



A parallel computing framework for performing structural-acoustic optimization with stochastic forcing

Micah R. Shepherd¹ · Robert L. Campbell¹ · Stephen A. Hambric¹

Received: 12 March 2019 / Revised: 20 June 2019 / Accepted: 13 August 2019 / Published online: 31 August 2019
© Springer-Verlag GmbH Germany, part of Springer Nature 2019

Abstract

Structural-acoustic optimization procedures can be used to find the optimal design for reduced noise or vibration in many real-world scenarios. However, the time required to compute the structural-acoustic quantity of interest often limits the size of the model. Additionally, structural-acoustic optimization using state-of-the-art evolutionary algorithms may require tens of thousands of system solutions, which add to the limitations for large full-scale systems. To reduce the time required for each function evaluation, parallel processing techniques are used to solve the system in a highly scalable fashion. The approach reduces the analysis time by solving the system using a frequency-domain formulation and distributing solution frequencies amongst processors to solve in parallel. To demonstrate, the sound radiated from a curved panel under the influence of a turbulent boundary layer is minimized in the presence of added point masses, which are varied during the optimization procedure. The total mass is also minimized and the Pareto front relating the trade-off between added mass and reduced noise is determined. Solver scaling information is provided that demonstrates the utility of the parallel processing approach.

Keywords Structural acoustics · Evolutionary strategy · High-performance computing

1 Introduction

Structural-acoustic optimization (SAO) techniques can be applied to minimize or maximize the acoustic performance of a structure. Efficient SAO strategies may be derived by combining modern search algorithms with efficient structural-acoustic analysis methods. The optimization algorithm searches a specified design space to locate an optimum set of design variables. The accuracy of the results depends on the fidelity of the structural-acoustic analysis and the robustness of the search algorithm (Belegundu et al. 1994).

Historically, optimization algorithms used some form of a gradient or Hessian operation to determine the search direction. While gradient-based algorithms can be very efficient on linear, convex, and unimodal problems, they typically do not perform well on discontinuous, multimodal,

or noisy objective functions. Evolutionary algorithms (EAs) on the other hand often perform well on highly complex or discontinuous functions, but usually at the cost of a large number of function evaluations (Coello Coello et al. 2007).

Many structural-acoustic problems involve solving the response at an array of frequencies. For SAO and other frequency-domain problems, the response must be collapsed into a single scalar value to use as the objective function (Marburg et al. 2016). This is often done by integrating the response over the frequency band of interest (Marburg 2002). The basic variables (material, geometry, or boundary conditions) that control the structural-acoustic behavior do not, however, affect the response at all frequencies the same. This causes the objective function to have a complicated, nonlinear dependence on the design variables making the objective function to be equally nonlinear and complex (Butkewitsch and Steffen 2001). Therefore, a heuristics-based optimizer is often better suited to perform SAO due to its ability to navigate the highly nonlinear search space.

A number of informative articles on SAO can be found in the literature, which optimize noise and/or vibration of structures ranging from simple beams/plates (Marburg et al. 2006; Jeon and Okuma 2008; Aiello and

Responsible Editor: Emilio Carlos Nelli Silva

✉ Micah R. Shepherd
mrs30@psu.edu

¹ The Pennsylvania State University, State College, PA, USA

Auremma 2018) to curved shells/cylinders (Johnson and Cunefare 2002; Shepherd and Hambric 2014), damping treatments (Wong 2016; Shepherd et al. 2016) and various automotive/aerospace vehicle components (Joshi et al. 2012; Yuksel et al. 2012). In general, one of the challenges of SAO is the computational cost associated with the sound radiation analysis, as noted by Cunefare (1994) and others. Because of the computationally intensive analysis, the number of design variables has typically been low, with notable exceptions being from Walsh et al. (2018)

As previously mentioned, EAs often require many hundreds or thousands of design evaluations which can lead to prohibitively long run times. This issue is amplified for problems when multiple objectives must be considered simultaneously in order to establish the trade-off between competing objectives (Coello Coello et al. 2007). In contrast, modern computing power has enabled researchers to solve larger structural-acoustic problems which may involve large finite element (FE) and/or boundary element (BE) analysis with complicated, partially correlated forcing functions (Hambric et al. 2010; Bonness et al. 2017). As models increase in size, the analysis time also increases. This can lead to compromises in mesh quality or a number of analysis frequencies in order to achieve feasible run times when doing SAO (Joshi et al. 2015). Naturally, this can introduce significant bias into the analysis and may lead to incorrect optima.

To reduce the computation time required to solve large-scale optimization problems, high-performance computing has been utilized to parallelize the multiobjective evolutionary algorithms (MOEAs). For example, master-slave approaches have been implemented where different slave processors independently evaluate each design in the generation (Durillo et al. 2008; Hadka and Reed 2015). Other strategies for parallelization include multi-master approach, island model, and the hybrid parallelization architecture (Hadka and Reed 2015; Reed 2016). Graphical processing units (GPUs) have also been used for speedups with MOEAs (de Souza and Ramirez Pozo 2014).

Alternatively, a parallel solution approach can be taken which parallelizes the function used to calculate the objective (Aage et al. 2015; Walsh and Aquino 2017). This approach is most appropriate when the required analysis is computationally expensive such as in large-scale structural or structural-acoustic problems. Parallelized computing strategies have been developed for a structural design using topology optimization where the matrix operations are performed in parallel. The most notable case has been the topology optimization of a full-scale aircraft wing, which was only possible due to the efficient use of high-performance computing during the analysis stage (Aage et al. 2017). While parallel approaches have been successfully applied to structural optimization problems,

parallelization has not been widely applied to structural-acoustic optimization problems, primarily due to the challenge posed by the solution matrices being a function of frequency. Additionally, the use of stochastic forcing functions, where the applied forces are partially correlated over the structure, complicates the parallelization procedure and must be carefully considered in the solution framework to reduce the computation time.

This paper describes the parallel computing framework for structural-acoustic optimization of structures excited by stochastic forcing functions. The analysis is performed in modal space where the modes are augmented with residual vectors and frequency-domain interpolation is utilized to reduce computation time. The solution matrices are then distributed to slave nodes according to a pre-set vector of analysis frequencies. Additionally, information passing between processors is reduced in order to lessen the overall evaluation time. The procedure is demonstrated on a curved panel subjected to a turbulent boundary layer flow (Shepherd and Hambric 2014). The panel is subjected to heavy fluid loading (, in water) and the radiated sound power is minimized by finding the optimal distribution of point masses while trading off overall system mass. This problem has been previously considered as a simple form of material tailoring (Constans et al. 1998). The overall scaling of the parallel procedure is also demonstrated.

