ENHANCEMENT OF COMPOSITE FIREWALL AIRBORNE INSULATION THROUGH LOCALLY RESONANT LOW FREQUENCY TUNED VIBRATION ABSORBERS

Youssef Gerges[1], Benjamin Morin[1], Marco Cardillo[1], Francisco Climent[1], Vanessa Cool[2,3], Elke Deckers[2,3], Claus Claeys[2,3]

[1]*Autoneum Management AG, Winterthur, Switzerland*
[2]*Department of Mechanical Engineering, KU Leuven, Belgium*
[3]*Flanders Make @ KU Leuven, Belgium*

Abstract

Improving low-frequency airborne insulation in automotive firewall panels is challenging because structure-borne transmission dominates below 500Hz, where traditional acoustic treatments are less effective. Measurements performed on a dash panel from a battery-electric-vehicle (BEV) revealed a pronounced transmission-loss (TL) dip near 500Hz, indicating a structural weakness linked to the dash panel dynamics. This work investigates tuned vibration absorbers (TVAs) as lightweight, locally resonant solutions to mitigate such narrowband deficiency. TVAs introduce sub-wavelength resonances that increase mechanical impedance and locally suppress bending-wave mobility, thereby enhancing low-frequency TL. A validated structural finite element (FE) model forms the basis for assessing TVA performance. Airborne TL predictions were further validated using a coupled vibro-acoustic FE model and measurements in coupled rooms. A first TVA layout, constrained by realistic mass and packaging limitations, was evaluated and yielded 3–6dB TL improvement at the 500Hz dip frequency without negative effects on adjacent bands. Finally, a Genetic Algorithm optimization step was performed to explore multi-frequency TVA arrangements under design constraints, demonstrating additional potential for targeted TL enhancements. These results confirm the feasibility of TVA-based countermeasures for automotive firewall panels and motivate further optimization-driven layout refinements.

Corresponding author: *youssef.gerges@autoneum.com.*

Keywords: *tuned vibration absorbers, airborne insulation, composite dash firewall panel, genetic algorithm optimization, battery electric vehicle*

1. Introduction

Improving low-frequency airborne insulation in automotive firewalls is challenging because structure-borne transmission dominates up to 500Hz, where conventional dash treatments offer limited benefit. As an illustrative example, the TL measured on a Glass-Fiber-Reinforced composite firewall panel from a battery-electric vehicle (BEV) showed a pronounced dip near 500Hz in both bare and trimmed configurations, confirming that the weakness originates from the panel's structural dynamics rather than the trim. This work investigates tuned vibration absorbers (TVAs) as lightweight, locally resonant solutions to such narrow-band TL deficiencies. By introducing sub-wavelength resonances, TVAs produce a sharp mechanical-impedance increase that reduces bending-wave mobility and enhances low-frequency TL [1]. Arranged periodically, they can in theory form locally resonant metamaterials with stop bands, but such ideal layouts are infeasible on firewalls due to geometric, functional, and mass constraints. This study therefore focuses on practical TVA placement strategies, supported by validated structural and acoustic FE models, and applies a genetic algorithm to explore multi-frequency TL improvements under realistic design limits. As the work is fully simulation-based, it represents a first assessment of TVA performance; a second phase will integrate TVAs into the dash panel for real-case measurements to validate the numerical predictions.

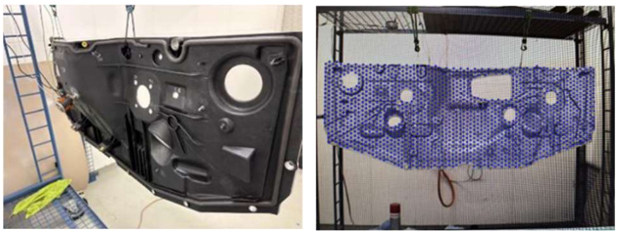

Figure 1. Structural vibration measurement setup. Left: Suspended dash panel. Right: Polytec 2D scanning points.

2. Structural FE Model Validation

A reliable FE model of the firewall panel is required before evaluating TVA effects. Experimental structural measurements and simulations were compared to validate the mobility response.

