stop global optimization techniques when they are (close to) globally optimal. In general, however, global optimization requires excessive computational resources, especially so in TO contexts.

The second category are called *exclusion methods*. If the class of numerically attractive but undesired local optima is known very well, then sometimes constraints can be included to remove those optima from the feasible design space. Arguably the simplest example of this in (dynamic) TO is the use of a volume constraint: without it, routines converge to a fully solid, mass-driven design. Nonetheless, exploiting dynamic behavior such as, e.g., resonances, by reducing mass and introducing compliance can lead to superior performance. In this context, a volume constraint is a numerical way to induce such behavior. A second example of an exclusion methods is the use of connectivity constraints. As highlighted in the review by Cool et al. [18], TO routines for vibro-acoustic applications often produce disconnected structures, i.e., structures with floating solid (or void) regions. Examples include broadband filter design by Dilgen and Aage [19], band gap maximization in Mindlin plates by Halkjær et al. [20] and sound transmission loss maximization of sandwich structures by Cool et al. [7]. Such disconnected designs can provide superior performance, making them numerically attractive local optima. However, they are often undesired, for either stiffness or manufacturability reasons. As such, applying a connectivity constraint makes disconnected designs infeasible and forces routines to produce more useful local optima. More generally, exclusion methods can completely prevent the appearance of undesired classes of local optima, provided that the class is understood well enough to come up with a suitable exclusion constraint.

The third category consists of *relaxation methods*. At its core, relaxation methods aim to solve a simpler but related optimization problem first, and use it as a starting point for the actual problem one wants to solve. In this sense, it is strongly related to the various continuation schemes encountered in TO. For example, classical density-based TO with the SIMP law often starts with a low penalization before increasing it to obtain a black-white solution. This is done because at low penalization, the problem is convex and the global optimum can be used as a proper starting guess for iterations at higher penalization and hence concavity. The idea of penalization continuation has a long history and stems back to early studies towards density-based compliance minimization. Arguably the earliest work is that of Allaire and Kohn in 1993 [21], which employs a grayness penalization term in the objective that increases once throughout the optimization. The design space is thus "relaxed" by allowing density values between 0 and 1, and then slowly transformed to the original black-white problem by increasing the penalization. This continuation is a general feature of relaxation methods. More recently, Clay et al. [22] used an objective consisting of an additional penalty term that encapsulates the error in the constitutive relations. The relaxation thus consists of the fact that the constitutive relations need not be satisfied at the optimization start. The optimization of the direct frequency response function involves choosing displacements, stresses and inertial forces on top of the usual density field. It thereby avoids optimization of a functional solely of the displacements, which yields superior designs.

As another example, Guo et al. [23] employ a second order smooth-extended technique in combination with the epsilon-relaxed method to allow for better truss layouts when considering local stress and buckling constraints. In general, relaxation is a powerful tool in optimization. Many techniques and variations exist and each requires careful tuning to ensure the initial relaxation is gradually restricted in a numerically stable way throughout the optimization process.
The fourth category are the *frequency shift* methods. In a sense, they are a subcategory of relaxation methods applicable specifically to dynamic TO problems. In the context of dynamic compliance minimization, it is known that TO generates designs by shifting the eigenfrequencies away from the loading frequency. If the loading frequency lies below an eigenfrequency, that eigenfrequency is shifted upwards by increasing the stiffness for the corresponding mode, and vice versa. This means that, for low loading frequencies, the result is a stiff, mass-driven design that shifts all eigenfrequencies upwards. This realization lead Olhoff and Du [24] to propose the *incremental frequency technique*, applying it to minimization of dynamic compliance for single- and bi-material structures. The technique uses a continuation scheme to vary the load frequency throughout the optimization process and obtain better local optima. The generalized incremental frequency method then repeats this process for a range of initial load frequencies, essentially obtaining a global optimization methodology from the first category. As shown by Silva et al. [10], the underlying problem tackled in [24] is that the dynamic compliance exhibits an alternation of resonances and anti-resonances. The resonances provide boundaries between which the optimizer gets stuck and the antiresonances provide local optima to which the optimizer converges. Guzman et al. [8] and Meng et al. [25] exploit this property by employing an antiresonance matching objective and an antiresonance constraint, respectively. For the purpose of this work, incremental frequency techniques can be seen as a way to generate solutions for frequency ranges where only mass-driven designs are found: by starting at a high loading frequency and then decreasing it to the desired one, more compliant structures that exploit dynamic behavior can be found. For example, Giannini et al. [9] obtain MEMS designs for a target frequency of 16 kHz and 24 kHz by starting from an adapted solution found for a 26 kHz target. However, a rigorous analysis of this technique is lacking in the literature. The above examples of frequency shift methods could be considered *direct* frequency shifts: the loading frequency is changed directly throughout the optimization. A more subtle way of frequency shift methods are what is here denoted as *indirect* frequency shift methods: through a continuation of, e.g., the geometry or mass of the setup, the eigenfrequencies of the initial design can be tuned below or above the loading frequency, without changing the latter. For example, if the design domain is an $L \times L$ square domain, one can artificially shift the eigenfrequencies up or down by decreasing or increasing L, respectively. Similar effects can be generated by artificially tuning the Young's modulus or mass density. As far as the authors are aware, there are as of yet no indirect frequency shift methods in the TO literature. As they have the same effect as direct frequency shift methods, they are not further studied in this work.

The above literature overview shows the clear existence of various classes of low-performing, undesired optima, e.g., disconnected or mass-driven designs, and the literature's strategies to mitigate convergence towards such optima. Yet, a quantifiable comparison between the distinguished strategy categories is still lacking. To provide such a comparison, this study considers the optimization of the sound transmission loss (STL) of a sandwich panel, shown in Figure 1, as a case study.
The STL of a uniform plate follows the mass law: a doubling of the frequency (an octave) or a doubling of the mass leads to an additional 6 dB. A sandwich panel, such as the 1D mass-spring-mass example of Figure 1, follows the mass law at low frequencies. However, at the decoupling frequency (3000 Hz) the top and bottom panel move in counter phase: pressure radiation is maximal and the STL reaches a minimum. Above the decoupling frequency, the STL increases by 18 dB per octave. As the next section will show, for the chosen application there exists a strongly attractive class of sandwich panel optima that behave like the mass law. This class consists of stiff, mass-driven structures that increase the decoupling frequency upwards as far as the given mass allows, by rigidly connecting the top and bottom plate. This defeats the purpose of using a sandwich panel and misses out on improved performance by, e.g., exploiting the decoupling and local resonance effects found by Cool et al. [7].
The outline of this manuscript is as follows. Section 2 lays out a baseline framework for sandwich panel STL optimization, based on the work of Cool et al. [7]. Then, Section 3 uses Monte Carlo sampling to illustrate

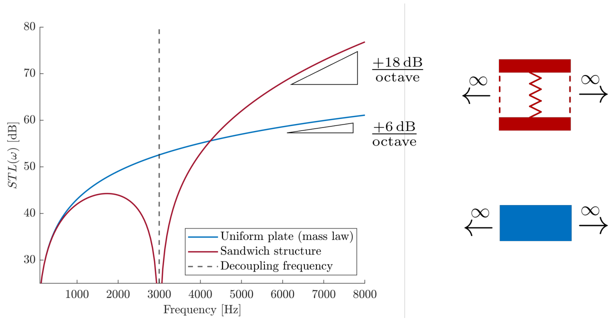

Figure 1: Sound transmission loss of a uniform plate (blue) and a sandwich panel connected with zero-mass spring (red). Both designs have the same mass.

and quantify the problem tackled in this work: the high probability that optimizers converge to low-quality optima. Next, Section 4 compares the three strategy classes: exclusion, frequency shift and relaxation. This is achieved by developing and implementing novel variants for each strategy class and testing them via the same Monte Carlo methodology. The first strategy (Section 4.1) is an exclusion method, which adopts an upper bound on the stiffness to make low-quality optima infeasible. The second strategy (Section 4.2) is a frequency shift method, i.e., a variant of the incremental frequency technique from Olhoff and Du [24], that pulls high-frequency solutions to lower frequency ranges. The third strategy (Section 4.3) is a relaxation method that softens the required design robustness at the start of the optimization process. Section 4.4 then bundles the conclusions for each strategy and, finally, Section 5 formulates the conclusions.

2. Optimization framework description

This work considers a 2D sandwich panel consisting of two plates with a vibro-acoustic core in between. The core consists of a repeated unit cell (see Fig. 2). Above and below the panel, infinite acoustic half spaces containing air (density $\rho_a = 1.225$ kg/m^3 and speed of sound $c_a = 340$ m/s) are present.

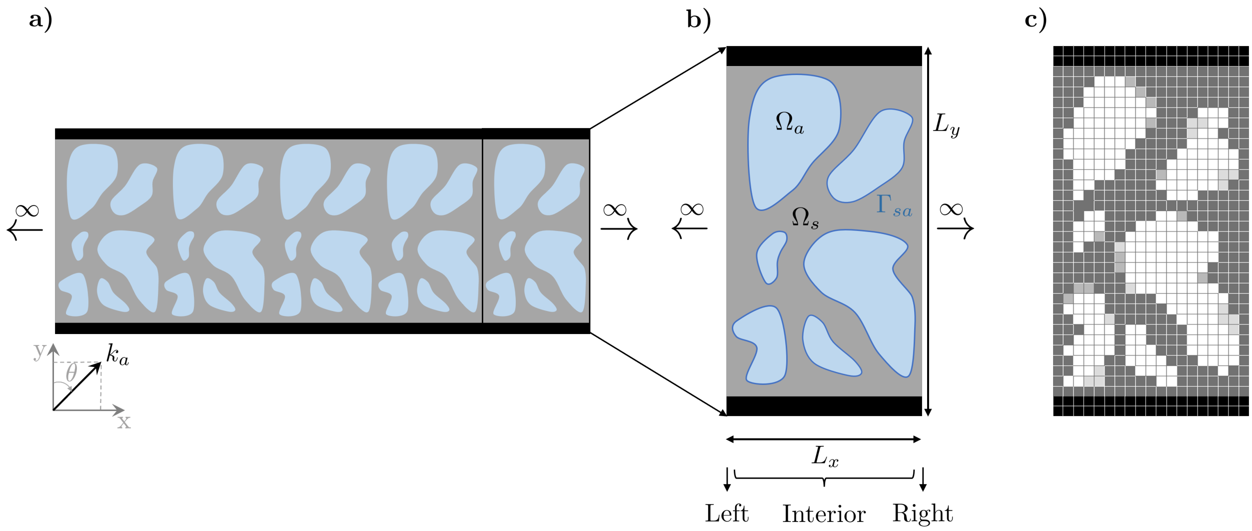

Figure 2: Problem setup. a) Infinite sandwich panel. b) Unit cell considered during optimization. c) Example of FE discretization. Taken and adapted from [7] with permission from the corresponding author.

An incident plane acoustic pressure wave (amplitude P_i, acoustic wave number $k_a = \omega/c_a$ and radial frequency ω (rad/s)) excites the panel at the bottom. This causes a transmitted wave at the top (amplitude

P_t) and a reflected wave at the bottom (amplitude P_r). In the Fourier domain, these three plane acoustic pressure waves are written as:

$$p_i(x,y) = P_i \mathrm{e}^{\mathrm{i}(-k_x x - k_y y)}, \quad p_t(x,y) = P_t \mathrm{e}^{\mathrm{i}(-k_x x - k_y y)}, \quad p_r(x,y) = P_r \mathrm{e}^{\mathrm{i}(-k_x x + k_y y)}, \tag{1}$$

where $k_x = k_a sin(\theta)$, $k_y = k_a cos(\theta)$ are trace wavenumbers and θ is the incidence angle. For ease of notation, the temporal factor $\mathrm{e}^{\mathrm{i}\omega t}$ is excluded. The exclusion of higher-order acoustic harmonics is made by assuming the sandwich panel is weakly periodic and the unit cell size is smaller than the wavelength [26]. The STL is used to asses the acoustic performance of the structure. It is defined as the ratio of the transmitted pressure amplitude to the incident pressure amplitude [27]:

$$\mathrm{STL}(\omega,\theta) = -10\log_{10}\left(\tau(\omega,\theta)\right), \qquad \tau(\omega,\theta) = \left|\frac{P_t}{P_i}\right|^2. \tag{2}$$

Here, τ denotes the sound power transmission coefficient.

The layout of the sandwich panel core is optimized with a TO framework. The baseline implementation is based on [7] and the following two sections elaborate on the simulation and optimization aspects of this considered approach, respectively.

2.1. Performance prediction

The sandwich panel is assumed to consist of an infinite repetition of the unit cell in the x-direction (see Fig. 2b). This unit cell has dimensions $L_x \times L_y$. The layout of the structural domain (Ω_s) and acoustic domain (Ω_a) is chosen by the optimization framework. The unit cell is discretized with the finite element (FE) method, employing a rectangular $N_x \times N_y$ grid for a total of N_e elements: bilinear square plane strain elements for the structural domain and linear elements for the acoustic domain. The assumption of time-harmonic motion leads to the following system of equations [27]:

$$\left(-\omega^2 \underbrace{\begin{bmatrix} \mathbf{M}_s & 0 \\ \mathbf{S}_u & \mathbf{M}_a \end{bmatrix}}_{\mathbf{M}} + \underbrace{\begin{bmatrix} \mathbf{K}_s & \mathbf{S}_p \\ 0 & \mathbf{K}_a \end{bmatrix}}_{\mathbf{K}}\right) \underbrace{\begin{bmatrix} \mathbf{u} \\ \mathbf{p} \end{bmatrix}}_{\mathbf{q}} = \underbrace{\begin{bmatrix} \mathbf{f}_s \\ \mathbf{f}_a \end{bmatrix}}_{\mathbf{f}} + \underbrace{\begin{bmatrix} \mathbf{e}_s \\ \mathbf{e}_a \end{bmatrix}}_{\mathbf{e}}, \tag{3}$$

where $\mathbf{M}, \mathbf{K}, \mathbf{S}$ are the mass, stiffness and coupling matrices and the subscripts s, a denote the structural and acoustic domain, respectively. Each node has three degrees-of-freedom (DOFs), namely the x- and y-displacement and the pressure. Conglomerated over all the nodes, these are denoted as the displacement vector $\mathbf{u}$ and the pressure vector $\mathbf{p}$. Together, they represent DOFs vector $\mathbf{q}$, which corresponds entry-wise to the nodal force vectors $\mathbf{f}$ and $\mathbf{e}$, i.e. the internal forces and external forces, respectively. Infinite periodicity is enforced with Bloch-Floquet boundary conditions [28]. For more details, the reader is referred first to Appendix A and the works of Boukadia et al. and Cool et al. [26, 7].