2 Structural-acoustic analysis

To compute the vibration response of a driven structure in physical coordinates, the damped, forced equation of motion must be solved:

$$\mathbf{M}\ddot{\mathbf{x}} + \mathbf{B}\dot{\mathbf{x}} + \mathbf{K}\mathbf{x} = \mathbf{F}e^{j\omega t}, \quad (1)$$

where \mathbf{M} , \mathbf{B} , and \mathbf{K} are the mass, viscous damping, and stiffness matrices, \mathbf{F} is the force vector, $\omega = 2\pi f$ is the angular frequency, f is frequency in Hz, and $j = \sqrt{-1}$. This equation is often solved using a discretized mesh of the structure of interest and the finite element (FE) method. The forcing function is approximated as time harmonic, allowing the velocity frequency response function (FRF) to be determined as

$$\mathbf{H}(\omega) = j\omega \left[-\omega^2 \mathbf{M} + j\omega \mathbf{B} + \mathbf{K} \right]^{-1}. \quad (2)$$

Fluid loading and complex impedance effects can be included by adding their respective matrices to the denominator of the right hand side of (2) and will be discussed subsequently.

Approximating the system as linear and transformation from physical to modal space allows significant computational savings because the number of required modal degrees of freedom is generally much less than the number

of physical degrees of freedom. Structural vibration modes are computed by solving the eigenvalue problem

$$\mathbf{K}\phi_n = \omega_n^2 \mathbf{M}\phi_n, \tag{3}$$

where the eigenvalues ω_n^2 are related to the natural frequencies of the structure (ω_n) and ϕ_n are the associated eigenvectors (i.e., unscaled mode shapes). The modal FRF equation can be obtained once the modes shapes are known:

$$\mathbf{h}(\omega) = j\omega \left[-\omega^2 \mathbf{m} + j\omega \mathbf{b} + \mathbf{k} \right]^{-1}. \tag{4}$$

In (4), $\mathbf{m} = \phi^T \mathbf{M} \phi$ is the generalized or modal mass, $\mathbf{b} = \phi^T \mathbf{B} \phi$ is the generalized or modal damping, and $\mathbf{k} = \phi^T \mathbf{K} \phi$ is the generalized or modal stiffness. When the mode shapes are mass-normalized, the modal mass matrix \mathbf{m} is the identity matrix and the modal stiffness matrix \mathbf{k} is a matrix containing the eigenvalues, ω_n^2 , on the diagonals and zero elsewhere.

Fluid loading effects can be determined by estimating frequency dependent acoustic impedance matrices, $Z(\omega) = R(\omega) + jX(\omega)$, using the lumped parameter boundary element method by Koopmann and Fahline (1997), where R is the resistive component of the impedance and X is the reactive component. The fluid impedance matrices describe how oscillations at each boundary element create acoustic pressures on the boundary surface. The acoustic impedance, which completely characterizes the additional mass and damping imparted by the fluid to a structure, can be transformed to modal space and incorporated directly into (4) (Fahline and Koopmann 1996). In contrast with the sparse matrices found in FE analysis, these matrices are dense (i.e., fully populated):

$$\mathbf{h}(\omega) = j\omega \left[-\omega^2 \mathbf{m} + j\omega \mathbf{b} + \mathbf{k} - j\omega \mathbf{r}(\omega) + \omega \mathbf{x}(\omega) \right]^{-1}. \tag{5}$$

The variables $\mathbf{r} = \phi^T R \phi$ and $\mathbf{x} = \phi^T X \phi$ are the modal resistance and reactance matrices of the fluid, respectively, and can be combined into a single complex function $\mathbf{z} = \mathbf{r} + j\mathbf{x}$. By including the external fluid loading matrices directly into the modal transfer function $\mathbf{h}(\omega)$, the in-vacuo modes can be used as the basis set for the analysis. Similarly, discrete and generalized impedances acting on the structure can be included in (4) once transformed into modal space. For problems involving structural damping, a complex stiffness matrix can be used $\tilde{\mathbf{K}} = \mathbf{K}(1 + j\eta)$, where η is the material loss factor.

Complex forcing functions may be incorporated into the analysis by computing the modal forcing function cross-spectral density (CSD) matrix:

$$\mathbf{G}_{ff} = \phi^T \mathbf{G}_{FF} \phi, \tag{6}$$

which describes the coupling between any matrix of external forces, \mathbf{G}_{FF} , and the vibration modes. Any

stochastic forcing function that is stationary and ergodic may be represented by \mathbf{G}_{FF} . The analysis presented herein uses turbulent boundary layer (TBL) forces created by in-flow and will be discussed in Section 2.1.

The modal amplitudes caused by the forcing function can then be computed to form a modal response CSD matrix

$$\mathbf{G}_{\psi\psi} = \mathbf{h}(\omega) \mathbf{G}_{ff} \mathbf{h}(\omega)^H. \tag{7}$$

This is the modal equivalent of the multiple input-multiple output problem arranged in matrix form (Bendat and Piersol 2000). $\mathbf{G}_{\psi\psi}$ represents the modal amplitudes created by the forcing function in a stochastic sense. Once $\mathbf{G}_{\psi\psi}$ is known, the radiated sound power spectral density can be computed using

$$\mathbf{G}_{Prad} = \sum_{m=1}^M \sum_{n=1}^M r_{mn} \mathbf{G}_{\psi_m \psi_n}, \tag{8}$$

where $r_{mn} = \phi_m^T R \phi_n$ is the modal resistance matrix and M is the number of retained modes. More details on this general analysis procedure can be found in Hambric et al. (2010)

2.1 Turbulent boundary layer forcing function

The forcing function matrix for turbulence-induced wall pressures resulting from flow over a structure can be computed using the product of a pressure auto-spectrum (G_{pp}) and a pressure cross-spectral function (Γ), which will subsequently be explained:

$$\mathbf{G}_{FF} = \mathbf{G}_{pp}(f) \Gamma(\xi, 2\pi f). \tag{9}$$

The pressure spectrum model used herein is a modified version of the Chase-Howe model (Howe 1998) that has been adjusted to account for the viscous dissipation range based on Lysak (2006), defined as

$$G_{pp}(f) = \frac{3\rho^2 u_\tau^4}{f^*} \left[\frac{(f/f^*)^2}{\{(f/f^*)^2 + \hat{\alpha}^2\}^{3/2}} \right] e^{-14fv/u_\tau^2}, \tag{10}$$

where f^* is the freestream velocity divided by the boundary layer displacement thickness (U/δ^*), ρ is fluid density, u_τ is the friction velocity, ν is the kinematic viscosity, and $\hat{\alpha}$ is the turbulence constant. All flow variables are defined in Table 1. Figure 1 shows the frequency dependence of the spectrum for flow at 5.14 m/s (10 knots). The peak in the spectrum occurs at 61.5 Hz.