2.1 Experimental Setup and structural model

The panel was suspended in near free-free conditions and excited using a shaker sweep from 50-1000Hz. Velocities were measured using a Polytec 2D scanning laser vibrometer over 1728 points (Figure 1). In parallel, a FE model of the firewall panel was built (visible in Figure 4 below). For the GFRP material constituting the panel, the following parameters were used: E=10GPa, density=1585kg/m^3, Poisson's ratio=0.3, damping loss factor = 0.03%.

2.2 Correlation Results

Space-average mobility FRFs match well up to 1000Hz (Figure 2), indicating that the model captures the main dynamics and is fit for purpose for TVA prediction.

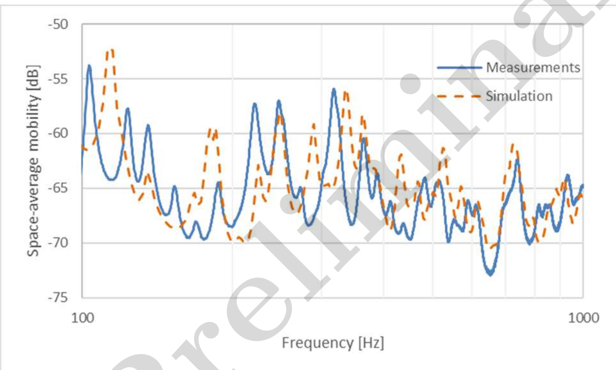

Figure 2. Space-average mobility (velocity/force): measurement (solid line) vs simulation (dashed line).

3. TL Simulation and Validation

The TL of the firewall was evaluated experimentally and numerically to provide a validated acoustic reference for TVA assessment.

3.1 Experimental Setup

Airborne TL was measured using a reverberant emission room coupled to a semi-anechoic receiving room. The dash was installed in a wooden mounting frame between rooms (Figure 3). 4 microphones were used to measure the spatial average of the RMS SPL in the source room, while the Microflown Scan & Paint (S&P) system [5] was used to measure transmitted sound intensity on the receiving side. From these quantities, TL was evaluated following the standard procedure of ISO 15186 [6].

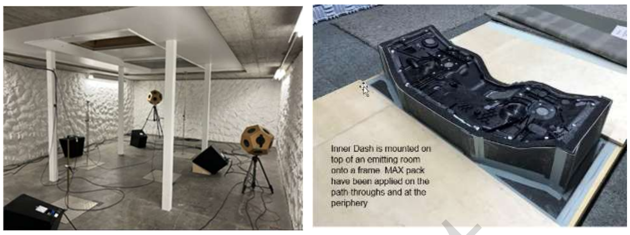

Figure 3. TL measurement setup. Left: Reverberant emission room. Right: dash panel facing the semi-anechoic receiving room.

3.2 TL FE-model Validation

The vibro-acoustic FE model (Figure 4) includes:

- the validated dash structural model,
- a front cavity approximating the diffuse field,
- a back cavity for the receiving room.

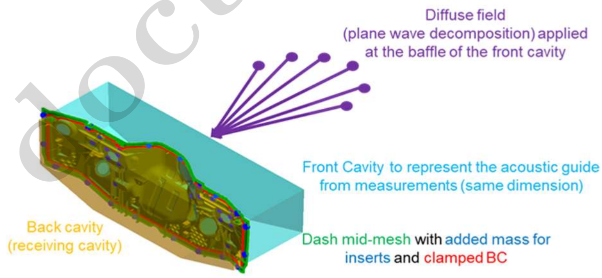

Figure 4. Transmission Loss FE simulation setup

The front cavity in Figure 4 represents the air volume created by the wooden mounting frame in the experimental setup and must be modeled explicitly. A diffuse-field boundary condition is applied on its outer surface via the plane-wave decomposition technique, i.e., the summation of numerous uncorrelated plane waves impinging from random directions [2]. The back cavity is coupled to a Perfectly Matched Layer (PML) domain, which models the receiving-side anechoic room and absorbs outgoing energy without reflection.