The unit cell has size 50×50 mm with a plate thickness of 5 mm. The element dimensions are 0.5×0.5 mm, for a total of 100×100 elements, to ensure sufficient accuracy for the acoustic and structural waves at the considered frequencies. The same materials as in [29] are considered. For the structural part, this is polymethyl methacrylate (PMMA) with Young's modulus $E = 4.85(1 + 0.05\mathrm{i})$ GPa where the imaginary part represents structural damping, density $\rho_s = 1188.35$ kg/m^3 and Poisson's ratio $\nu = 0.31$, all taken from [30]. The core and surrounding fluid have the same properties, except for a small ratio of acoustic damping in the core through a complex speed of sound: $\rho_a = 1.225$ kg/m^3 and $c_a = 340(1 + 2 \cdot 10^{-4}\mathrm{i})$ m/s.

2.2. Optimization framework

This section presents the baseline TO framework used to design the sandwich panel core.

2.2.1. Optimization variables

The unit cell core is optimized with a density-based TO framework. Each FE of the unit cell is controlled with an optimization variable $\boldsymbol{\xi}$ for a total of N_e optimization variables. Projection and filtering techniques map the optimization variable field $\boldsymbol{\xi}$ to the physical density field $\boldsymbol{\xi}_P$. An element i contains fluid material if $\xi_{i,P} = 0$, solid material if $\xi_{i,P} = 1$ and an interpolation of fluid and solid if $0 < \xi_{i,P} < 1$. The physical density field $\boldsymbol{\xi}_P$ thus determines the material properties of the elements, the resulting stiffness and mass matrices to execute the FE analysis and ultimately the values of the objective and constraints. The solid-fluid boundaries during optimization are dealt with via the vibro-acoustic coupling method of Jensen [31]. Appendix A provides more information on the projection and filtering techniques.

Instead of a single physical design field, the robust formulation of [32] with three physical designs is employed: an eroded (e), a blueprint (b) and a dilated (d) design. The minmax formulation, discussed further below, will ensure the optimizer prioritizes the worst-performing of the three designs. To further improve numerical stability, the double filtering approach of [33], with filter radii R_1 and R_2, is used. Two corresponding smoothed Heaviside projections are used with steepness variables β_1 and β_2.

2.2.2. Objective function

For a given design (eroded, blueprint or dilated), the performance indicator is the STL in a frequency range of interest $\Delta\omega = [\omega_-, \omega_+]$, defined and numerically approximated as follows:

$$STL(\Delta\omega, \theta) = \frac{\int_{\omega_-}^{\omega_+} STL(\omega, \theta)\,\mathrm{d}\omega}{\int_{\omega_-}^{\omega_+} \mathrm{d}\omega} \approx \frac{1}{\omega_+ - \omega_-} \sum_{\omega_i \in \Delta\omega} a_{\omega_i} STL(\omega_i, \theta), \tag{4}$$

where a_{ω_i} are midpoint integration weights. In this paper, $\Delta\omega$ is sampled with 5 ω_i. The incidence angle θ is kept at $0°$ and, for brevity, omitted in the rest of this paper. For numerical stability, the performance indicator is normalized as follows:

$$J_{b/e/d} = 1 - \frac{STL_{b/e/d}(\Delta\omega)}{C}, \tag{5}$$

where $b/e/d$ denotes either the blueprint, eroded or dilated design and C is a normalization constant. Here, $C = 120$ and the frequency range width $\omega_+ - \omega_-$ is $500\,\mathrm{Hz} = 2\pi 500\,\frac{\mathrm{rad}}{\mathrm{s}}$. In the remainder of this work, all numerical values of ω are specified in Hz without the implied 2π scaling factor.

2.2.3. Constraints

Three constraints are used: 1) a bound constraint on the optimization variables $\boldsymbol{\xi}$, 2) a volume constraint and 3) a connectivity constraint. The optimization variables are bounded to lie in $[0, 1]$, so $0 \leq \boldsymbol{\xi} \leq 1$. The volume $v_{d,\mathrm{P}}$ is computed as in [7] and the maximum volume is kept constant at $V = 0.5V_{tot}$, with V_{tot} the total volume. The connectivity constraint differs from [7]: the static top load is replaced by the self-weight load to reduce the appearance of disconnected structures in the core [18]. The compliance under this self-weight θ_{sw} is constrained to be $\mu_{sw} = 15$ times smaller than the self-weight compliance for a fully solid design, denoted by $\hat{\theta}_{sw}$. Appendix A contains computational details regarding self-weight.

2.2.4. Formulation

The baseline implementation employs the following minmax formulation:

$$\begin{cases} & \min\limits_{\boldsymbol{\xi} \in \mathbb{R}^{N_e}} \quad \max(J_b, J_e, J_d) \\ \text{s.t.} & J_{\mathrm{vol}} = v_{d,\mathrm{P}}/V - 1 \leq 0 \\ & J_{\mathrm{conn}} = \theta_{\mathrm{sw}}/\mu_{\mathrm{sw}}\hat{\theta}_{\mathrm{sw}} - 1 \leq 0 \\ & 0 \leq \boldsymbol{\xi} \leq 1. \end{cases} \tag{6}$$

Here, the volume and connectivity constraints are normalized to obtain J_{vol} and J_{conn}, respectively.

2.2.5. Convergence, continuation and adaptive move limit parameters

The optimization is carried out with the Method of Moving Asymptotes [34] (MMA) with the asymptote control parameters chosen as $s_{init} = 0.2$, $s_{decr} = 0.65$, $s_{incr} = 1.06$ and the outer move limit at $\mu = 0.1$.
All variations of the optimization routine in this work make extensive use of continuation, i.e., the change of hyperparameters ($\beta,\ \Delta\omega,\ \ldots$) throughout the optimization process to gradually drive the solution towards an optimum. The baseline implementation uses classical β continuation: β_1 is increased from $\beta_1 = 1$ to, at most, $\beta_1 \approx 64$, through 19 increases by a factor of 1.2. After every increase of β_1, β_2 is set to $\beta_1/2$. The projection parameters η_e, η_b and η_d are kept constant at 0.4, 0.5 and 0.6. The smoothing parameters R_1 and R_2 are also kept constant at 14 and 7 element widths, respectively.
The strategies of Section 4 use continuation of additional parameters. For all variations of the TO algorithm, the hyperparameters only change after an intermediate convergence criterion is satisfied. The routine thus runs in stages with constant hyperparameters. This study considers convergence to be 1) when the constraints are satisfied (or their absolute change is less than 0.01 between iterations) and 2) when the last 5 iterations saw a change of less than 0.01 dB in the smallest STL. When convergence is reached and there are no more hyperparameters that must be changed, the routine terminates.

On top of the outer move limit μ imposed internally by MMA, the bound constraints $0 \leq \boldsymbol{\xi} \leq 1$ are adapted to reflect an additional move limit μ_2 such that $\Delta\boldsymbol{\xi} = \max(\boldsymbol{\xi}_i - \boldsymbol{\xi}_{i-1}) < \mu_2$. To ensure numerical stability, μ_2 changes throughout the optimization process, irrespective of if convergence is reached. The details of the μ_2 updating scheme are explained in Appendix A.

3. Problem description

This section illustrates (Section 3.1) and then quantifies (Section 3.2) the studied problem.

3.1. Problem illustration

Figures 3 and 4 show two optimization runs with the baseline implementation, targeting the frequency range [2000 Hz, 2500 Hz] by starting from a random initial guess. A clear difference is visible. Whereas Figure 4 shows a successful run that achieves a high-performing design, Figure 3 shows a run that remains stuck in a low-performing local optimum. Almost no STL improvement is achieved. Instead, the final result is stiff to ensure that it behaves like a single plate, with performance comparable to the mass law.
The unsuccessful run starts from a random initial guess (Figure 3c, top left). The first 100 iterations follow a repeating pattern: a violation of the volume constraint (Figure 3b) leads to a quick recovery to a feasible design, an $STL(\Delta\omega)$ (Figure 3a) that flatlines and finally a β_1 increase that again causes a violation of the volume constraint. During this process, the core of the sandwich panel gets gradually less grey and hence less stiff, causing a decrease of the $STL(\Delta\omega)$ because the decoupling frequency shifts leftwards. The remainder of the optimization process sees a counteraction of this effect: the panel is connected with a rigid, vertical beam that increases the core's stiffness, shifts the decoupling frequency upwards and causes an $STL(\Delta\omega)$ gain. The final design is thus very stiff and does not exploit the flexibility allowed by the rather loose connectivity constraint.
The successful run starts similarly. The first 150 iterations see the same softening of the core followed by a stiffening due to the formation of a rigid, vertical connection. After iteration 150, the successful run visibly diverges: the vertical connection becomes more and more askew (Figure 4c) until the connectivity constraint is active (Figure 4b). The decoupling frequency shifts to more than 1000 Hz below the targeted frequency range (Figure 4d) and the performance ($STL(\Delta\omega)$, Figure 4a) increases by 20 dB. At about iteration 400, this increase flatlines and subsequent β_1 increases result in a crisp and high-performing design shown in the bottom right of Figure 4c. The remaining iterations (200 to 245) see an almost constant $STL(\Delta\omega)$ and an $STL(\omega)$ that follows the mass law (see Figure 3d). On the other hand, the successful run sees the lowest $STL(\Delta\omega)$, the eroded $STL_e(\Delta\omega)$, collide with the second-lowest blueprint $STL_b(\Delta\omega)$ at about iteration 175. The $STL(\Delta\omega)$ then flatlines and iteration 200 sees three fast β increases. However, contrary to the unsuccessful run, the connection between the top and bottom plate is slightly askew and $STL(\omega)$ (Figure 4d)

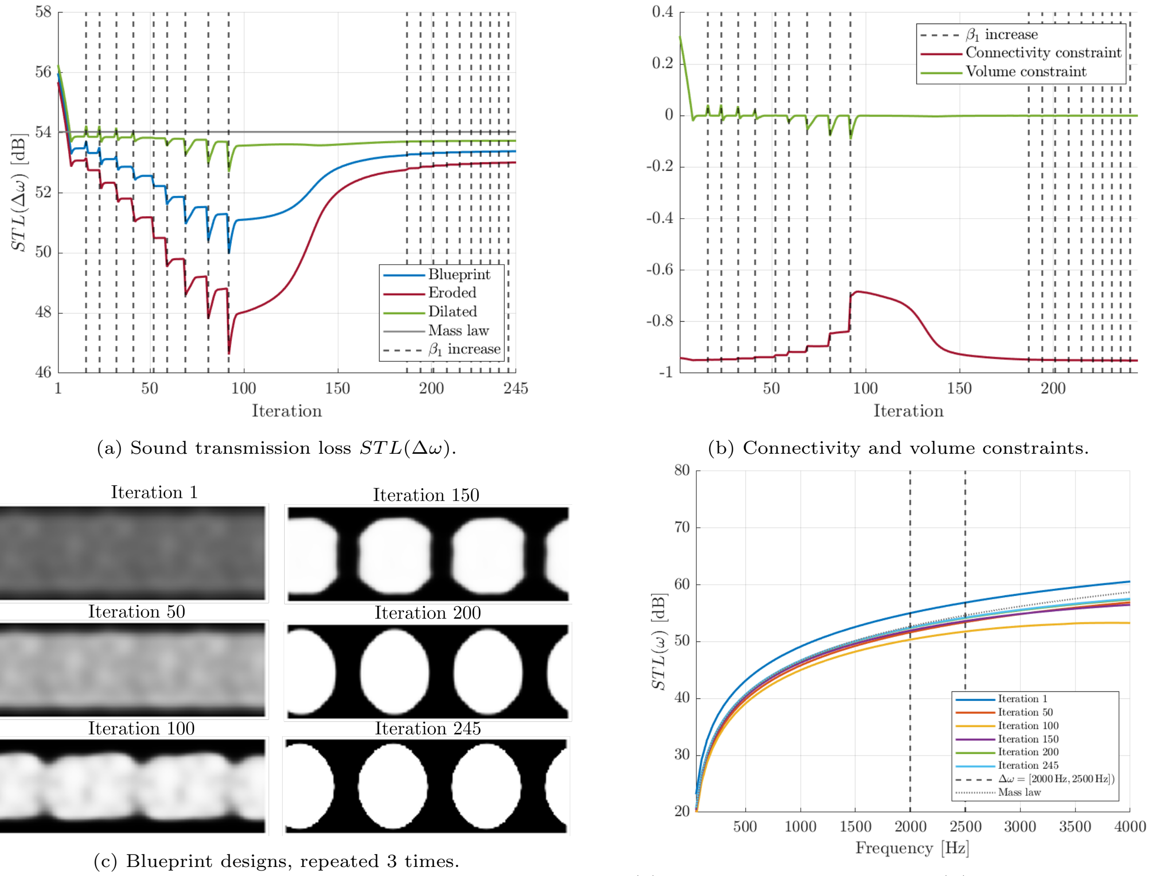

(a) Sound transmission loss $STL(\Delta\omega)$.