The cross-spectrum Γ defines the partially correlated regions of pressure over the structure and is often referred to as a coherence function. A well-known TBL coherence function model was proposed by Corcos (1963) and modified by Mellen (1990) to be

$$\Gamma(\xi_1, \xi_2, \omega) = e^{-\sqrt{(\beta_1|\omega\xi_1/U_c|^2 + \beta_2|\omega\xi_2/U_c|^2)}} e^{j\omega\xi_1/U_c}, \tag{11}$$

Table 1 TBL flow parameters

Density (ρ)	1000 kg/m ³
Friction velocity (u_τ)	0.1542 m/s
Freestream velocity (U)	5.14 m/s (10 knots)
Convective velocity (U_c)	3.598 m/s
Boundary layer disp thickness (δ^*)	0.0022 m
Kinematic viscosity (ν)	1.15e-6
Streamwise decay constant (β_1)	0.11
Spanwise decay constant (β_2)	0.7
Turbulence constant ($\hat{\alpha}$)	0.019

where ξ is the separation distance (the streamwise direction denoted with subscript 1 and spanwise with subscript 2) between all points on the structure, β is the decay constant (streamwise and spanwise), and U_c is the convective flow velocity.

2.2 Numerical implementation

In general, this structural-acoustic analysis procedure is independent across all frequencies and is therefore straightforward to implement as a loop over a pre-determined frequency range which may span hundreds or thousands of Hertz. However, storage requirements for the acoustic resistance and reactance matrices, as well as the forcing function matrix, can be very large if they are computed a priori for every frequency in the analysis set since the matrices are fully populated. Because these matrices vary with frequency in a relatively smooth manner, they can be computed and stored at a much-reduced set of frequencies and interpolated to a frequency of interest. This idea was originally proposed for acoustic matrices in Benthien (1989). As an example of the smoothness in the frequency of these matrices, the modal force for TBL flow is shown for a single mode in Fig. 2.

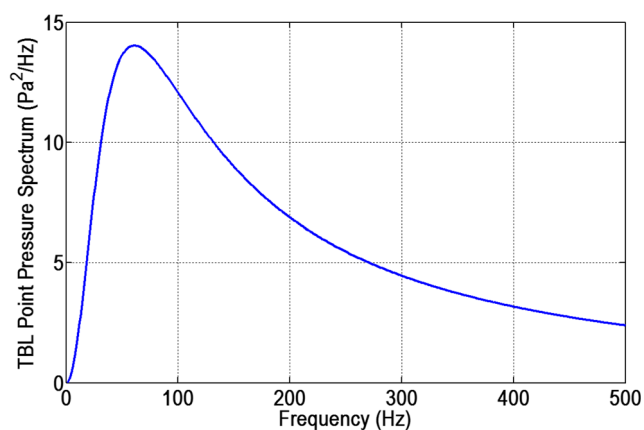


Fig. 1 The point pressure spectrum (modified Chase model) using flow parameters found in Table 1

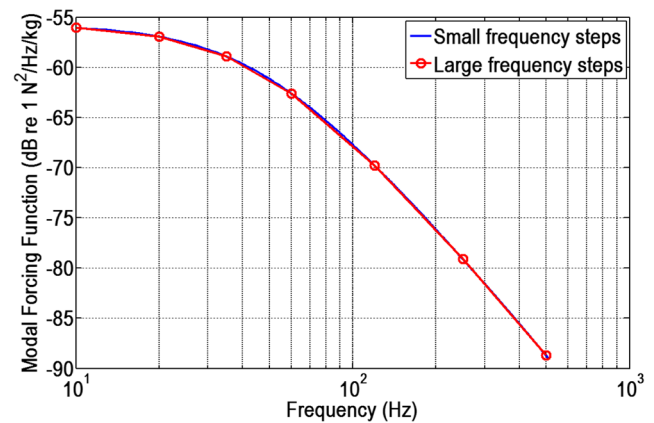


Fig. 2 The modal forcing function is smooth and can be accurately interpolated to the analysis frequency of interest using only a small of frequencies

Typical frequency-domain, structural-acoustic analyses using the formulations described herein would load several acoustic and forcing function matrices into memory, interpolate to the desired solution frequency, and solve the governing system of equations. This process would repeat for all desired solution frequencies, removing some and loading new acoustic and forcing function matrices into memory for interpolation as the sweep through frequencies progresses. For optimization problems, as described herein, many thousand solutions at each of these frequencies may be required. To avoid repetitive reading and interpolation of the matrices, the problem can be divided amongst processors and have each processor load and interpolate the matrices only once at the start of the simulation. As long as the basis set does not change during the simulation, the modal acoustic and forcing function matrices are static.

Using this approach, the number of total processors used for the analysis is determined by the number of analysis frequencies. All necessary matrices are loaded by the master node and interpolated incrementally to each frequency used in the analysis. Each processor is then assigned a frequency to solve and sent via standard message passing interface (MPI) methods its corresponding matrices (mode shape, boundary element matrices, forcing function and geometry files), storing them in RAM. Depending upon the number of modes in the system and the available RAM per node, the number of processors per node may be under-subscribed to ensure sufficient RAM is available for the solutions. Alternatively, and for very large problems, the solution could be further divided using MPI and parallel solution techniques to solve the resulting linear system of equations. This, however, is not within the scope of the present work but is of interest for future work. The use of modal space and residual vectors enables relatively large solutions (thousands of modes) on computer systems which have modest RAM levels (256 GB).

Load balancing is automatic, and RAM usage is essentially identical for all processors when using the approach described herein because each processor solves the same system of equations using a direct approach, with only matrix terms differing between frequencies/processors. For each generation of the optimization, new design variables (in the form of added impedances) are proposed and sent to each individual processor, the analysis is performed, and scalar results are returned to the master node. The net data transfer between processors using this approach is relatively small, consisting of a few scalar values for all transfers subsequent to the initial matrix transfer. The master node compiles results of all processors and returns the objective values to the optimizer. This process is summarized in Fig. 3 in the form of a flowchart.

A key enabler to these simulations is the use of residual vectors at locations where an external impedance is applied, such as an added point mass, to account for truncated modes in the conversion from physical to modal space (Roy and Girard 2005). Residual vectors allow for the same basis set to be used for varying added impedances, which allows for the acoustic and forcing function matrices to be computed and stored by each processor. Without residual vectors, system normal modes (and hence modal acoustic and function matrices) would need to be recomputed for each selection of added impedance as defined by the optimizer.