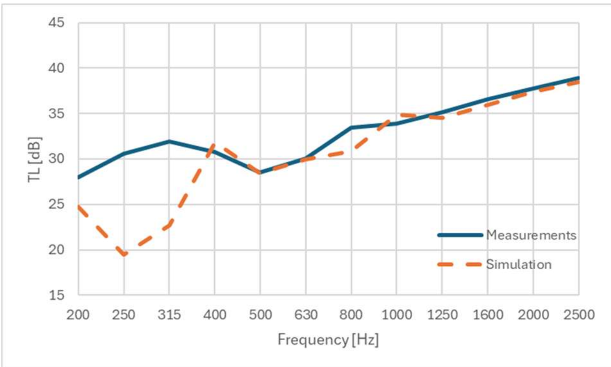

Figure 5. TL comparison in one third octave band: measurement (solid line) vs FE simulation (dashed line).

Simulated TL agrees well with measurements above 300 Hz (Figure 5); remaining low-frequency deviations stem from mounting constraints and imperfect diffusivity in the physical setup. The measured 500 Hz TL dip, linked to the panel's curvature-driven stiffness, validates 500 Hz as the primary target frequency for TVA tuning.

4. TVA Principles and Design

4.1 Principles and Unit-Cell Behavior

A TVA is a structural resonator tuned to a target frequency. Near resonance, it exhibits a strong dynamic mass effect that produces a large impedance peak on the host structure and locally suppresses vibration. When multiple resonators are distributed on a plate, their interaction with the host structure can generate a stop band in the bending-wave dispersion curve.
The Modal Effective Mass (MEM) is a key indicator of stop-band width: higher MEM yields stronger inertial coupling and broader attenuation, essential for lightweight metamaterial-type layouts [3].

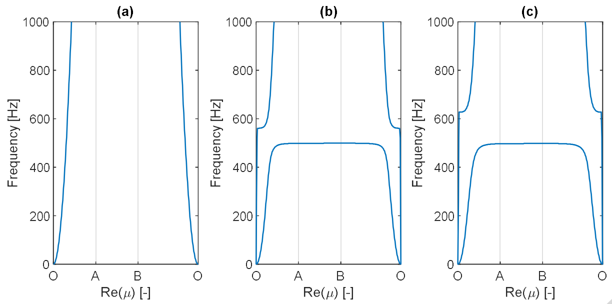

Figure 6. Dispersion curve of a host plate with and without TVAs (tuned at 500 Hz) for different MEM: (a) no TVA, (b) 40% MEM with 60 Hz stop band, (c) 90% MEM with 130 Hz stop band.

Figure 6 illustrates the MEM impact on the dispersion curve with a flat 30×30 mm equivalent unit cell; computational guidelines for dispersion curve plotting are provided in [4]. Case (a) shows the host structure without TVA, where bending waves propagate over the entire frequency band. Cases (b) and (c) use 3.7 g resonators tuned at 500 Hz: with 40% MEM (b) the stop band spans 500–560 Hz, and with 90% MEM (c) it spans 500–630 Hz. The resonators are modelled as 1D discrete spring-mass systems, where the mass at the resonator equals the MEM fraction of its total weight and the remaining mass is lumped at the attachment point.

4.2 Design Constraints and Frequency Tuning

The bare dash weighs 5.5 kg; to limit added mass to ~10%, the number of resonators and the lattice pitch must be constrained. A pitch of 60 mm was selected as a compromise between lightweighting and stop-band robustness. Multiple frequencies can be targeted using resonators tuned to specific bands, achieved in practice by adjusting the mass coupled to each resonator. Heavier units are used only in selected areas to avoid a large increase in added mass.

This tradeoff motivates the optimization framework developed later. The concept is to deploy resonator patches in localized areas of the dash that can realistically accommodate TVAs, with each patch assigned a tuning frequency determined through optimization to target specific bands. The patch pitch is chosen so that, at least within the patch, resonator spacing remains subwavelength with respect to the bending wave targeted at that frequency, assuming an equivalent flat host material.

5. TVA Integration and Assessment

The first TVA layout comprised 70 resonators arranged with an average spacing of about 60 mm, applied primarily to flat regions of the dash surface (Figure 7). The TVAs were modeled as 1D discrete spring-mass systems tuned at 500 Hz, acting normal to the panel, with a 3% damping ratio. They were distributed substantially over the entire firewall surface wherever possible and all tuned to the same frequency. Each resonator weighs 3.5 g, for a total added mass of 245 g, only about 4.5%. This first assessment aimed to evaluate the improvement potentially obtainable via TVAs at the 500 Hz target frequency.