(b) Connectivity and volume constraints.

(c) Blueprint designs, repeated 3 times.

(d) Sound transmission loss $STL(\omega)$ of blueprint designs.

Figure 3: Baseline optimization run yielding a mass-driven optimum, for $\Delta\omega = [2000\,\text{Hz}, 2500\,\text{Hz}]$.

shows the appearance of the decoupling frequency inside the targeted frequency range. The remainder of the optimization process then sees the connection between top and bottom plate become 1) more askew and 2) notched just before it reaches the top plate. This is accompanied by a further reduction of the decoupling frequency until it lies more than 1000 Hz below the targeted frequency range.

The high-performing design exploits two mechanisms: 1) a reduction in the decoupling frequency to exploit the strong 18 dB per octave STL increase that occurs afterwards and 2) the transformation of the vertical motion of the top plate, which causes the transmitted pressure wave, to a horizontal motion that does not transmit pressure waves. As a result, the STL of the high-performing design lies 20 dB or about 35% above the mass law in the targeted frequency range. This is much better than the low-performing, mass-driven optimum of Figure 3. As such, the goal of this work is to prevent the appearance of mass-driven optima in favor of those that exceed the mass law.

3.2. Problem quantification

The runs of Figures 3 and 4 differ only in the random initial guess. Yet, this relatively minor difference results in one very good and one very bad optimum. To quantify this behavior, this section employs a Monte Carlo sampling methodology to compute the "chance at a high-performing optimum" P_{HP}.

In total, 16 frequency ranges $\Delta\omega$ between 50 Hz and 8000 Hz are considered: each has a width of 500Hz except the first range, [50Hz, 500Hz], to prevent analysis at 0Hz. For each frequency range, the baseline

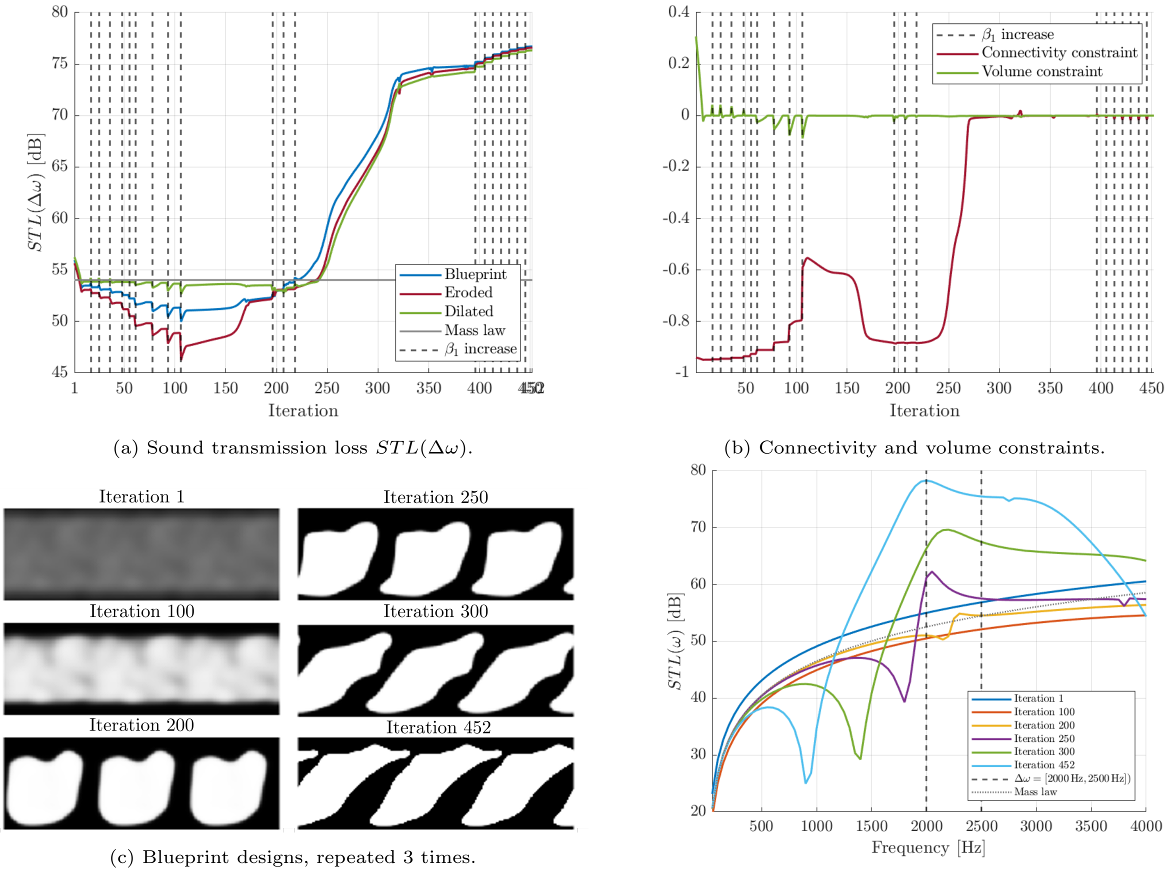

(a) Sound transmission loss $STL(\Delta\omega)$.

(b) Connectivity and volume constraints.

(c) Blueprint designs, repeated 3 times.

(d) Sound transmission loss $STL(\omega)$ of blueprint designs.

Figure 4: Baseline optimization run yielding a high-performing optimum, for $\Delta\omega = [2000\,\text{Hz}, 2500\,\text{Hz}]$.

implementation is run $N = 20$ times starting from a random initial guess $\boldsymbol{\xi}_i$, $i \in [1, \ldots, 20]$. These initial guesses were generated a priori using Matlab's `rand` function.

After the Monte Carlo sampling and subsequent optimization process finishes, the performance $STL^* = \min(STL_b(\Delta\omega), STL_e(\Delta\omega), STL_d(\Delta\omega))$ is compared to the performance of the mass law $STL_{ml}(\Delta\omega)$, i.e.,the performance of a single plate with equal mass. If $STL^* \geq 1.1 \times STL_{ml}$, then it is considered to be a *high-performing* optimum. In contrast, if $STL^* < 1.1 \times STL_{ml}$, then it is seen as a *low-performing* optimum. Note that although the cutoff of 10% above the mass law is arbitrary, subsequent results will show that the division between high- and low-performing optima is relatively insensitive to the exact cutoff value. In other words, the two categories are easily separable for most frequency ranges.
For a frequency range $\Delta\omega$, the number of high-performing optima, N_{HP}, is compared to the number of low-performing optima, N_{LP}, to provide an estimate $\tilde{P}_{\text{HP}}$:

$$P_{\text{HP}} \approx \tilde{P}_{\text{HP}} = 100 \frac{N_{HP}}{N_{LP} + N_{HP}} \, [\%]. \tag{7}$$

Computational costs limit this work to 20 samples per frequency range, which in turn causes relatively high error bars for $\tilde{P}_{\text{HP}}$. Therefore, frequency ranges $\Delta\omega$ are classified as follows:

- $\tilde{P}_{\text{HP}} \leq \sim 20\%$: low-performing region,
- $\tilde{P}_{\text{HP}} \geq \sim 80\%$: high-performing region,

- $\tilde{P}_{\mathrm{HP}} \in]\ \sim 20\%,\ \sim 80\%[$: transition region.

For more details regarding the Monte Carlo sampling strategy, the reader is referred to Appendix B.

Figure 5 shows the performance of the optima found by applying the sampling methodology to the baseline implementation. Looking at $\tilde{P}_{\mathrm{HP}}$, the problem at hand is stark: for $\Delta\omega$ below 2000 Hz, *all* of the optima found are low-performing and for $\Delta\omega$ above 4500 Hz, $\tilde{P}_{\mathrm{HP}}$ frequently dips below 50%. The frequency ranges can be divided into three regions: a wide low-performing region for all frequency ranges below 2500 Hz, followed by a high-performing region for $2500\,\mathrm{Hz} \leq \Delta\omega \leq 4500\,\mathrm{Hz}$ and then a transition region for $\Delta\omega \geq 4500\,\mathrm{Hz}$.
Figure 6 shows a selection of the optimal designs found by the Monte Carlo sampling. For every $\Delta\omega$, it visualizes two low-performing designs: the worst one in red and the best one in yellow. Additionally, for the high-performing optima, it shows the worst one in green and the best one in blue. A closer look at this figure reveals that the low-performance and transition regions are caused by strongly attractive classes of optima. The low-performing region is filled with variants on the same design, namely a rigid vertical connection between top and bottom plate, which shows the strong attraction of mass-driven optima in that region. Similarly, $\Delta\omega$ ranges between 2000 Hz and 5000 Hz see a strong attraction of another, high-performing type of design: an askew connection between top and bottom plate. For $\Delta\omega = [2000\,\mathrm{Hz}, 2500\,\mathrm{Hz}]$, this design is the (only) high-performing design. For $\Delta\omega$ ranges in $[3000\,\mathrm{Hz}, 4500\,\mathrm{Hz}]$, the design appears as the worst high-performing design, with Figure 5 showing that its performance is inversely proportional with $\Delta\omega$. Finally, the design reappears as a *low*-performing design for $\Delta\omega = [4500\,\mathrm{Hz}, 5000\,\mathrm{Hz}]$. This proves the existence of a second family of optima, the askew connections, whose performance both exceeds and trails the mass law based on the targeted frequency range.
In addition to the mass-driven optima, the optimal design space seems to contain many more distinct families, each with their own performance benefits and trade-offs. Thus, escaping from an optimum belonging to a "bad" family to reach a different family is crucial.

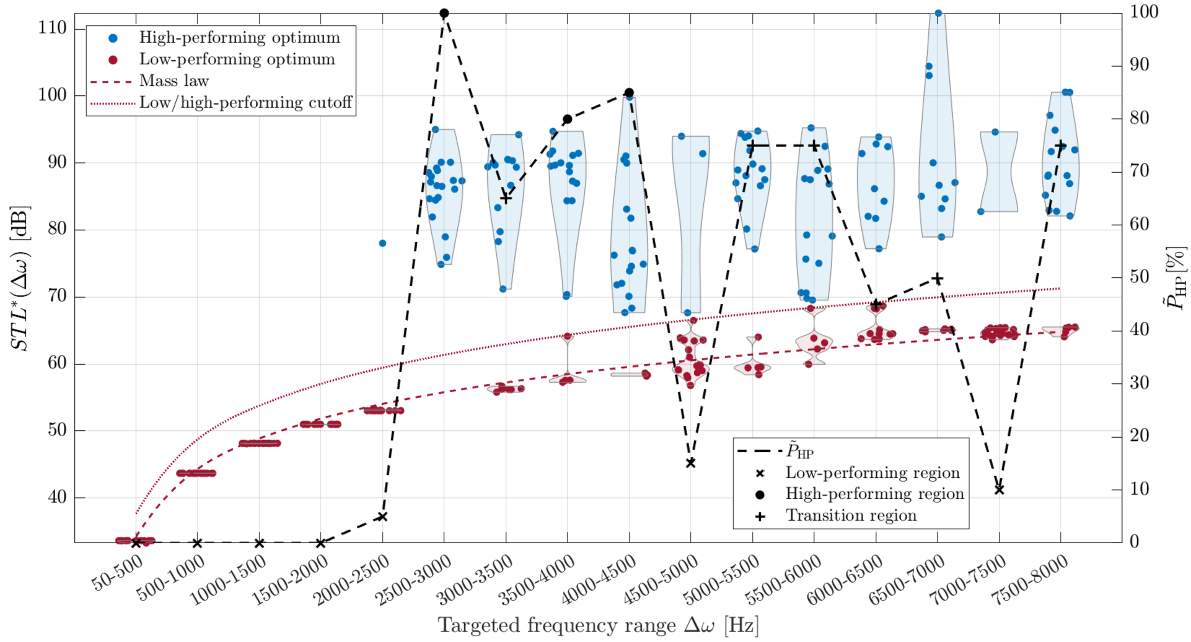

Figure 5: Results of applying the Monte Carlo sampling methodology to the baseline implementation. Performances STL^* of the optimized designs (left y-axis) and estimated chance at a high-performing optimum (right y-axis) as a function of the optimized frequency range $\Delta\omega$ (x-axis). Note that all scatter plots were made with the Violin function, which adds a small, artificial frequency shift to samples for enhanced visibility [35].