As a final note, this parallel analysis procedure would be advantageous for other types of analysis involving

Table 2 Dimensions and material properties of the curved panel

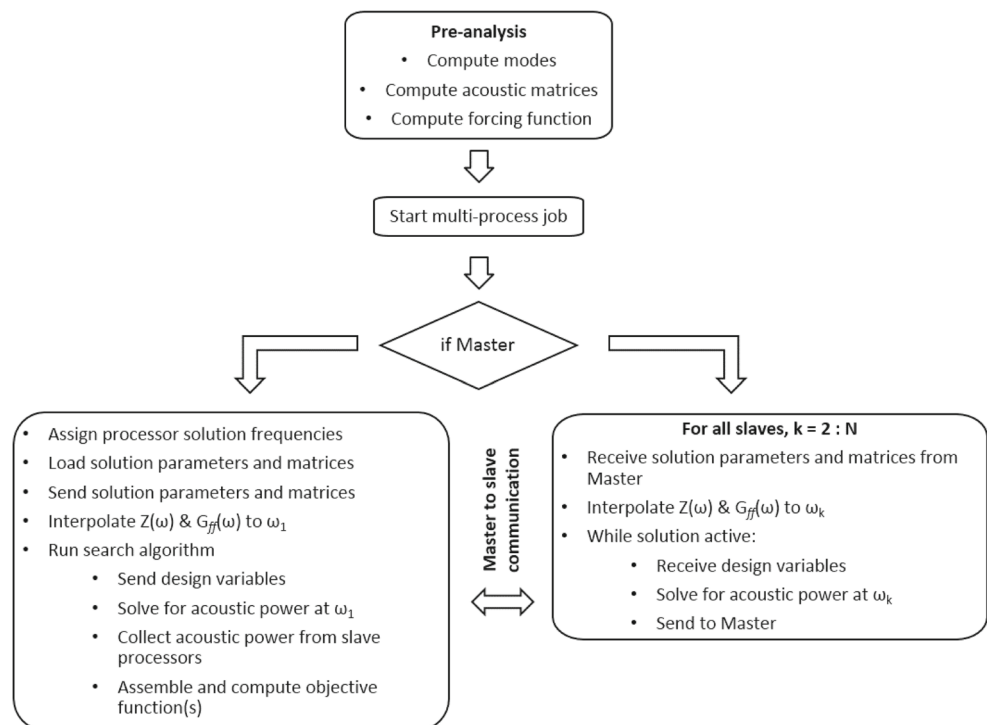
Streamwise length	2.086 m
Spanwise length	0.8 m
Thickness	0.0318 m
Radius of curvature	9.96 m
Modal loss factor	0.02
Young’s modulus	195 GPa
Density	7700 kg/m ³
Poisson’s ratio	0.28

large numbers of repetitive calculations such as uncertainty analysis (Wixom et al. 2019).

3 Curved panel optimization

To demonstrate this parallel implementation for SAO, a curved panel in heavy fluid is used (Shepherd and Hambric 2014). The panel is made of steel and meshed with a 32×84 set of rectangular, linear shell elements. The spanwise edges of the panel are pinned, while the streamwise edges are free. The dimensions of the panel are shown in Table 2 and a basic diagram of the panel is shown in Fig. 4. Twelve discrete masses, fixed in space, are used as the design variables, similar to work by Ratle and Berry (1998). Each mass is modeled as a discrete impedance $C(\omega) = j\omega M_i$

Fig. 3 A master node controls all aspects of the optimization job including loading input matrices, controlling the information exchange between nodes, and running the search algorithm. Slave nodes are assigned a specific analysis frequency and solve for the sound power at that frequency when passed a set of design variables



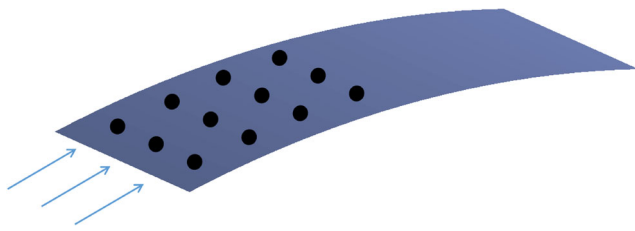


Fig. 4 The curved panel used in the optimization. The arrows indicate the direction of the flow while the black dots represent discrete masses which vary during the optimization. The short edges perpendicular to the flow are simply supported while the long edges parallel to the flow are free

and transformed to modal space, $\mathbf{c} = \phi^T \mathbf{C} \phi$. Equation (4) then becomes

$$\mathbf{h}(\omega) = j\omega \left[-\omega^2 \mathbf{m} + \tilde{\mathbf{k}} - j\omega \mathbf{z}(\omega) + \mathbf{c}(\omega) \right]^{-1}, \quad (12)$$

where structural damping is used and viscous damping is neglected. When computing $\mathbf{c}(\omega)$, the portions of the mode shape matrix which do not have an attached mass are zero and can be reduced out. For this analysis, the discrete masses were bounded between 1 kg and 10 kg.

Fifty modes of the panel, computed using the commercial software NX/Nastran and shown in Fig. 5, are reformatted for use in the optimization loop. During the optimization, the acoustic matrices are assumed to not vary significantly since the base structure does not change and the total response changes are small compared with the wavelengths of interest. This results in significantly reduced computational times since the acoustic matrices are only computed

once and stored (Fahnline et al. 2006; Shepherd and Hambric 2012). Additionally, the Mellen coherence function was found to be in the low wavenumber regime (Shepherd and Hambric 2014; Bonness et al. 2017) for 5 m/s flow and therefore the effective correlation area is given by

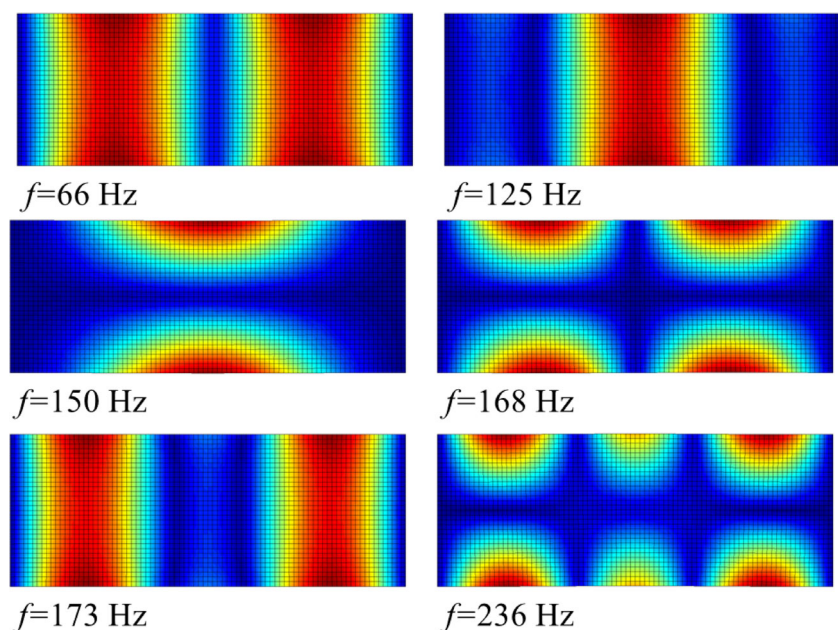
$$\frac{\pi}{2} \left(\frac{\beta_1^3}{(1 + \beta_1^2)^{3/2}} \right) \left(\frac{2U_c}{\beta_1 \omega} \right) \left(\frac{2U_c}{\beta_2 \omega} \right). \quad (13)$$

Table 2 contains the flow parameters used in the simulation.