Figure 7. Resonators distribution (red dots) on the firewall panel with corresponding repartition into patches (yellow rectangles).

Two cases were evaluated, with MEM of 40% and 90%, yielding +3dB and +6dB TL improvements, respectively, at 500 Hz, with no adverse side effects in adjacent bands (Figure 8).

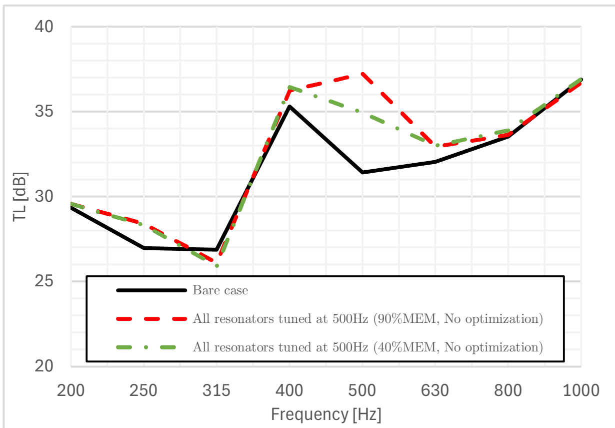

Figure 8. Transmission Loss comparison between bare panel (solid line), TVA-treated panel with 90% MEM (dashed line) and TVA-treated panel with 40% MEM (dashed dot line) using FE simulation.

These results confirm the effectiveness of the TVAs. At the same time, they also leave doubt about whether the layout considered may not be overengineered, i.e., whether similar results could be obtained by tuning a smaller number of resonators to the target frequency of 500Hz, and making use of the remaining resonators to improve TL at other frequencies. This motivates the analysis of multi-frequency layouts guided by optimization under realistic constraints, as described in the next section.

6. Genetic Algorithm Methodology

The optimization framework employs a Genetic Algorithm (GA) designed for discrete multi-valued design spaces. GAs are stochastic optimization methods inspired by natural evolution, first introduced by Holland [7]. In the present implementation, each chromosome represents a candidate design whose genes correspond to discrete design variables. The dash panel carries 70 resonators divided into 10 patches (Figure 7), and each patch can be assigned to one of three configurations: no resonators (0), resonators tuned to 250 Hz (1), or resonators tuned to 500 Hz (2). All TVAs assume a MEM of 90%, as investigated in the previous section. The algorithm minimizes an objective function based on the RMS velocity response of the panel across third-octave bands. This vibration-based approach was chosen for computational efficiency, cutting evaluation time from 12 hours (transmission loss simulation) to 30 minutes per design; its validity relies on the strong correlation between panel velocity and transmission loss. For each third-octave band i, a normalized velocity reduction is computed as:
$\Delta_i = (V_{bare,i} – V_{opt,i})/V_{bare,i}$, where $V_{bare,i}$ and $V_{opt,i}$ are the RMS velocities of the bare and optimized panels. This normalization makes results comparable across bands: $\Delta_i > 0$ indicates improvement (velocity reduction), $\Delta_i < 0$ degradation. The fitness function is then defined as: $F = (1/b) \times \Sigma\, \Delta_i \times P$, where b is the number of third-octave bands studied and P is a penalty factor discouraging inconsistent designs: $P = 1 - 0.25 \times N$, with N the number of bands where the optimized design underperforms the bare case ($\Delta_i < 0$). This ensures that designs with more than four underperforming bands yield negative fitness, favoring configurations that improve performance across the whole frequency range rather than excelling at some frequencies while degrading others.
The algorithm implements a generational evolutionary strategy combining several complementary techniques to balance exploration and exploitation while preventing premature convergence [8].