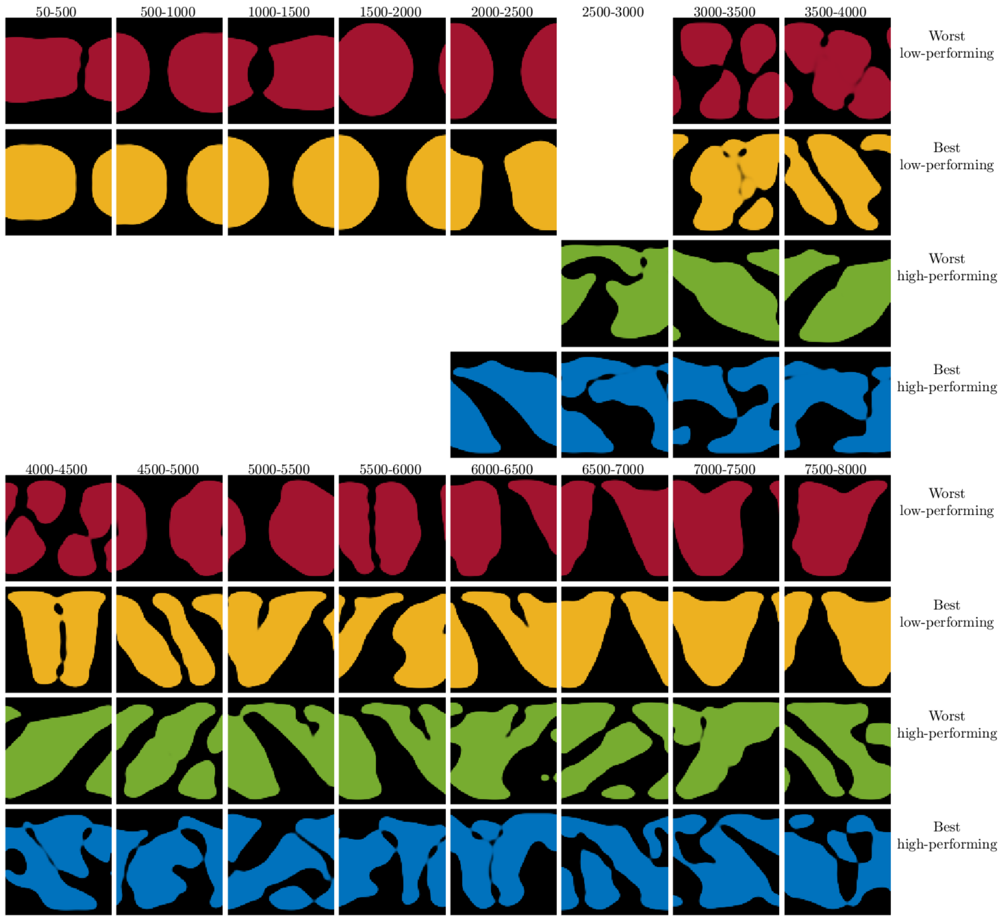

Figure 6: Worst- and best-performing designs of the low- and high-performing optima for the baseline implementation.

3.3. Explaining the numerical attractiveness of mass-driven optima

Based on the above observations, the baseline routine's strong tendency to generate mass-driven optima in the low-performing region follows from three causes.

1. **Funneling effect**: every random initial guess $\boldsymbol{\xi}_i$ is pushed towards a rigid, straight plate connection by the connectivity constraint and the β increases that drive the design towards 0-1.
2. **Low attraction of askew connections due to symmetry**: in order to move from a mass-driven optimum to one that transforms horizontal into vertical motion, the design must change from a straight to a slanted connection. This means rotating the straight connection, either clockwise or counterclockwise. Due to the plane incidence wave arriving at an angle of $\theta = 0°$, both rotations have the same effect on the objective function. This favors further convergence towards a mass-driven optimum.
3. **Low attraction of resonance peaks due to flatness of the local optima**: modal density generally increases with frequency. At high frequencies, this introduces many steep descent directions that an optimizer can exploit. It is thus more likely to fall into one of these narrow valleys during the optimization process.

It is currently unclear which of these causes is most critical and if any additional causes are present. However, more to the point, any of these causes can explain why the optimizer often converges to mass-driven optima, which aids in understanding and devising strategies to prevent this from happening.

4. Investigation of strategies for finding a high-performing optimum

This section investigates three strategy categories for improving the chance of a high-performing local optimum. The sampling method introduced in the previous section is reused to ensure an equal comparison.

The strategies are judged against the following four criteria, in decreasing order of importance.

1. The **chance at a high-performing optimum**, approximated by $\tilde{P}_{\mathrm{HP}}$.
2. **The average STL of the high-performing optima**, $STL_{\mathrm{HP}} = \sum_{i=1}^{N_{\mathrm{HP}}} STL_i(\Delta\omega)/N_{\mathrm{HP}}$. This captures the goal that high-performing optima should still be as high-performing as possible. Note that STL_{HP} is undefined for the $\Delta\omega$ ranges without high-performing optima ($N_{HP} = 0$).
3. The **bias** towards a (predefined) class of optima. For example, a strategy that uses an initial guess based on previously found optima scores poorly on this criterion. This criterion is judged qualitatively by interpreting the produced designs.
4. The **computation time**. Because the proposed strategies do not increase the time per optimization iteration, this criterion is expressed by the average number of optimization iterations for the high-performing optima $N_{\mathrm{iter}}^{\mathrm{HP}} = \sum_{i=1}^{N_{\mathrm{HP}}} N_{\mathrm{iter},i}/N_{\mathrm{HP}}$. If $N_{\mathrm{HP}} = 0$, then $N_{\mathrm{iter}}^{\mathrm{HP}}$ is undefined.

This section is built up as follows. Section 4.1 studies the effect of exclusion methods that aim to make low-quality local optima infeasible. For the application in this study, an upper bound on the stiffness is introduced to make rigid connections infeasible and force the optimizer to produce compliant structures. Section 4.2 investigates frequency shifts: by starting from a high targeted frequency range and then decreasing it slowly, the greater density of high-performing optima at high frequencies is transplanted to low frequencies. Section 4.3 examines a set of relaxation strategies that delay the required robustness against design changes. Finally, Section 4.4 provides a comparison of the studied strategies, highlighting advantages and disadvantages along all four considered criteria.

4.1. Lower bounds on the connectivity constraint

This section explores the use of a lower bound, $J_{\mathrm{conn}}^{\mathrm{min}}$, on the connectivity constraint to ensure (see Section 2.2.3) that $-1 < J_{\mathrm{conn}}^{\mathrm{min}} \leq J_{\mathrm{conn}} \leq 0$. Physically, this means that the optimal structure is *required* to be less stiff against a self-weight load. The motivation for this strategy is the observation that high-performing optima have an active connectivity constraint ($-|\epsilon| < J_{\mathrm{conn}} < 0,\ \epsilon \approx 0$) whereas low-performing optima are mass-driven and hence very rigid. By requiring compliance, the low-performing optima are made infeasible. Therefore, this method is an exclusion method.
Arguably, this method is a priori violating the first secondary goal: by restricting the connectivity, stiff solutions are excluded and the optima are biased to be (more) compliant. If stiff, high-performing optima exist, then this strategy will never discover them.
Three variants of the proposed exclusion (E) strategy are considered.

- E_1: A constant, loose bound $J_{\mathrm{conn}}^{\mathrm{min}} = -0.5$.
- E_2: A constant, tight bound $J_{\mathrm{conn}}^{\mathrm{min}} = -0.05$.
- E_3: An adaptively increasing bound, starting at $J_{\mathrm{conn}}^{\mathrm{min}} = -1$ and going to -0.05, denoted as $J_{\mathrm{conn}}^{\mathrm{min}} \to -0.05$.

For the last variant, E_3, changes in $J_{\mathrm{conn}}^{\mathrm{min}}$ are triggered after $\beta_1 > 15$, allowing for spontaneous convergence to high-performing optima since these (see Section 3.1) usually occur when $\beta_1 < 15$. Changes in $J_{\mathrm{conn}}^{\mathrm{min}}$ are applied by looking at the current value of J_{conn}: $J_{\mathrm{conn}}^{\mathrm{min}} = \min_{n\in\mathbb{N}^+} -0.05n \quad \text{s.t.} \quad J_{\mathrm{conn}}^{\mathrm{min}} > J_{\mathrm{conn}}$. Thus, if the optimizer finds an optimum with an active connectivity constraint ($J_{\mathrm{conn}} \approx 0$) then $k = 1$ and $J_{\mathrm{conn}}^{\mathrm{min}}$ is set immediately to -0.05.

Figure 7 provides a comparison of the Monte Carlo sampling results for each variant. Figure 7a visualizes $\tilde{P}_{HP}$ for the baseline and all three exclusion strategy variants. Variant E_1 both reduces the low-performing

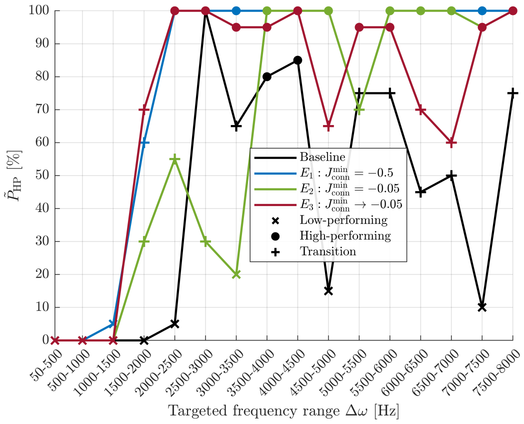

(a) Chance of obtaining a high-performing optimum $\tilde{P}_{\mathrm{HP}}$.

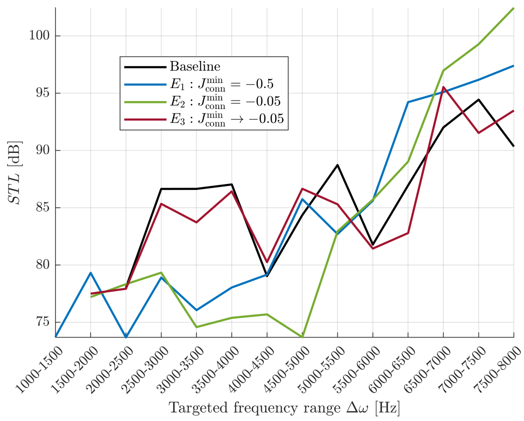

(b) Average STL^* of the high-performing optima for all targeted frequency ranges.

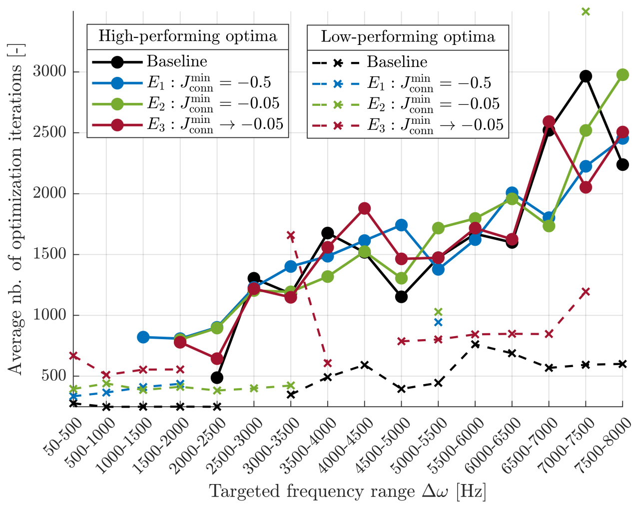

(c) Average number of optimization iterations for the low- and high-performing optima.

Figure 7: Overview of the methodology's results for the considered exclusion strategies, E_1 to E_3.

region by 1000 Hz and achieves $\tilde{P}_{HP} = 100\%$ for all targeted frequency ranges above 2000 Hz except for [5000 Hz, 5500 Hz]. Variant E_2 performs similar to E_1 above 3500 Hz but replaces the baseline's high-performing region around 2500 Hz with a transition region. Variant E_3 does the opposite: it performs equally well as E_1 below 3500 Hz but has some transition regions above 3500 Hz. Thus, regarding the first criterion, exclusion variants can clearly outperform the baseline implementation, with variant E_1 outperforming E_2 and E_3.

The conclusion is more mixed when considering the second criterion, i.e., the STL^* of the high-performing optima. Figure 7b shows that variant E_3 has comparable performance to the baseline: the average STL^* of the high-performing optima is not far from the baseline. In contrast, variants E_1 and E_2 differ significantly depending on the frequency range: around 3000 Hz the high-performing optima are $5 - 10$ dB worse but at high frequencies they start outperforming the baseline by the same margins.
Regarding the third criterion, i.e., bias, every exclusion variant a priori has a strict bias away from stiff designs. Between the variants, E_1 has a looser bound on the maximum stiffness than E_2 and E_3 and thus

less bias.

Regarding the fourth and final criterion, computation time, theoretically none of the strategies incur a much higher computational cost per iteration. Furthermore, Figure 7c shows that, except for some outliers due to a low number of samples, the number of optimization iterations is always comparable to the baseline. This number tends to increase linearly with the frequency. There is no immediate explanation for this behavior, but it is hypothesized that the larger modal density yields more low-performing local optima that the optimizer must travel around.

4.2. Frequency range continuation

This section investigates the direct frequency shift strategy by changing the targeted frequency range $\Delta\omega$ during the optimization process. The motivation for this strategy is that Section 3.2 reveals a low-performing region occurs at low frequencies. A natural idea is then to start at a higher frequency range $\Delta\omega^* = [\omega_- + \omega^*, \omega_+ + \omega^*]$, with $\omega^* > 0\,\text{Hz}$ the excess frequency, and slowly remove the excess after a high-performing optimum is found. If the high-performing, high-frequency optimum is not "lost" during this process, the larger $P_{\mathbf{HP}}$ at high $\Delta\omega$ is transplanted to the low-frequency, low-performing region.

Four variants of the frequency shift (F) strategy are tested: $\text{F}_n : \omega^* = 1000n\,\text{Hz}, n \in [1, \ldots, 4]$. The TO routine starts with β continuation until an intermediately converged, high-performing design is found: $STL^*(\Delta\omega + \omega^*) \geq 1.15 \times STL_{ml}(\Delta\omega + \omega^*)$. Note that the mass law also depends on the targeted frequency range. If this criterion is met, or if $\beta_1 > 50$, frequency continuation is initiated. The excess frequency ω^* is reduced by 100 Hz (or 1/5th of the width of the considered frequency range) every time convergence is achieved. After $\omega^* = 0$, β continuation resumes.