Two competing objective functions were defined as total mass and the integrated pressure power spectral density (He et al. 2017; He and Sun 2018). Recently performed multiobjective optimization of an elastic beam for noise reduction and used similar objectives. The noise objective was determined by integrating the radiated sound power between 10 Hz and 490 Hz, a procedure similar to other recent works (Shepherd and Hambric 2014; Klaerner et al. 2017). Since the noise objective will be many orders of magnitude smaller than the range of the mass objective, it was normalized by a factor of 1e-10. This helps the optimizer to make equal progress in reducing both objectives.

Non-uniform frequency resolution was used in order to achieve sufficient resolution at the low-frequency resonance peaks without over resolving the high-frequency resonance peaks. Since the number of required processors depends on the number of solution frequencies, using a larger frequency spacing requires fewer processors. However, if the resonance peaks are not adequately resolved with at least five points per half-power bandwidth, the objective function, which is an integration of power over

Fig. 5 Curved panel mode shapes and natural frequencies. The red regions represent high deflection (positive and negative), while the blue regions represent no motion



frequency, will not accurately represent the response and the optimization results may be erroneous. Therefore, a non-uniform frequency vector was determined which would adequately resolve the response at all frequencies.

3.1 Pareto optimality and search methods

Two objectives are considered in this study, and the optimization problem is formally stated as

$$\min\{\mathbf{F}(\mathbf{x})\}, \quad \mathbf{x} \in \Omega.$$

$$\mathbf{x}_i^L \leq \mathbf{x}_i \leq \mathbf{x}_i^U$$

$\mathbf{F} = (F_1, F_2)$ is the objective function vector while the design variable vector $\mathbf{x} = (x_1, x_2, \dots, x_N)$ represents N real-valued design variables within the design space Ω . The lower and upper bounds are set by \mathbf{x}_i^L and \mathbf{x}_i^U , respectively. The objective functions F_1 and F_2 are radiated noise and total mass, respectively. For the study shown here, no constraints are included such that the feasible region is equivalent to Ω . To address the trade-offs between the competing objectives, the concept of Pareto dominance must be introduced. A vector of design variables x_a is said to dominate another design variable vector x_b when $\mathbf{F}(x_a) \leq \mathbf{F}(x_b)$. If x_a is superior in one objective but not all objectives, then x_a is considered non-dominated (Coello Coello et al. 2007). When x_a is not superior for any objective, it is dominated within Ω . The set of all non-dominated points is known as the Pareto front and defines the trade-off between all objective functions being considered. Estimating the Pareto front is particularly useful when the objective function weighting is not well-defined and/or preferences have not yet been identified.

Since the objective function is expected to be highly nonlinear and discontinuous, where gradient methods typically do not perform well, an evolutionary algorithm was used to perform the search. The search algorithm used for this study is a multiobjective evolutionary algorithm (MOEA) developed by Hadka and Reed (2013). This MOEA, referred to by its developers as Borg, is an auto-adaptive, optimization framework for use on multiobjective, multimodal problems. The goal of Borg is to capture or closely approximate the Pareto front on real-valued test problems by utilizing robust selection, crossover, and mutation operators to minimize multiple, competing objectives. It has been successfully applied to difficult optimization problems in many fields, including water resource management and vibration reduction (Hadka and Reed 2015; Giuliani et al. 2018; McCormick and Shepherd 2018).

Borg has been shown to be a robust MOEA which outperforms many other MOEAs on standard demonstration

functions (Hadka and Reed 2013). While the details of Borg can be obtained from Hadka and Reed, several important features will be summarized. First, it utilizes ϵ -dominance to obtain both convergence and diversity in the multiobjective search. ϵ -dominance divides the search space into hypercubes with ϵ , which can vary for each objective, defining the characteristic lengths of the hypercube. When a dominant solution is detected, it is added to the archive of dominant solutions. Solutions in the archive set the population size and are randomly selected for recombination to create future generations. Additionally, the ϵ "box" is used to define the minimum threshold for improvement.

Second, Borg utilizes restarts to revitalize the search and escape local optima when search stagnation is detected, as defined by the minimum threshold for improvement ϵ . During a restart, the population size is adapted in order to remain proportional to the archive size. The tournament selection size is also adjusted to keep an elitist selection (Hadka and Reed 2013). Additionally, the population is flushed and refilled with solutions from the archive, as well as randomly selected and mutated archive solutions. Since stagnation triggers randomized restarts, Borg will never stop evaluating the design space but will continue to fill in points along the Pareto front. For this reason, a stop criterion based on a number of function evaluations is used. If this criterion is not considered sufficient, a restart capability is used to continue the search.

Finally, Borg adaptively selects recombination operators based on their success, an idea originally proposed by Vrugt and Robinson (2007) but improved upon in the Borg. Traditionally, an evolutionary algorithm will utilize a single, pre-set operator involving a set of crossover and/or mutation rules to determine the design variables for the following generation. However, it is difficult to know a priori if the pre-set operator criteria are well suited for the problem being considered (Hadka and Reed 2013). Additionally, multiple regions of the search space may be better suited for different types of operators. To overcome this, a set of six recombination operators are used in Borg, with each operator given a certain probability of being used. As the search progresses, the probability of using each operator is updated based on the number of dominant solutions produced by that operator. Thus, the operators which produce good results are rewarded with a higher probability of usage in future generations. This allows Borg to adjust its recombination to best search the design space in future generations. The recombination operators used in Borg are simulated binary crossover, differential evolution, parent-centric crossover, unimodal normal distribution crossover, simplex crossover, and uniform mutation. The first four recombination operators listed here also include polynomial

mutation (Hadka and Reed 2013). A schematic for the search procedure is shown in Fig. 6.

It should be noted that the Borg architecture is already well suited for parallel computing in terms of evaluating the independent design variables (Hadka and Reed 2015). However, the bulk of computation time for most structural-acoustic problems, such as the curved panel shown here, is spent computing the objective function and the time savings is most significant when performed as has been outlined. If a cluster is available with a very large number of processors, the Borg parallelization could be merged with the procedure from Section 2.2 for additional speed ups.