6.1 Population Initialization - Covering Arrays

The quality of the initial population significantly affects GA performance. Random initialization offers no guarantee of systematic exploration and may miss important regions of the design space in early generations. To address this, the algorithm uses pairwise covering array seeding [9], which ensures that every combination of values for any pair of design variables appears at least once in the initial population. For n design variables with v possible values each, this covers all $n(n\text{-}1)v^2/2$ pairwise combinations, 405 in the present application (n=10, v=3). The array is built with a greedy algorithm: random candidates are generated in batches, and the one covering the most uncovered pairs is iteratively selected until full coverage is reached. Combinatorial-testing studies have shown that pairwise coverage effectively detects interaction effects, and this principle transfers well to evolutionary optimization by ensuring diverse initial exploration of variable interactions.

6.2 Selection and Elitism

Parent selection uses tournament selection, which balances selection pressure with computational efficiency. In each tournament, k individuals are randomly sampled and the one with highest fitness is chosen as a parent. The tournament size k controls selection pressure: smaller values (k=2) maintain genetic diversity, while larger values increase convergence speed. The algorithm also implements elitism [11], preserving a fraction of the highest-fitness individuals unchanged into the next generation (here 10% of the population) ensuring that the best discovered solutions are never lost to stochastic effects.

6.3 Uniform Crossover

The recombination operator significantly influences the algorithm's exploration behavior. This implementation employs uniform crossover, where each gene is independently inherited from either parent with 50% probability. This maximizes recombination potential, allowing any combination of parental alleles to appear in an offspring. For the optimization problems considered in this work, where design variables represent independent engineering parameters with no inherent ordering, uniform crossover provides good exploration characteristics.

6.4 Temperature-Based Adaptive Mutation

Mutation introduces new genetic material and prevents convergence to local optima. The algorithm uses a temperature-controlled adaptive mutation strategy inspired by simulated annealing [10]. A temperature parameter T, with $T_{max} = 1$, controls the mutation rate p_m through $p_m = p_{m,min} + (p_{m,max} - p_{m,min})*T$, with typical values $p_{m,min} = 0.02$ and $p_{m,max} = 0.30$. At each generation, temperature decreases by a cooling factor, $T_{t+1} = \alpha T_t$ with α typically around 0.95, naturally shifting from exploration to exploitation. The key innovation is adaptive reheating: when fitness stagnates for several consecutive generations or population diversity falls below a threshold (typically 40% unique individuals), temperature rises to escape local optima. This dual-trigger mechanism, responding to both stagnation and diversity loss, provides robust adaptation without manual tuning. When mutation occurs, the new value is drawn uniformly from the three alternatives, ensuring that computational effort produces actual genetic variation.

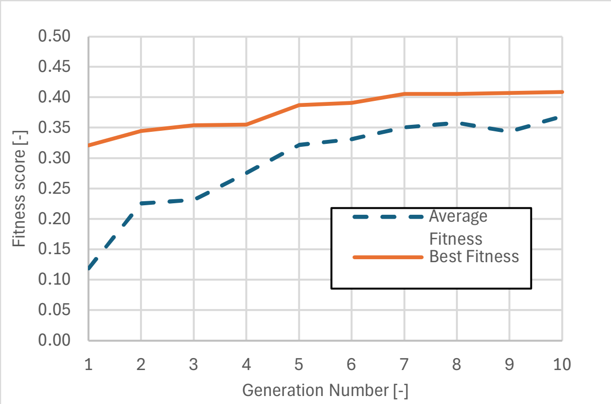

Figure 9. Evolution of the average and best fitness

7. Results and Discussion

7.1 Optimization convergence

The genetic algorithm ran for 10 generations with a population of 30, giving 300 total fitness evaluations. Of these, 158 unique configurations were tested (52.7% exploration efficiency); the rest were duplicate designs that naturally emerge as the population converges toward high-fitness regions. Figure 9 shows the evolution of best and average fitness, which progressed through three distinct phases. In the initial exploration phase (generations 1–4), diversity stayed high (87–100%) as the pairwise covering array initialization systematically explored the design space, and best fitness improved from 0.321 to 0.355 (10.6%). The second phase (generations 5–7) transitioned toward exploitation, with diversity falling to 47% as the population converged; this phase yielded the largest gain, the best solution rising from 0.355 to 0.406 (14.4%). At generation 8, the temperature controller detected stagnation and triggered a reheating event, temporarily raising the mutation rate to renew exploration. The final phase (generations 8–10) balanced continued exploitation with this reheating-induced exploration, discovering a marginally improved solution with fitness 0.409.