Figure 8 shows the results of the four frequency continuation strategies.
Figure 8a allows to evaluate the first criterion, $\tilde{P}_{HP}$, for each frequency shift variant. The baseline's low-performing region from $\Delta\omega = [50\,\text{Hz}, 500\,\text{Hz}]$ to $[2000\,\text{Hz}, 2500\,\text{Hz}]$ is decreased, depending on the size of the excess frequency ω^*. That is, F_1 reduces it by its excess frequency of 1000 Hz and for variants F_2-F_4 only $\Delta\omega = [50\,\text{Hz}, 500\,\text{Hz}]$ remains low-performing. This reduction of the low-performing region is part of a larger shift of the baseline's $\tilde{P}_{HP}(\Delta\omega)$ curve. For frequency shift variant F_i, one sees the shift $\tilde{P}^{F_i}_{HP}(\Delta\omega) \approx \tilde{P}^{BL}_{HP}(\Delta\omega + \omega^*)$. Only when a high-performing optimum is "lost" during the frequency continuation steps does this relation cease to hold. This is mostly the case for the lowest frequency range $\Delta\omega = [50\,\text{Hz}, 500\,\text{Hz}]$, which implies that this is most likely a physical limit and not a numerical one. In other words, with the current setup there are likely no high-performing optima in this frequency range. Finally, given the conclusion above, Figure 8a reveals the existence of a high-performing region above 8000 Hz, starting most likely at $\Delta\omega = [8500\,\text{Hz}, 9000\,\text{Hz}]$.

Evaluation of the second criterion via Figure 8b shows that frequency shift methods generally outperform the baseline at higher frequencies (above ±4000 Hz). At low frequencies, variants F_1 and F_2 also see the leftward shift of the baseline's excellently performing designs (±85 dB) between 2500 and 4000 Hz, causing them to outperform variants F_3 and F_4 below 3000 Hz.
Regarding computation time, measured via the number of optimization iterations shown in Figure 8c, a clear and marked increase with respect to the baseline is noticeable. This is due to the additional continuation steps: variant F_n sees $10n$ frequency continuation steps on top of the usual β continuation, each of which requires convergence. This also explains why F_4 has the largest overall computation cost of all variants. The size of the increase is also noticeable: a doubling or even tripling of the computational cost is not unusual. Speedups could possibly be achieved through 1) more advanced ω^* reductions, e.g., based on the accompanying STL reduction, and 2) loosened convergence tolerances during frequency continuation steps. These are, however, out of the scope of this work. The final note regarding the computational cost of frequency shift strategies is that it is not always worth it: no high-performing optima were found for

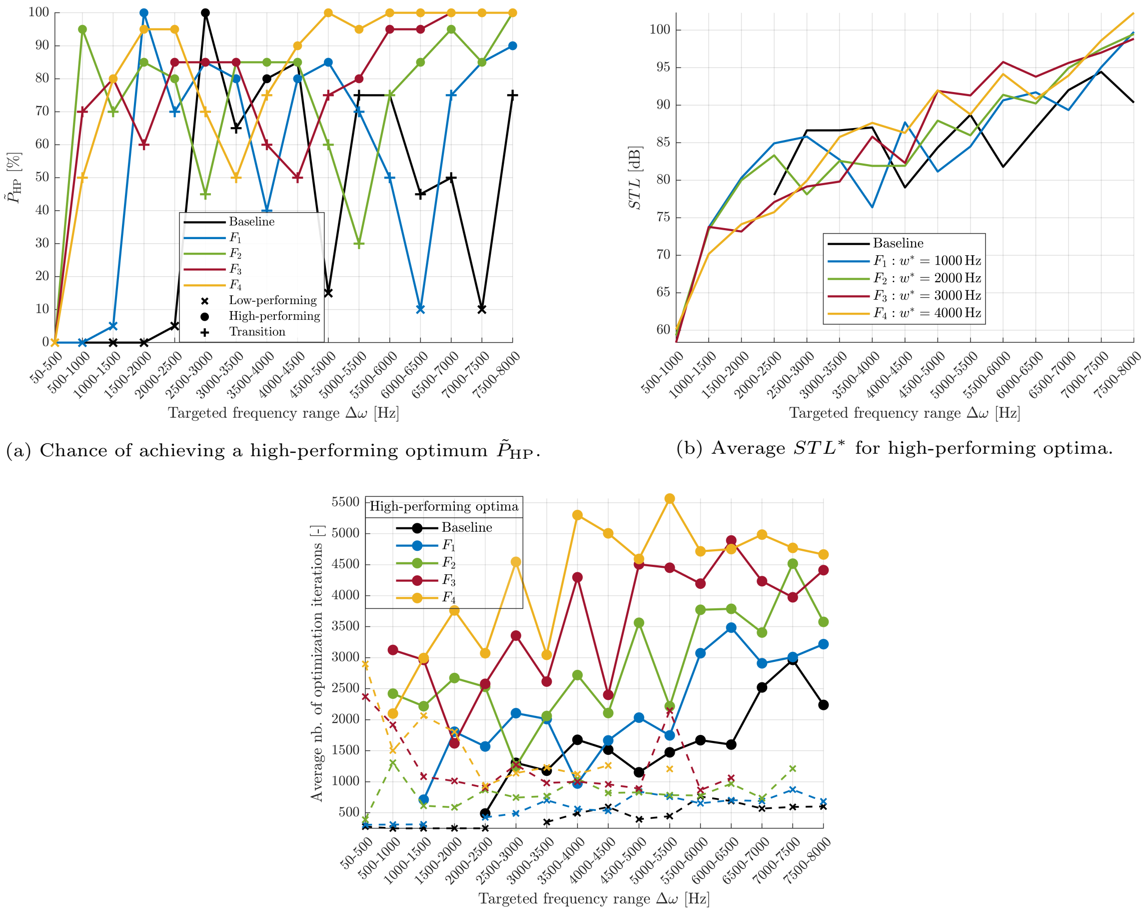

(a) Chance of achieving a high-performing optimum $\tilde{P}_{\mathrm{HP}}$.

(b) Average STL^* for high-performing optima.

(c) Average number of optimization iterations for the low- and high-performing optima. Low-performing optima curves are dashed.

Figure 8: Overview of the methodology's results for the four variants of the frequency shift strategy, F_1-F_4.

$\Delta\omega = [50\,\mathrm{Hz}, 500\,\mathrm{Hz}]$ but the number of iterations more than quadrupled for variants F_3 and F_4.

Regarding bias, frequency continuation strategies aim to draw high-frequency solutions towards lower frequencies and hence bias towards high-frequency optima. However, this is not guaranteed to be successful. Figure 9 shows an optimization run with variant F_4, targeting $\Delta\omega = [50\,\mathrm{Hz}, 500\,\mathrm{Hz}]$. Four stages can be discerned. In the first stage (iteration $1 - 665$), the usual optimization process is visible (cf. Figure 4): stagnation for a few β increases until a straight connection is formed (iteration 150) and then sharp improvement when this connection becomes slanted (iteration 400). The STL improvement triggers the mass law criterion $STL^* \geq 1.15 \times STL_{ml}$, which causes a frequency continuation step at iteration 665. At this iteration, the design contains a double askew connection with a decoupling frequency below 2000 Hz. The second stage (iteration $665 - 2500$) sees frequency continuation steps that pull down the targeted frequency range from around 4000 Hz to around 1000 Hz. The peak of the $STL(\omega)$ curves move correspondingly to lie within this frequency band and the decoupling frequency is also pulled down to remain below the targeted frequency band. This keeps the $STL(\Delta\omega)$ above the mass law until about iteration 2500, where the decoupling frequency is only narrowly below $\Delta\omega$. The third stage (iteration $2500 - 3000$) then combats this performance drop. The two slanted connections melt together into a connection with a tuned rotational inertia. This again pulls down the decoupling frequency, causing great performance improvement. However, stage four (iteration $3000 - 3337$) sees frequency continuations for which no corresponding decoupling frequency drop can be found. The performance again drops to the mass law and the routine terminates with several β increases. The optimization process of Figure 9 thus shows a fallback to mass-law performance, not once but twice. In conclusion, frequency shift strategies can bias towards suboptimal designs because they start from high-frequency solutions that cannot be adapted to fit low-frequency targets. Smaller and more careful frequency continuation could combat the fallback to the mass law, but not the underlying bias towards solutions that exploit the decoupling frequency. In other words, the considered frequency shift variants are unlikely to produce high-performing designs with a decoupling frequency *above* $\Delta\omega$.

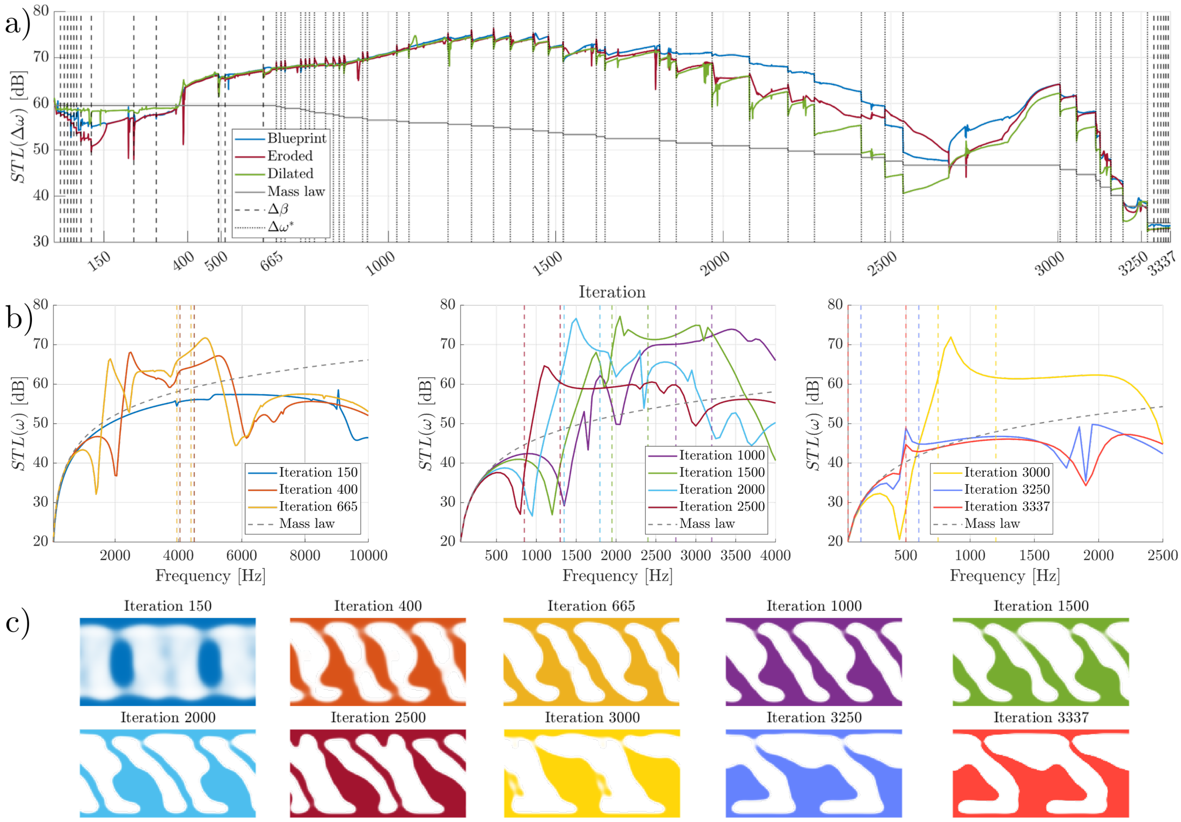

Figure 9: Optimization process overview of a run targeting $\Delta\omega = [50\,\text{Hz}, 500\,\text{Hz}]$ with frequency continuation ($\omega^* = 4000\,\text{Hz}$). a) Progress of $STL_b(\Delta\omega)$, $STL_e(\Delta\omega)$ and $STL_d(\Delta\omega)$. b) $STL_b(\omega)$ for selected blueprint designs throughout the optimization. The vertical lines denote the targeted $\Delta\omega$ range for the iteration with the same color. c) Selected blueprint unit cell layouts, repeated twice.

Overall, frequency shift strategies are great at reducing low-performing regions at low frequencies. However, the right value for the excess frequency ω^* is not known a priori and a wrong choice can even lead to a lowered P_{HP}, such as when variant F_2 replaces the baseline's high-performing region at $\Delta\omega = [2500\,\text{Hz}, 3000\,\text{Hz}]$ with a transition region. This is doubly regrettable since frequency shifting incurs a significantly larger computational cost since every frequency continuation step requires convergence, increasing the number of iterations.

4.3. Design robustness and formulation

The third strategy tested in this work involves the relaxation of the design robustness. This is motivated by Figure 10, which shows the $STL(\omega)$ curves of the blueprint, eroded and dilated designs of some of the intermediate solutions of the optimization run of Figure 4. All $STL(\omega)$ curves show a trough followed by a peak. The eroded designs are shifted leftwards due to their lower mass and vice versa for the dilated designs. For iteration 200, this causes the blueprint design to have a peak and trough in the middle of the targeted frequency range, whereas the eroded and dilated designs have their peaks and troughs left and right of that range, respectively. The design change from iteration 200 to iteration 300 sees an increase of $STL(\Delta\omega)$ for all designs, but this is not self-evident because 1) the trough of the dilated design comes into the targeted frequency range while 2) (the tail end of) the peak of the eroded design moves out. Both effects cause a decrease of the performance of the worst design. However, the decoupling effect compensates for this and still causes an overall improvement.