3.2 Curved panel results

Optimization for the curved panel was run on a Linux cluster having 864 2.8 GHz AMD processors, 48 processors per compute node and 256 GB RAM per processor. A frequency vector of 405 frequencies ranging from 10 to 490 Hz was used divided between 405 processors and run for 100,000 function evaluations. The resulting Pareto front, which illustrates the trade-off between mass and radiated

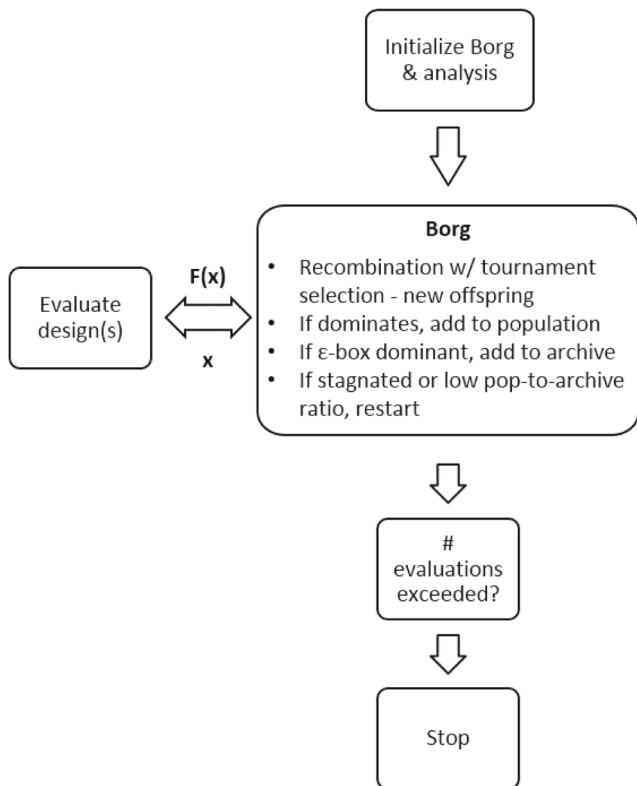


Fig. 6 The multiobjective optimizer loop. Only the number of function evaluations and the desired resolution (ϵ) along the Pareto front are required as inputs

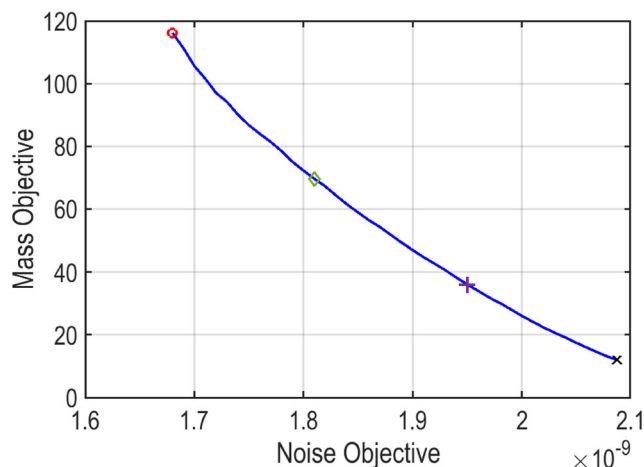


Fig. 7 The Pareto front illustrates the trade-off between radiated noise and total added mass. The four marks indicate optimal designs with different preferences. The red circle emphasizes reduced noise while the black x emphasizes reduced mass

noise reduction, is shown in Fig. 7. On the right end of the curve, the mass is minimized and the highest noise value is obtained. On the left end of the curve, the noise is minimized and the mass objective is highest.

The radiated noise spectrum is shown for the two end point designs in the Pareto front in Fig. 8. As expected, the optimal design for maximum radiated noise reduction is the case with the most added mass. While the total reduction over the 10–490 Hz range is approximately 1.0 dB, reductions at individual peaks are each different, as listed in Table 3. A frequency which experiences a slight increase in level is indicated in Table 3 by a negative sign.

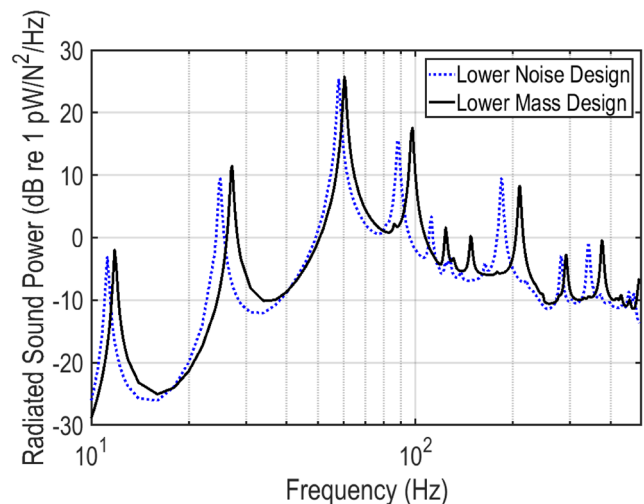


Fig. 8 The radiated noise for the two endpoint designs from the Pareto front. These cases represent the lower mass design (dotted line) and the lower noise design (solid line)

Table 3 Reduction in radiated noise between the two endpoint designs in the Pareto front

Approx. frequency (Hz)	Noise reduction (dB)
11	1.1
25	2.1
58	0.5
88	2.0
215	-1.2
10-490	1.0

The values of the masses for all points along the Pareto front are shown in Fig. 9. On the left side, labeled low noise, the masses are at or near the upper limit of 10 kg. As they shift down to their lowest value, they are clustered in four groups which follow similar behavior. The clustering occurs due to the sensitivity of each specific mass location with respect to reducing radiated noise. The mass values are shown in their respective locations in Fig. 10 for four optimal designs equally spaced along the Pareto front (see Fig. 7). The top figure illustrates the far-left point on the Pareto front, while the bottom figure illustrates the far-right point. The masses are colored according to their values and illustrate the regions of the panel that are most and least sensitive to added mass and radiated noise.

In Fig. 10, the clustering of design variables seen in Fig. 9 will be examined. The three masses closest to the center of the panel change to the lower limit of mass (1 kg) the fastest indicating the center positions are the least sensitive to mass. The masses along the edge locations are slightly more sensitive. The first three masses along the center of the panel are shown to have the highest sensitivity to mass

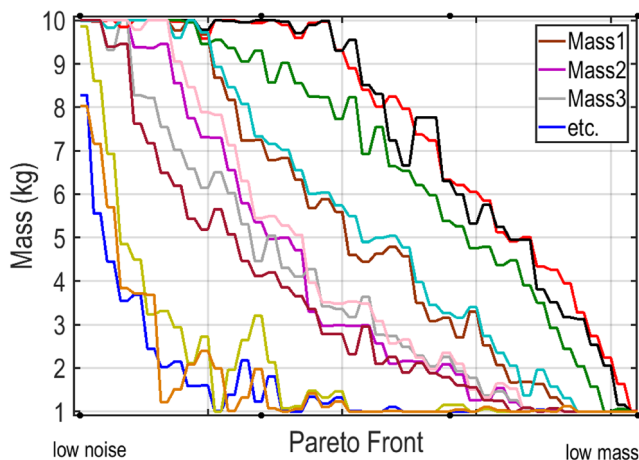


Fig. 9 The optimal mass values for each point along the Pareto front. The mass cluster together in four groups