Table 1 summarizes the generation-by-generation statistics. The algorithm identified nine successive improvements to the best fitness, demonstrating consistent progress without premature convergence, and the temperature-controlled mutation mechanism maintained sufficient diversity to keep discovering improved solutions in later generations.

7.2 Optimized Design Configuration

The optimization converged to a consistent design pattern across multiple high-performing individuals. The final generation showed strong consensus on most design variables, with eight of ten showing greater than 90% agreement. The optimal configuration is [2, 1, 2, 1, 2, 2, 1, 2, 2, 2] (Figure 10), meaning patches [1, 3, 5, 6, 8, 9, 10] are tuned at 500Hz and patches [2, 4, 7] at 250Hz, achieving a fitness of 0.409. Here each 500 Hz resonator weighs 3.5g and each 250Hz resonator 14g, giving a total added mass of 476g. The converged design represents a 27.4% improvement in fitness over the initial best individual while increasing the dash weight by less than 10% relative to the bare panel. This substantial improvement validates the genetic algorithm's effectiveness in navigating the complex, multi-dimensional design space.

Figure 10. Patches tuning after optimization.

7.3 Computational Efficiency

The GA optimization required 300 fitness evaluations distributed across 10 generations. Each fitness evaluation involved a finite element simulation of the panel's vibration response, with an average computation time of approximately 30 minutes per evaluation. The total optimization wall-clock time was approximately 150 hours, demonstrating that the genetic algorithm approach is computationally tractable for practical engineering design problems.

The pairwise covering array initialization contributed to the algorithm's efficiency by ensuring systematic exploration of variable interactions from the outset. Combined with the adaptive mutation mechanism, this allowed the algorithm to identify high-quality solutions within a relatively modest number of evaluations compared to exhaustive search, which would require 3^{10} = 59'049 evaluations to test all possible configurations.

7.4 Transmission Loss Improvement

The GA's fitness function used to identify the optimal TVA arrangement was based on the resonators' effect on the RMS vibration of the firewall panel. The primary objective, however, was to enhance the dash panel TL at two critical frequencies: 250Hz and 500Hz.

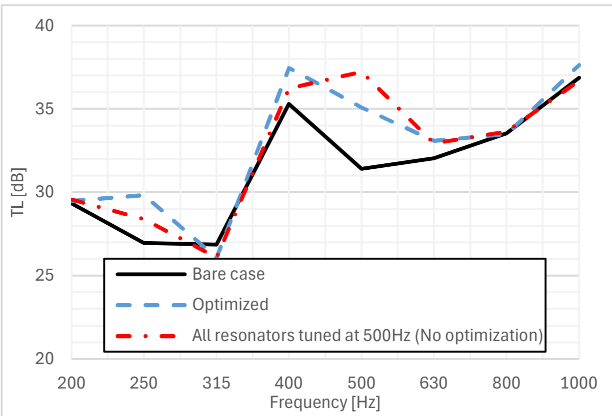

Figure 11. TL comparison between bare panel (solid line), panel treated with all TVAs tuned at 500 Hz (dashed dot

line), and optimized configuration (dashed line) across the 200–1000Hz frequency range in the case of 90% MEM.

Table 1. Optimization progress per generation.

Generation	Temperature	Best Fitness	Average Fitness	Diversity (%)
1	0.80	0.321	0.119	100
2	0.76	0.344	0.226	93
3	0.72	0.354	0.232	90
4	0.69	0.355	0.276	87
5	0.65	0.387	0.322	70
6	0.62	0.391	0.331	63
7	0.59	0.406	0.351	47
8	0.76	0.406	0.358	43
9	0.72	0.407	0.343	50
10	0.69	0.409	0.369	43