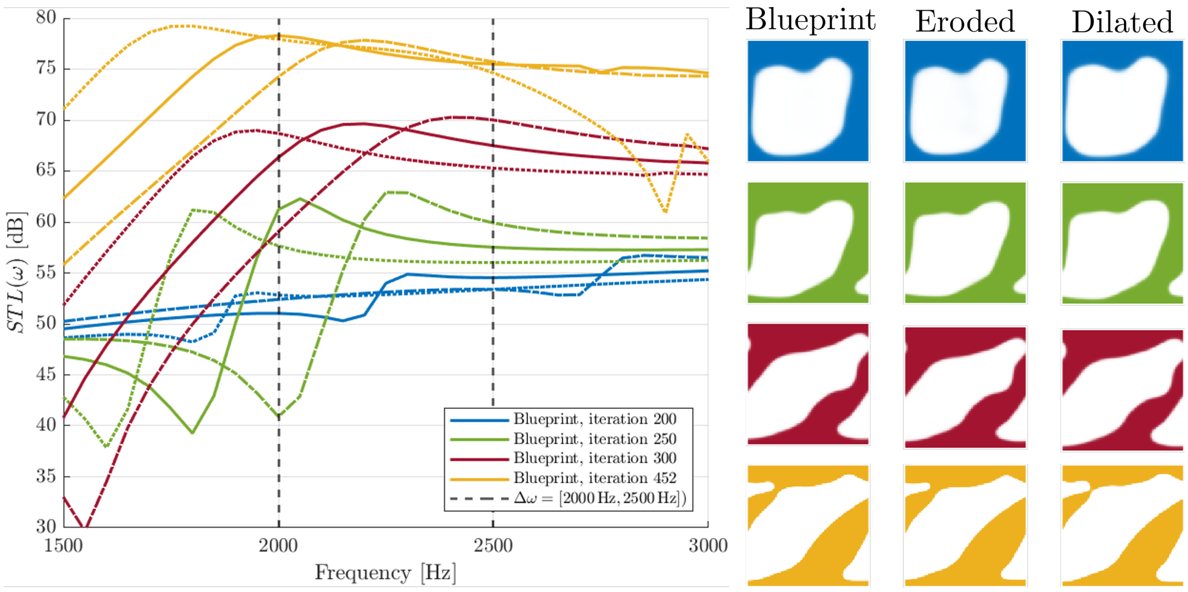

Figure 10: Eroded, blueprint and dilated designs (right) and corresponding $STL(\omega)$ curves around the targeted frequency range $\Delta\omega$ for the successful optimization run of Figure 4. Eroded and dilated STL curves are dotted and dashdotted, respectively.

Based on these findings, it is hypothesized that the tendency of the minmax formulation to improve the worst-performing design causes a bias towards designs for which $STL_b(\Delta\omega) \approx STL_e(\Delta\omega) \approx STL_d(\Delta\omega)$, which is largely true for mass-driven optima. Furthermore, a minmax formulation disallows large improvements of one design if they are accompanied by even a slight performance reduction of the worst design. Designs are prone to this phenomenon if the targeted frequency range lies exactly between the STL decoupling peaks of the eroded and dilated designs, which might explain the baseline implementation's tendency to generate askew connections.

Two solutions are proposed. First, the minmax formulation can be replaced by an aggregated formulation, i.e., a formulation where the sum of the blueprint, eroded and dilated performance is optimized.

$$\min_{\boldsymbol{\xi}\in\mathbb{R}^{N_e}} \quad J_b + J_e + J_d, \tag{8}$$

where the three constraints are omitted for brevity. The aggregated formulation thus allows improvements in $\max STL_i(\Delta\omega)$, $i \in [\mathrm{b,e,d}]$ as long as it compensates for the decrease in $\min STL_i(\Delta\omega)$, $i \in [\mathrm{b,e,d}]$.

Second, the design robustness can be delayed. That is, only the blueprint design is considered at the start of the optimization and the eroded and dilated designs are introduced *after* a high-performing optimum is found. This introduction of the eroded and dilated designs is done gradually by setting $\eta_e = \eta_b - \Delta\eta$, $\eta_d = \eta_b + \Delta\eta$ and slowly increasing $\Delta\eta$ from 0 to its default value of 0.1. The volume and connectivity constraint are still computed on the original dilated and eroded designs, respectively, to retain consistency with the baseline implementation.

The two proposed solutions are now combined to form three different relaxation (R) strategies, R_1, R_2 and R_3.
Strategy R_1 starts with the aggregated formulation, maximizing the sum of the performances of the blueprint, eroded and dilated designs. To ensure consistency with the baseline implementation, the minmax formulation is reintroduced after the last β continuation step. Strategy R_1 is denoted by "$\sum \rightarrow$ minmax".
Strategy R_2 starts by considering only the blueprint STL. Then, design robustness (with low $\Delta\eta = 0.02$) with the aggregated formulation is introduced after the criterion of Section 4.2 is met: a converged design with $STL^* \geq 1.15 \times STL_{ml}$ or $\beta > 50$. Next, the design robustness increases gradually: $\Delta\eta$ goes from 0.02 to 0.1 in four steps of 0.02, with each step requiring convergence (see Section 2.2.5). Finally, just like strategy A, the formulation switches from aggregated to minmax after the last β increase. Strategy R_2 is denoted by "$\sum \rightarrow \Delta\eta \uparrow \rightarrow$ minmax".
Strategy R_3 employs the same $\Delta\eta$ continuation strategy as strategy R_2 but uses the minmax approach throughout the whole optimization process. All three strategies end with the same design robustness and formulation as the baseline approach, allowing for a fair comparison. Furthermore, strategy R_1 differs from strategy R_2 only in the $\Delta\eta$ continuation and strategy R_2 differs from strategy R_3 only in the formulation. This allows for an unbiased conclusion as to the effect of 1) the design robustness and 2) the formulation on the chance at a high-performing optimum.
Strategy R_3 is denoted by "minmax $\rightarrow \Delta\eta \uparrow$".
Figure 11 shows the results of the three relaxation strategy variants. Figure 11a shows that although none of the three variants are able to decrease the width of the low-frequency, low-performing region, both R_2 and R_3 yield a continuous high-performing region above 3000 Hz. On the other hand, R_1 shows no improvement in $\tilde{P}_{HP}$.
Figure 11b shows that the performance of the high-performing optima is best for variant R_3, which always outperforms the baseline, albeit sometimes narrowly. Variant R_2 has, on average, comparable performance to the baseline: slightly worse at low frequencies and slightly better at higher frequencies. Finally, variant R_1 is often at least 5 dB worse than the baseline.

Figure 11c shows the cost of the relaxation strategy variants. Whereas R_1 generally requires less iterations than the baseline, R_2 requires slightly more and R_3 sometimes even sees more than a doubling of the number of iterations. This is caused by the larger number of continuation steps: every increase of η requires convergence before another step is taken.
No clear, definitive conclusions regarding bias are apparent. A study of the optimal designs reveals a wide array of types and mechanisms. Clearly, the initial absence of design robustness of variants R_2 and R_3 allow them to better navigate towards better designs, but a more detailed classification and analysis is outside the scope of this work.
Overall, it is clear that relaxation strategies can strongly improve both the chance at a high-performing optimum and the performance of those optima. Of the three variants, R_3 comes out on top. This variant sets $\Delta\eta = 0$ until a high-performing optimum is found. The associated computational cost increases due to the additional continuation steps, but only by at most a factor of two, making it more than worth it.

4.4. Comparison

The three strategy classes (exclusion, frequency shift and relaxation) are now compared for each of the four evaluation criteria. A summary is provided by Table 1.

Regarding P_{HP}, only frequency shift techniques are able to significantly reduce the low-performing region at low frequencies until a physical limit is reached, namely, the decoupling frequency that cannot be further reduced. Exclusion strategies are also able to (slightly) reduce this region, whereas relaxation regions do not have an affect. However, regarding the transition regions, the conclusion is mostly reversed: for poor choices of ω^*, frequency shift techniques seriously degrade P_{HP}. On the other hand, exclusion and relaxation strategies turn these regions into high-performing regions.

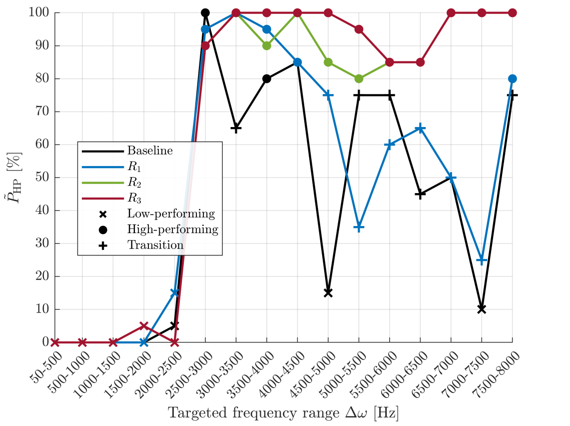

(a) Chance of achieving a high-performing optimum $\tilde{P}_{\mathrm{HP}}$ (y-axis) as a function of the optimized frequency range $\Delta\omega$ (x-axis).

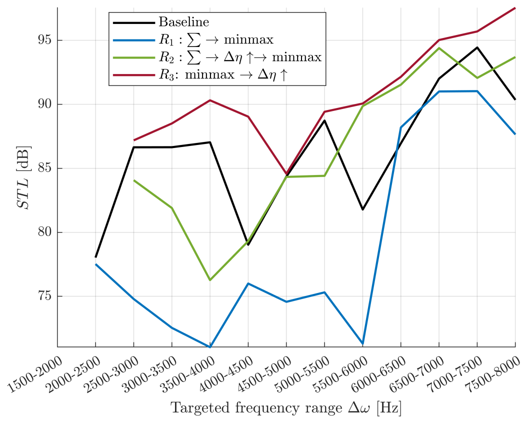

(b) Average STL^* for low- and high-performing optima for all targeted frequency ranges.

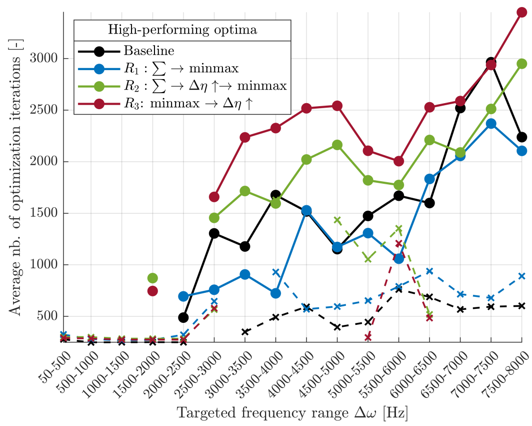

(c) Average number of optimization iterations for the low- and high-performing optima.

Figure 11: Monte Carlo sampling results for the relaxation strategy variants, R_1, R_2 and R_3.

Regarding STL_{HP}, exclusion and frequency shift strategies have a mixed effect: improvement at high frequencies and vice versa for low frequencies. However, relaxation strategy variant R_3 causes a general improvement of the STL across the whole studied frequency range.

Regarding bias, it is clear that exclusion strategies bias towards compliant designs since stiff designs are a priori excluded. Frequency shift strategies bias towards high-frequency solutions adapted for the target frequency and relaxation strategies only show a bias away from the undesired mass-driven designs.
Finally, regarding computation cost, exclusion and relaxation strategies do not significantly increase the computation time. On the other hand, frequency shift techniques can double or even triple the number of optimization iterations. Careful convergence and continuation management is proposed to mitigate this disadvantage.

5. Conclusions

This work presented an overview and comparison of strategies that aim to maximize P_{HP}, the probability that a topology optimization routine converges to a high-performing optimum. As a representative benchmark, the study is applied to the STL optimization of a sandwich plate, due to the general property of such

Table 1: Strategy comparison across the chosen evaluation criteria.

Strategy	Criteria			
	P_{HP}	STL_{HP}	Bias	Computational cost
Exclusion	Turns transition regions into high-performing regions and reduces low-performing region.	Improvement at higher $\Delta\omega$, decrease at lower $\Delta\omega$.	Towards compliant designs.	No difference with baseline.
Frequency shift	Reduces low-performing region at low $\Delta\omega$. Can replace high-performing regions with transition regions for poor ω^* choice.	Improvement at higher $\Delta\omega$, decrease at lower $\Delta\omega$.	Towards optima of higher $\Delta\omega$.	Often doubles or triples cost, depending on setup.
Relaxation	Turns transition regions into high-performing regions. No influence on low-performing, low-frequency region.	Can improve performance across all frequencies (variant R_3).	/	Modest (usually $< 2\times$) increase in cost.

vibro-acoustic topology optimization problems to converge to low-performing, mass-driven solutions. Two distinct subproblems were identified: a low-frequency region completely absent of any high-performing optima ($P_{HP} = 0\%$) and a mid-frequency region where P_{HP} was often below 50%. For each strategy classes, novel variant were developed for the chosen benchmark, tested and compared to see if they could improve upon this behavior.

The first class are exclusion strategies. They add constraints to the optimization problem to make the class of low-performing optima infeasible. In general, this requires detailed knowledge of the optimum class one aims to avoid. For the considered benchmark, it was found that low-performing optima maximize the sandwich panel stiffness and so an upper bound was placed on this value. The second class contains frequency shift methods: by first optimizing for a frequency range higher than the target and then gradually reducing the excess frequency, high-performing solutions are transplanted to regions that lack such designs.
The third strategy are relaxation methods, i.e., methods that first solve a “simpler” formulation and use the solution as a starting point for the actual problem. In this work, this was implemented by initially relaxing the requirement for design robustness.