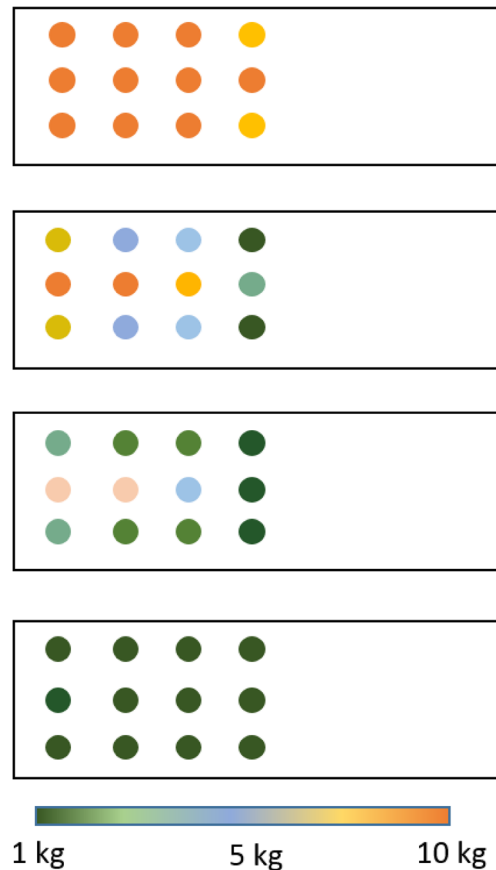


Fig. 10 Topview of the curved panel with four design points along the Pareto front (see Fig. 7). The top design illustrates lower noise, while the bottom design is for lower mass. The two middle designs show the trade-off between mass and noise

since they stay at their upper limit the longest and drop to the lower limit the latest.

3.3 Timing study

A timing study was then performed using different numbers of processors with the results shown in Table 4. A 480-length vector of frequencies was used in the timing study. For a single compute node (48 processors for this cluster),

Table 4 Timing study

Number of nodes	Number of processors	Evaluation time
1	1	10 h 42 min 16 s
1	48	34 min 10 s
2	96	18 min 20 s
5	240	9 min 28 s

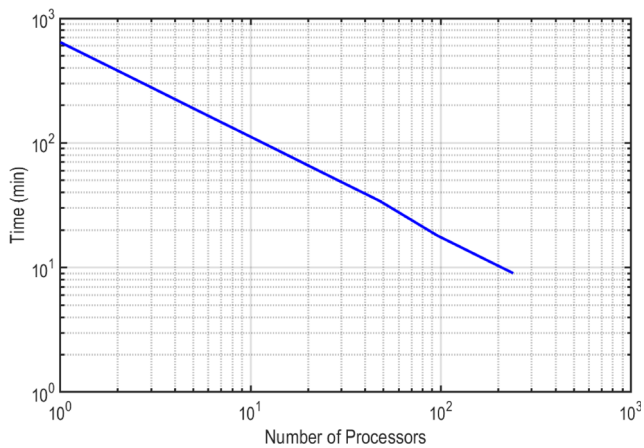


Fig. 11 Compute time for the SAO process to run with 100,000 function evaluations

the compute time went from over 10 h on a single processor to 34 min. The computation time decreases to less than 10 min when 5 compute nodes are used.

The scaling is shown graphically in Fig. 11 and illustrates nearly logarithmic behavior. The speed up is also shown in Fig. 12 when compared against the ideal condition. As the time of calculation and communication becomes more balanced, the scaling should be more linear with a number of processors and the speedup is expected to move toward the ideal speedup condition. However, if the problem size becomes very large, memory issues would become more important and the initial loading and interpolation of the input matrices may need to be performed incrementally. This has been and will continue to be considered in future work.

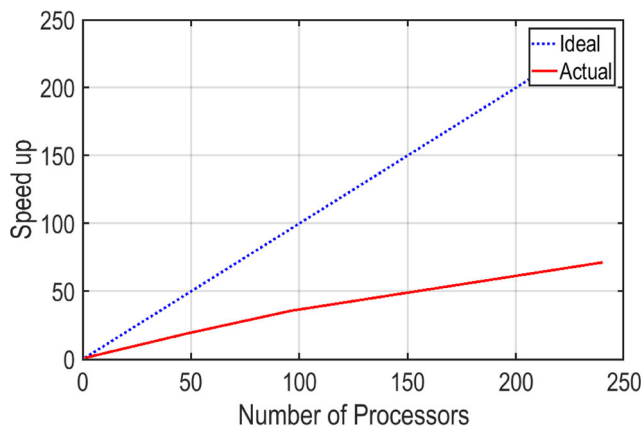


Fig. 12 The speed up as more processors are used in the optimization process. The ideal speedup is shown as the dotted line

4 Conclusions

A parallel SAO procedure has been introduced to reduce the noise of vibrating structures excited by complex forcing functions such as turbulent flow. The use of high-performance computing coupled with residual vectors and modal techniques reduces the analysis time such that multiobjective optimization can be performed on large-scale problems. The search algorithm is run on a master processor which coordinates the parallel solution between the master and slave processors. The procedure is demonstrated on a curved underwater panel with attached masses.

The trade-off between reduced noise and reduced mass was determined for a curved underwater panel excited by turbulent boundary layer flow. The middle region of the panel was found to be most sensitive to reducing noise using added masses. By utilizing 240 compute processors on a computer cluster, a 10+-h job on a single processor was reduced to under 10 min (67× speed up). This savings in computation time allows for noise and vibration concerns to be addressed in the design stage.

5 Replication of results

To replicate the results, the Borg search algorithm can be downloaded at <http://borgmoea.org/>.

Acknowledgments The authors would like to thank Dr. Peter Lysak for his recommendations regarding the forcing functions used in this paper.

Compliance with ethical standards

Conflict of interest The authors declare that they have no conflict of interest.