Figure 11 compares the TL spectra of the bare panel (no resonators), the optimized configuration, and the non-optimized configuration (all resonators tuned at 500Hz) across 200–1000Hz. At 250Hz, the optimized design achieves a TL of about 29.5dB versus 26.5dB for the bare case, an improvement of 3.0dB. At 500Hz, TL increases from about 31.5dB to 35.0dB, a 3.5dB improvement. These gains are significant: matching them by adding non-structural mass would require increasing the panel mass by about 2.3kg (roughly 40% of its total mass). The optimized configuration represents a balanced trade-off compared to a dash fully covered with resonators tuned to 500Hz, which achieves a higher improvement of about 5dB at that frequency but only 1.5dB at 250Hz. By distributing resonators across both tuning frequencies, the optimized design delivers more uniform improvement at both targets.
The TL curves show that the optimization targeted the specified frequencies without significantly degrading performance elsewhere: a very mild degradation (<1 dB) appears only at 315Hz, while at 400Hz the optimized design improves by about 2dB, and performance above 630Hz remains unchanged. Overall acoustic performance is clearly superior for the intended application, where the 250–500Hz range is of primary concern.

Conclusions

This paper presented a custom genetic algorithm developed for the vibroacoustic optimization of a vehicle dash panel with discrete-valued design variables. The first part of the work validated the FE model of the bare dash panel, essential for any reliable assessment of TVAs. The GA integrates pairwise covering array initialization for systematic exploration, uniform crossover, tournament selection with elitism, and a temperature-controlled adaptive mutation mechanism inspired by simulated annealing. This scheme balances exploration and exploitation by gradually reducing mutation rates while triggering reheating when stagnation or diversity loss is detected, preventing premature convergence. The optimization targeted transmission loss at the dash panel resonance frequencies of 250Hz and 500Hz. Over 10 generations with 300 fitness evaluations, the algorithm identified an optimal configuration achieving improvements of 3.0dB and 3.5dB, respectively. Robust convergence, with nine successive fitness improvements and strong consensus on the final design, validates the approach for discrete vibroacoustic optimization problems. As the present work is entirely simulation-based, it constitutes a first assessment of the TVA concept on automotive firewalls. Future work will validate both the TVA effect and the optimized layouts through full-scale measurements, confirming the predicted improvements

Acknowledgements

The authors would like to thank Federico Di Marco, Claudio Bertolini and Jan Horak from Autoneum Management AG for their valuable support and insightful contributions throughout this work.

References

[1] Claeys, Claus, et al. "On the potential of tuned resonators to obtain low-frequency vibrational stop bands in periodic panels," *Journal of Sound and Vibration*, vol. 332, no. 6, pp. 1418-1436, 2013.

[2] Free Field Technologies (FFT), "Actran 2024.2 - User's Guide Volume 2", Hexagon.

[3] Cool, Vanessa, et al. "Topology optimization of locally resonant metamaterials for high modal effective mass," *Structures*, vol. 85, pp. 111192, 2026.

[4] Cool, Vanessa, et al. "A guide to numerical dispersion curve calculations: Explanation, interpretation and basic Matlab code," *Mechanical Systems and Signal Processing*, vol. 215, pp. 111393, 2024.

[5] Comesaña, Daniel Fernández, et al. "Scan and paint: theory and practice of a sound field visualization method," *International Scholarly Research Notices,* vol. 2013, no. 1, pp. 241958, 2013.

[6] ISO 15186-1:2000, "Acoustics – Measurements of sound insulation in buildings and of building elements using sound intensity", 2000.

[7] J.H. Holland, "Adaptation in Natural and Artificial Systems". Ann Arbor, MI: University of Michigan Press, 1975.

[8] A.E. Eiben and J.E. Smith, "Introduction to Evolutionary Computation", 2nd ed. Berlin: Springer, 2015.

[9] C. Nie and H. Leung, "A survey of combinatorial testing," ACM Computing Surveys, vol. 43, no. 2, pp. 11:1–11:29, 2011.

[10] S. Kirkpatrick, C. D. Gelatt, and M. P. Vecchi, "Optimization by simulated annealing," Science, vol. 220, no. 4598, pp. 671–680, 1983.

[11] DeJong, K. "An Analysis of the Behavior of a Class of Genetic Adaptive Systems. Ph.D. Dissertation", Department of Computer and Communication Sciences, University of Michigan, Ann Arbor, 1975.