The three strategy classes were compared by applying a Monte Carlo sampling method: for a prespecified set of frequency ranges and a random initial guesses, the optimization routine was run and the final design was recorded. The results were used to see if strategies could 1) increase P_{HP}, 2) increase the STL of the high-performing optima, 3) do so without introducing bias and 4) do so at a reasonable computational cost. It was found that frequency shift techniques are particularly suitable to increase P_{HP} in the low-frequency region, but at a doubling or even tripling of the computational cost and a bias towards high-frequency designs. On the other hand, exclusion and relaxation techniques can increase P_{HP} to 100% in the mid-frequency region for only a limited increase in computation time.

While the optimal strategy is problem-dependent, the comparison provided by this work offers researchers the tools to select and adapt suitable methods for their application. Moreover, although not explored in this study, combining multiple strategies may hold promise for further improving performance. For example, a frequency shift strategy to deal with low-frequency regions could be combined with a relaxation method for the mid-frequency region.

Acknowledgements

Resources and services were provided by the VSC (Flemish Supercomputer Center), funded by the Research Foundation - Flanders (FWO) and the Flemish Government. The research of V. Cool (fellowship no. 1213925N) is funded by a grant from the Research Foundation - Flanders (FWO). Internal Funds KU Leuven are gratefully acknowledged for their support. The Research Foundation - Flanders (FWO) is gratefully acknowledged for its support through research grant no. G0F9922N.

Appendix A. Topology optimization framework

This appendix provides further information regarding the optimization framework.

From optimization variables to physical density fields

The routine uses $\boldsymbol{\xi}$ internally as design variables. Via a sequence of smoothing and projection filters, a physical design field is obtained. The smoothing filter [36] is defined as follows:

$$\tilde{\xi}^e = \frac{\sum_{i=1}^{N_e} w(\boldsymbol{x}^e - \boldsymbol{x}^i)\xi^i}{\sum_{i=1}^{N_e} w(\boldsymbol{x}^e - \boldsymbol{x}^i)}, \qquad w(\boldsymbol{x}^e - \boldsymbol{x}^i) = \max\left(0, R - ||\boldsymbol{x}^e - \boldsymbol{x}^i||_2\right), \tag{A.1}$$

with $\boldsymbol{x^k} = (x^k, y^k)$ the center position of element k. The smoothing filter averages the element design variables over a radius R. When this radius overlaps with the periodic boundaries at $x = 0$ and $x = L_x$, the densities of elements at opposite sides of the mesh are also averaged to ensure designs are independent of the unit cell choice. The Heaviside projection [37] pushes the variables towards 0 (fluid) and 1 (solid). It is defined by:

$$\bar{\xi}^e = \frac{\tanh(\eta\beta) + \tanh((\xi^e - \eta)\beta)}{\tanh(\eta\beta) + \tanh((1 - \eta)\beta)}, \tag{A.2}$$

where η is the projection level and β is the projection strength.

Five distinct designs are produced using the above smoothing and projection filters. It is based on the double filtering technique of [33] and the robust formulation of [32] and consists of the following operations:

$$\boldsymbol{\xi} \xrightarrow[R_1]{\text{Eq. } (A.1)} \tilde{\boldsymbol{\xi}} \xrightarrow[\eta_1,\beta_1]{\text{Eq. } (A.2)} \bar{\tilde{\boldsymbol{\xi}}} \xrightarrow[R_2]{\text{Eq. } (A.1)} \tilde{\bar{\tilde{\boldsymbol{\xi}}}} \xrightarrow[(\eta_b,\eta_e,\eta_{e_2},\eta_d,\eta_{d_2}),\beta_2]{\text{Eq. } (A.2)} (\bar{\tilde{\bar{\tilde{\boldsymbol{\xi}}}}}_b, \bar{\tilde{\bar{\tilde{\boldsymbol{\xi}}}}}_e, \bar{\tilde{\bar{\tilde{\boldsymbol{\xi}}}}}_{e_2}, \bar{\tilde{\bar{\tilde{\boldsymbol{\xi}}}}}_d, \bar{\tilde{\bar{\tilde{\boldsymbol{\xi}}}}}_{d_2}) = (\boldsymbol{\xi}_{b,\mathrm{P}}, \boldsymbol{\xi}_{e,\mathrm{P}}, \boldsymbol{\xi}_{e_2,\mathrm{P}}, \boldsymbol{\xi}_{d,\mathrm{P}}, \boldsymbol{\xi}_{d_2,\mathrm{P}}), \tag{A.3}$$

where $\eta_e = \eta_b - \Delta\eta$, $\eta_d = \eta_b + \Delta\eta$, $\eta_{e_2} = \eta_b - \Delta\eta_2$, $\eta_{d_2} = \eta_b + \Delta\eta_2$. The physical blueprint (b), eroded (e) and dilated (d) design fields are used for performance computation and design robustness. The value of η is either kept constant at 0.1 or varied throughout the optimization process until it reaches 0.1. The additional eroded and dilated designs, e_2 and d_2, are used for the connectivity and volume constraints, respectively. $\Delta\eta_{e_2} = 0.4$ and $\Delta\eta_{d_2} = 0.6$ and they remain constant. The filter radii R_1 and R_2 are kept constant at 7 mm (14 element widths) and 3.5 mm (7 element widths), respectively. Finally, to ensure the required plate thickness of 5 mm, the top and bottom row of 10 elements are fixed at $\xi_i = 1$.

Obtaining the system matrices

The RAMP interpolation method of [38] is used in conjunction with a linear interpolation to obtain the elemental Young's modulus E^e, solid density ρ_s^e, fluid bulk modulus κ^e and fluid density ρ_a^e:

$$\begin{cases} E^e(\xi^e_\mathrm{P}) = E_v + \frac{\xi^e_\mathrm{P}}{1+q(1-\xi^e_\mathrm{P})}(E - E_v), \\ \rho^e_s(\xi^e_\mathrm{P}) = \rho_v + \xi^e_\mathrm{P}(\rho_s - \rho_v), \\ \frac{1}{\kappa^e(\xi^e_\mathrm{P})} = \frac{1}{\kappa} + \xi^e_\mathrm{P}(\frac{1}{\kappa_r} - \frac{1}{\kappa}), \\ \frac{1}{\rho^e_a(\xi^e_\mathrm{P})} = \frac{1}{\rho_a} + \frac{\xi^e_P}{1+q(1-\xi^e_\mathrm{P})}(\frac{1}{\rho_r} - \frac{1}{\rho_a}). \end{cases} \tag{A.4}$$

Here, E_v, ρ_v, κ_r and ρ_r are artificial material properties that prevent numerical singularities. The mass, stiffness and coupling matrices of Eq. (3) are computed by summing over all element contributions and using the elemental material properties of Equation (A.4):

$$\begin{aligned}
\mathbf{K}_s &= \sum_{e=1}^{N_e} E^e(x,y)\mathbf{K}_s^e, & \mathbf{M}_s &= \sum_{e=1}^{N_e} \rho_s^e(x,y)\mathbf{M}_s^e, \\
\mathbf{K}_a &= \sum_{e=1}^{N_e} \frac{1}{\rho_a^e(x,y)}\mathbf{K}_a^e, & \mathbf{M}_a &= \sum_{e=1}^{N_e} \frac{1}{\kappa^e(x,y)}\mathbf{M}_a^e, \\
\mathbf{S}_p &= \sum_{e=1}^{N_e} (1-\xi_{\mathrm{P}}^e)\mathbf{S}^e, & \mathbf{S}_u &= \sum_{e=1}^{N_e} \xi_{\mathrm{P}}^e(\mathbf{S}^e)^T.
\end{aligned} \tag{A.5}$$

The method by Jensen [31] is used to model the vibro-acoustic coupling ($\mathbf{S}_u$, $\mathbf{S}_p$, $\mathbf{M}_u$ and $\mathbf{M}_p$) in the core.

Periodicity

Bloch-Floquet periodicity is enforced as follows. Through the periodicity matrix $\mathbf{\Lambda} = \mathbf{\Lambda}(\lambda_x)$, where $\lambda_x = e^{-\mathrm{i}\mu_x} = e^{-\mathrm{i}k_x L_x}$, the nodal DOFs relate to the periodic DOFs vector $\hat{\mathbf{q}}$ as $\mathbf{q} = \mathbf{\Lambda}\hat{\mathbf{q}}$. As shown in Fig. 2b, the periodic DOFs vector only contains the left and interior DOFs. This results in the following modified system of equations (see [39]):

$$(\hat{\mathbf{K}} - \omega^2\hat{\mathbf{M}})\hat{\mathbf{q}} = \hat{\mathbf{e}}, \qquad \hat{\mathbf{K}} = \mathbf{\Lambda}^H\mathbf{K}\mathbf{\Lambda}, \qquad \hat{\mathbf{M}} = \mathbf{\Lambda}^H\mathbf{M}\mathbf{\Lambda}, \qquad \hat{\mathbf{e}} = \mathbf{\Lambda}^H\mathbf{e}. \tag{A.6}$$

STL computation

The STL of the infinitely periodic structure (Equation (2)) is computed with WFE method [40, 7, 26]. The surrounding acoustic halfspaces are modeled with the analytical approach of Yang et al. [41], with continuity conditions to transform structural displacements into pressure information.

Self-weight connectivity constraint

Connectivity between the top and bottom plate is enforced through the use of a connectivity constraint based on self-weight. As mentioned in [18], this has a reasonably small computational cost (in practice much smaller than the cost of computing the STL) and produces designs relatively insensitive to the μ_{sw} hyperparameter, here kept constant at 15.

The following derivation considers only the structural domain, assumes periodicity conditions are dealt with using the periodicity matrix $\mathbf{\Lambda}$ and is based on [42]. The self-weight force is applied downwards (the $-y$-direction) while the bottom of the domain is constrained. The force is computed using the blueprint design and applied to the additional eroded design to have a slightly more strict connectivity constraint. Periodicity boundary conditions are applied in the x-direction, leading to:

$$\begin{cases} \hat{\mathbf{K}}_{e_2,s}\hat{\mathbf{u}}_{sw} = \hat{\mathbf{F}}_{sw} = -\mathbf{N}\boldsymbol{\xi}_{b,P} & \\ \hat{\mathbf{u}}_{sw}(x) = \hat{\mathbf{u}}_{sw}(x+nL_x) & n = \mathbb{Z} \\ \hat{\mathbf{u}}_{sw}(x) = 0 & x \in \Gamma_b \end{cases}, \tag{A.7}$$

where $\mathbf{N} \in \mathbb{R}^{N_s/2\times N}$ maps element design variables to structural nodal forces in the $-y$-direction and $\mathbf{u}_{sw} = \mathbf{\Lambda}\hat{\mathbf{u}}_{sw}$ is the resulting displacement vector under self-weight. The self-weight compliance θ_{sw} is then given by:

$$\theta_{sw} = \hat{\mathbf{u}}_{sw}^T\hat{\mathbf{F}}_{sw}. \tag{A.8}$$

Derivatives

The optimization problem is solved using the Method of Moving Asymptotes [34]. Since this is a gradient-based solver, derivatives for the objectives and constraints are required. For the STL derivatives with respect to the physical design fields, the reader is referred to appendix C1 of [7] for the derivation of the adjoint equations and derivatives. The self-weight compliance is self-adjoint and thus requires no additional linear system solve. A derivation of the connectivity constraint derivative can be found in appendix B.4 of [18], with the only change that the stiffness matrix depends on the blueprint but the force on the eroded design. The derivative of the volume constraint with respect to the physical design field $\boldsymbol{\xi}_{d_2,P}$ is straightforward. Additionally, all derivatives require the use of the chain rule to incorporate the smoothing and projection filters to obtain the derivative with respect to the optimization variable vector $\boldsymbol{\xi}$. The derivation for this operation can be found in appendix C.2 of [7].

Optimization problem formulation

The minmax formulation is solved by rewriting it to a formulation with bounds on the STL:

$$\begin{cases} \min\limits_{\boldsymbol{\xi}\in\mathbb{R}^{N_e},z} \quad z \\ \text{s.t.} \quad -J_b - z \leq 0 \\ \quad -J_e - z \leq 0 \\ \quad -J_d - z \leq 0 \\ \quad v_{d_2,\mathrm{P}}/V - 1 \leq 0 \\ \theta_{\mathrm{sw}}/\mu_{\mathrm{sw}}\hat{\theta}_{\mathrm{sw}} - 1 \leq 0 \\ \quad z \geq 0, \quad 0 \leq \boldsymbol{\xi} \leq 1, \end{cases} \tag{A.9}$$

For brevity, the volume and stiffness constraints are omitted here. The aggregated approach works out-of-the-box and thus requires no extra explanation.

The above formulation corresponds to the formulation of the Method of Moving Asymptotes [34] with $a_0 = 1$ and $a_i = 1$ for the STL constraints and $a_i = 0$ for the connectivity and volume constraints. The constraint penalization constants are $c_i = 1000$ and $d_i = 1$.

Updating the outer move limit

The outer move limit μ_2 restricts the maximum change of the optimization variables between iterations. It is updated due to the following two observations. First, when μ_2 is set too low, the routine converges in a stable but slow manner. Second, when μ_2 is set too high, the routine experiences numerical oscillations: STL values vary drastically between iterations. Thus, a value for μ_2 that is as high as possible while ensuring numerical stability, is desirable. To this end, initially and after every convergence cycle μ_2 is set to 0.05. STL oscillations are detected by checking if $STL_b(\Delta\omega)$, $STL_e(\Delta\omega)$ or $STL_d(\Delta\omega)$ go from $> 0.01\,\mathrm{dB}$ to $< 0.01\,\mathrm{dB}$. If such oscillations occur for five consecutive iterations, μ_2 is set to $\Delta\boldsymbol{\xi}/2$. On the other hand, if $\min(STL_b(\Delta\omega), STL_e(\Delta\omega), STL_d(\Delta\omega))$ increases by more than $0.02\,\mathrm{dB}$ for five consecutive iterations then μ_2 is set to $\min(1.5 * \mu_2, \mu_2^{\max}])$, where $\mu_2^{\max}$ is an adaptive upper bound set to the lowest μ_2 value that caused oscillations for the current stage of hyperparameters.