References

- Aage N, Andreassen E, Lazarov BS (2015) Topology optimization using PETSc: an easy-to-use, fully, parallel, open source topology optimization framework. *Struct Multi Discip Optim*, 51
- Aage N, Andreassen E, Lazarov BS, Sigmund O (2017) Giga-voxel computational morphogenesis for structural design. *Nature*, 550
- Aiello R, Auriemma F (2018) Optimized vibro-acoustic design of suspended glass panels. *Struct MultiDiscip Optim* 58:2253–2268
- Belegundu AD, Salagame RR, Koopmann GH (1994) A general optimization strategy for sound power minimization. *Struct MultiDiscip Optim* 8(2–3):113–119
- Bendat JS, Piersol AG (2000) *Random data*. Wiley, New York
- Benthien GW (1989) Application of frequency interpolation to acoustic-structure interaction problems, Naval Ocean System Center TR-1323

- Bonness WK, Fahline JB, Lysak PD, Shepherd MR (2017) Modal forcing functions for structural vibration from turbulent boundary layer flow at low and high speed. *J Sound Vib* 395:224–239
- Butkewitsch S, Steffen V (2001) A case study on frequency response optimization. *Int J Solids Str* 38:1737–1748
- Coello Coello CA, Lamont GB, Van Veldhuizen DA (2007) *Evolutionary algorithms for solving multi-objective problems*. Springer, New York
- Constans EW, Koopmann GH, Belegundu AD (1998) The use of modal tailoring to minimize the radiated sound power of vibrating shells: theory and experiment. *J Sound Vib* 217(2):335–350
- Corcus GM (1963) Resolution of pressure in turbulence. *J Acoust Soc Am* 35(2):192–199
- Cunefare KA (1994) Optimization techniques in structural acoustic design. *J Acoust Soc Am*, 95
- de Souza MZ, Ramirez Pozo AT (2014) A GPU implementation of MOEA/D-ACO for the multiobjective traveling salesman problem. In: *Proceedings of the 2014 Brazilian conference on intelligent systems*
- Durillo JJ, Nebro AJ, Luna F, Alba E (2008) A study of master-slave approaches to parallelize NSGA-II. In: *Proceedings of 2008 IEEE international symposium on parallel and distributed processing*
- Fahline JB, Koopmann GH (1996) A lumped parameter model for the acoustic power output from a vibrating structure. *J Acoust Soc Am* 100(6):3539–3547
- Fahline JB, McDevitt TE, Whitney EJ, Capone DE (2006) Structural-acoustic tailoring of metal structures by laser free-forming. *Noise. Control Eng J* 54(2):124–129
- Giuliani M, Quinn J, Herman J, Castelletti A, Reed PM (2018) Scalable multi-objective control for large scale water resources systems under uncertainty. *IEEE Trans Cont Syst Tech* 26(4)
- Hadka D, Reed P (2013) Borg: an auto-adaptive many-objective evolutionary computing framework. *Evol Comp* 21(2):231–259
- Hadka D, Reed P (2015) Large-scale parallelization of the borg multiobjective evolutionary algorithm to enhance the management of complex environmental systems. *Env Model & Soft* 69:353–369
- Hambric SA, Boger DA, Fahline JB, Campbell RL (2010) Structure- and fluid-borne acoustic power sources induced by turbulent flow in 90° piping elbows. *J Fl Str* 26:121–147
- He MX, Sun JQ (2018) Multi-objective structural-acoustic optimization of beams made of functionally graded materials. *Comp Str* 185
- He MX, Xiong FR, Sun JQ (2017) Multi-objective optimization of elastic beams for noise reduction. *J Vib Acs*, 139
- Howe MS (1998) *Acoustics of fluid-structure interaction*. Cambridge Univ Press, Cambridge
- Jeon JY, Okuma M (2008) An optimum embossment of rectangular section in panel to minimize noise power. *J Vib Acoust* 130:021012-1–021012-7
- Johnson WM, Cunefare KA (2002) Structural acoustic optimization of a composite cylindrical shell using FEM/BEM. *J Vib Acoust* 124:410–414
- Joshi P, Mulani SB, Slemple WCH, Kapania RK (2012) Vibro-acoustic optimization of turbulent boundary excited panel with curvilinear stiffeners. *J Aircraft* 49(1):52–65
- Joshi P, Mulani SB, Kapania RK (2015) Multi-objective vibro-acoustic optimization of stiffened panels. *Struct MultiDiscip Optim* 51:835–848
- Klaerner M, Wuehrl M, Kroll L, Marburg S (2017) FEA-based methods for optimising structure-borne sound radiation. *Mech Sys Sig Proc* 89:37–47
- Koopmann GH, Fahline JB (1997) *Designing quiet structures*, Academic Press
- Lysak PD (2006) Modeling the wall pressure spectrum in turbulent pipe flows. *J Fluids Eng* 128:216–222
- Marburg S (2002) Developments in structural-acoustic optimization for passive noise control. *Archives of Computational Methods in Engineering State of the art reviews* 9(4):291–370
- Marburg S, Dienerowitz F, Fritze D, Hardtke HJ (2006) Case studies on structural-acoustic optimization of a finite beam. *Acta Acustica united with Acustica* 92(3):427–439
- Marburg S, Shepherd MR, Hambric SA (2016) *Structural-acoustic optimization, engineering vibroacoustic analysis: methods and applications*, Wiley
- McCormick CA, Shepherd MR (2018) Optimal design and position of an embedded one-dimensional acoustic black hole. In: *Proceedings of internoise 2018*. Chicago
- Mellen RH (1990) On modeling convective turbulence. *J Acoust Soc Am* 88(6):2891–2893
- Ratle A, Berry A (1998) Use of genetic algorithm for the vibroacoustic optimization of a plate carrying point masses. *J Acoust Soc Am* 104(6):3385–3397
- Reed P (2016) *High performance computing with the MOEA framework and ignite*, CreateSpace Independent Publishing Platform, USA
- Roy N, Girard A (2005) Impact of residual modes in structural dynamics, spacecraft structures. *Mater Mech Test* 2005:581
- Shepherd MR, Hambric SA (2012) An approach for structural-acoustic optimization of ribbed panels using component mode synthesis. In: *Proceedings of internoise 2012/ASME NCAD IN12-592*
- Shepherd MR, Hambric SA (2014) Minimization the acoustic power radiated by a fluid-loaded curved panel excited by turbulent boundary layer flow. *J Acoust Soc Am* 136(5):2575–2585
- Shepherd MR, Fuertado FA, Conlon SC (2016) Multi-objective optimization of acoustic black hole vibration absorbers. *J Acoust Soc Am* 140(3):EL227–EL230
- Vrugt JA, Robinson BA (2007) Improved evolutionary optimization from genetically adaptive multimethod search. *Proc Natl Acad Sci* 104(3):708–711
- Walsh TF, Aquino W (2017) Massively parallel structural acoustics for forward and inverse problems. *J Acoust Soc Am* 143:3588
- Walsh TF, Hammetter C, Sinclair MB, Shaklee-Brown H, Bishop J, Aquino W (2018) Design, optimization and fabrication of mechanical metamaterials for vibration control. *J Acoust Soc Am* 143:1917
- Wixom AS, Walters GS, Martinelli SL, Williams DM (2019) Generalized polynomial chaos with optimized quadrature applied to a turbulent boundary layer forced plate. *J Comp Nonlinear Dyn* 14:021010-1–021010-9
- Wong WO (2016) Optimal design of a hysteretic vibration absorber using fixed-points theory. *J Acoust Soc Am* 139(6):3110–3115
- Yuksel E, Kamci G, Basdogan I (2012) Vibro-acoustic design optimization study to improve the sound pressure level inside the passenger cabin, vol 134

Publisher's note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.