All in all, this updating scheme allows for an automatic, cheap and intelligent tuning of the move limit μ_2, thus preventing persistent numerical oscillations and enabling fast improvements in the objective function.

Appendix B. Monte Carlo sampling methodology

This section elaborates on the details regarding the Monte Carlo used to gauge the chance at a high-performing optimum.

Number of optimization runs

All optimizations are carried out on the Flemish Supercomputer. A total of 16 frequency ranges, ranging from $[50\,\mathrm{Hz}, 500\,\mathrm{Hz}]$ to $[7500\,\mathrm{Hz}, 8000\,\mathrm{Hz}]$ are considered. All ranges have a width of 500 Hz, except the first to prevent STL evaluation at 0 Hz. As there are a total of 11 strategies considered in this work (including the baseline), a total of $11 \times 20 \times 16 = 3520$ optimization runs were run.

Statistical error

The standard error of proportion when $N_H = N_L = 10$ and $\tilde{P}_{\mathrm{HP}} = 50\%$ is around $\mathrm{SE} = \sqrt{\frac{\tilde{P}_{\mathrm{HP}}(1-\tilde{P}_{\mathrm{HP}})}{N}} \approx 11\%$, meaning the true chance P_{HP} lies within the 95% confidence interval $50\% \pm 1.96 * \mathrm{SE} = [28\%, 72\%]$.

Discussion on disconnected designs

Despite the connectivity constraint, a minority of the designs shown in this manuscript exhibit disconnected regions (see e.g., Figure 6). More specifically, some beams contain greyscale density values because the optimizer relies on the compliance of this beam to ensure it oscillates at the right frequency. This is thus a way for the optimizer to generate sub-lengthscale beams. In order to prevent this artifact type, additional greyness penalization must be added. As perfect connectivity is not the focus of this study, this is not further explored.

References

[1] M. P. Bendsøe, N. Kikuchi, Generating optimal topologies in structural design using a homogenization method, Computer Methods in Applied Mechanics and Engineering 71 (2) (1988) 197–224. doi:https://doi.org/10.1016/0045-7825(88)90086-2.
URL https://www.sciencedirect.com/science/article/pii/0045782588900862

[2] A. R. Díaaz, N. Kikuchi, Solutions to shape and topology eigenvalue optimization problems using a homogenization method, International Journal for Numerical Methods in Engineering 35 (7) (1992) 1487–1502. arXiv:https://onlinelibrary.wiley.com/doi/pdf/10.1002/nme.1620350707, doi:https://doi.org/10.1002/nme.1620350707.
URL https://onlinelibrary.wiley.com/doi/abs/10.1002/nme.1620350707

[3] A. Lewis, M. Overton, Eigenvalue optimization, Acta Numerica 5 (1996) 149–190. doi:10.1017/S0962492900002646.

[4] H. S. Koh, G. H. Yoon, Large-scale topology optimization for dynamic problems using a repetitive substructuring approach, Structural and Multidisciplinary Optimization 67 (33) (2024). doi:https://doi.org/10.1007/s00158-024-03759-1.

[5] M. Xiao, J. Ma, X. Gao, P. Breitkopf, B. Raghavan, W. Zhang, L. Cauvin, P. Villon, Primal–dual on-the-fly reduced-order modeling for large-scale transient dynamic topology optimization, Computer Methods in Applied Mechanics and Engineering 428 (2024) 117099. doi:https://doi.org/10.1016/j.cma.2024.117099.
URL https://www.sciencedirect.com/science/article/pii/S0045782524003554

[6] Q. Li, O. Sigmund, J. S. Jensen, N. Aage, Reduced-order methods for dynamic problems in topology optimization: A comparative study, Computer Methods in Applied Mechanics and Engineering 387 (2021) 114149. doi:https://doi.org/10.1016/j.cma.2021.114149.
URL https://www.sciencedirect.com/science/article/pii/S0045782521004801

[7] V. Cool, O. Sigmund, N. Aage, F. Naets, E. Deckers, Vibroacoustic topology optimization for sound transmission minimization through sandwich structures, J. of Sound and Vibr. 568 (2024) 117959, https://doi.org/10.1016/j.jsv.2023.117959.

[8] D. Giraldo Guzman, C. Lissenden, P. Shokouhi, M. Frecker, Topology optimization design of resonant structures based on antiresonance eigenfrequency matching informed by harmonic analysis, Journal of Mechanical Design 145 (10) (2023) 101704. doi:10.1115/1.4062882.
URL https://doi.org/10.1115/1.4062882

[9] D. Giannini, N. Aage, F. Braghin, Topology optimization of mems resonators with target eigenfrequencies and modes, European Journal of Mechanics - A/Solids 91 (2022) 104352. doi:https://doi.org/10.1016/j.euromechsol.2021.104352.
URL https://www.sciencedirect.com/science/article/pii/S0997753821001182

[10] O. M. Silva, M. M. Neves, A. Lenzi, A critical analysis of using the dynamic compliance as objective function in topology optimization of one-material structures considering steady-state forced vibration problems, Journal of Sound and Vibration 444 (2019) 1–20. doi:https://doi.org/10.1016/j.jsv.2018.12.030.
URL https://www.sciencedirect.com/science/article/pii/S0022460X18308563

[11] E. Andreassen, F. Ferrari, O. Sigmund, A. R. Diaz, Frequency response as a surrogate eigenvalue problem in topology optimization, International Journal for Numerical Methods in Engineering 113 (8) (2018) 1214–1229. doi:https://doi.org/10.1002/nme.5563.

[12] L. Herrmann, O. Sigmund, V. M. Li, C. Vogl, S. Kollmannsberger, On neural networks for generating better local optima in topology optimization, Structural and Multidisciplinary Optimization 67 (11) (NOV 2024). doi:10.1007/s00158-024-03908-6.

[13] H. Assimi, A. Jamali, A hybrid algorithm coupling genetic programming and nelder-mead for topology and size optimization of trusses with static and dynamic constraints, Expert Systems with Applications 95 (2018) 127–141. doi:10.1016/j.eswa.2017.11.035.

[14] D. Giannini, M. Schevenels, E. P. Reynders, Rotational and multimodal local resonators for broadband sound insulation of orthotropic metamaterial plates, Journal of Sound and Vibration 547 (2023) 117453. doi:https://doi.org/10.1016/j.jsv.2022.117453.
URL https://www.sciencedirect.com/science/article/pii/S0022460X22006368

[15] S. Zhang, J. A. Norato, Finding better local optima in topology optimization via tunneling, in: Proceedings of the ASME International Design Engineering Technical Conferences and Computers and Information in Engineering Conference, ASME, Design Engn Div; ASME, Comp & Informat Engn Div, 2018, aSME International Design Engineering Technical Conferences (IDETC) / Computers and Information in Engineering Conference (CIE), Quebec City, CANADA, AUG 26-29, 2018.

[16] P. Liu, X. Zhang, Y. Luo, Topological design of freely vibrating bi-material structures to achieve the maximum band gap centering at a specified frequency, Journal of Applied Mechanics-Transaction of the ASME 88 (8) (AUG 1 2021). doi:10.1115/1.4050621.

[17] A. Dalklint, R. E. Christiansen, O. Sigmund, On the potential use of performance bounds based on semidefinite programs for topology optimization (2025).

[18] V. Cool, O. Sigmund, N. Aage, A practical review on promoting connectivity in topology optimization, Struct. and Multidiscipl. Opt. 68 (73), https://doi.org/10.1007/s00158-025-04004-z (2025).

[19] C. B. Dilgen, N. Aage, Topology optimization of transient vibroacoustic problems for broadband filter design using cut elements, Finite Elements in Analysis and Design 234 (2024) 104123. doi:https://doi.org/10.1016/j.finel.2024.104123.
URL https://www.sciencedirect.com/science/article/pii/S0168874X24000179

[20] S. Halkjær, O. Sigmund, J. S. Jensen, Maximizing band gaps in plate structures, Structural and Multidisciplinary Optimization 32 (2006) 263–275. doi:https://doi.org/10.1007/s00158-006-0037-7.
[21] R. Kohn, G. Allaire, Topology optimization and optimal shape design using homogenization, Kluwer, 1993, pp. 207–218.
[22] C. Sanders, J. Norato, T. Walsh, W. Aquino, An error-in-constitutive equations strategy for topology optimization for frequency-domain dynamics, Computer Methods in Applied Mechanics and Engineering 372 (DEC 1 2020). doi:10.1016/j.cma.2020.113330.
[23] X. Guo, G. Cheng, K. Yamazaki, A new approach for the solution of singular optima in truss topology optimization with stress and local buckling constraints, Structural and Multidisciplinary Optimization 22 (5) (2001) 364–372. doi:10.1007/s00158-001-0156-0.
[24] N. Olhoff, J. Du, Generalized incremental frequency method for topological designof continuum structures for minimum dynamic compliance subject to forced vibration at a prescribed low or high value of the excitation frequency, Structural and Multidisciplinary Optimization 54 (5, SI) (2016) 1113–1141. doi:10.1007/s00158-016-1574-3.
[25] F. Meng, L. Meng, J. Wang, J. Zhu, B. P. Wang, S. Yuan, W. Zhang, Topology optimization for minimum dynamic compliance using an antiresonant frequency constraint, Structural and Multidisciplinary Optimization 67 (9) (SEP 2024). doi:10.1007/s00158-024-03878-9.
[26] R. F. Boukadia, E. Deckers, C. Claeys, M. Ichchou, W. Desmet, A wave-based optimization framework for 1D and 2D periodic structures, Mech. Syst. and Signal Process. 139 (2020) 106603, https://doi.org/10.1016/j.ymssp.2019.106603.
[27] F. Fahy, P. Gardonio, Sound and structural vibration: radiation, transmission and response, 2nd edition, Elsevier, UK, 2007.
[28] F. Bloch, Über die quantenmechanik der elektronen in kristallgittern, Z. für Phys. 52 (7) (1929) 555–600, https://link.springer.com/article/10.1007/BF01339455.
[29] V. Cool, C. Claeys, H. Denayer, F. Naets, E. Deckers, Experimental validation of a topology optimized 2d sandwich panel for sound transmission minimization, Mechanical Systems and Signal Processing 211 (2024) 111197, https://doi.org/10.1016/j.ymssp.2024.111197.
[30] N. de Melo Filho, L. Van Belle, C. Claeys, E. Deckers, W. Desmet, Dynamic mass based sound transmission loss prediction of vibro-acoustic metamaterial double panels applied to the mass-air-mass resonance, J. Sound Vibr. 442 (2019) 28–44, https://doi.org/10.1016/j.jsv.2018.10.047.
[31] J. Jensen, A simple method for coupled acoustic-mechanical analysis with application to gradient-based topology optimization, Struct. and Multidiscipl. Opt. 59 (5) (2019) 1567–1580, https://doi.org/10.1007/s00158-018-2147-4.
[32] F. Wang, B. S. Lazarov, O. Sigmund, On projection methods, convergence and robust formulations in topology optimization, Struct. and Multidiscipl. Opt. 43 (6) (2011) 767–784, https://doi.org/10.1007/s00158-010-0602-y.
[33] R. E. Christiansen, B. S. Lazarov, J. S. Jensen, O. Sigmund, Creating geometrically robust designs for highly sensitive problems using topology optimization, Struct. and Multidiscipl. Opt. 52 (4) (2015) 737–754, https://doi.org.10.1007/s00158-015-1265-5.
[34] K. Svanberg, The method of moving asymptotes—a new method for structural optimization, Int. J. for Num. Methods in Eng. 24 (2) (1987) 359–373, https://doi.org/10.1002/nme.1620240207.
[35] H. Hoffmann, Violin plot, https://www.mathworks.com/matlabcentral/fileexchange/45134-violin-plot, accessed: 2025-03-24.
[36] B. Bourdin, Filters in topology optimization, Int. J. Num. methods Eng. 50 (9) (2001) 2143–2158, https://doi.org/10.1002/nme.116.
[37] J. K. Guest, J. H. Prévost, T. Belytschko, Achieving minimum length scale in topology optimization using nodal design variables and projection functions, Int. J. Num. Methods Eng. 61 (2) (2004) 238–254, https://doi.org/10.1002/nme.1064.
[38] M. Stolpe, K. Svanberg, An alternative interpolation scheme for minimum compliance topology optimization, Struct. and Multidiscipl. Opt. 22 (2) (2001) 116–124, https://doi.org/10.1007/s001580100129.
[39] M. Hussein, M. Leamy, M. Ruzzene, Dynamics of phononic materials and structures: Historical origins, recent progress, and future outlook, Appl. Mech. Rev. 66 (4) (2014) 040802, https://doi.org/10.1115/1.4026911.
[40] Y. Yang, B. R. Mace, M. J. Kingan, Prediction of sound transmission through, and radiation from, panels using a wave and finite element method, J. Acoust. Soc. Am. 141 (4) (2017) 2452–2460, https://doi.org/10.1121/1.4977925.
[41] Y. Yang, B. Mace, M. Kingan, Vibroacoustic analysis of periodic structures using a wave and finite element method, J. of Sound and Vibr. 457 (2019) 333–353, https://doi.org/10.1016/j.jsv.2019.06.009.
[42] K. E. Swartz, D. A. Tortorelli, D. A. White, K. A. James, Manufacturing and stiffness constraints for topology optimized periodic structures, Struct. and Multidiscipl. Opt. 65 (129), https://doi.org/10.1007/s00158-022-03222-z (2022).