



Generalized Bloch mode synthesis for accelerated calculation of elastic band structures



Dimitri Krattiger, Mahmoud I. Hussein*

Department of Aerospace Engineering Sciences, University of Colorado Boulder, CO 80309, USA

ARTICLE INFO

Article history:

Received 4 August 2017

Received in revised form 20 November 2017

Accepted 10 December 2017

Available online 15 December 2017

Keywords:

Phononics

Elastic metamaterials

Band structure

Dispersion curves

Bloch mode synthesis

Component mode synthesis

ABSTRACT

The Bloch mode synthesis (BMS) model-reduction method adapts component mode synthesis techniques to unit-cell problems in order to obtain a reduced-order model that quickly produces band-structure frequencies for any wave vector, or vice versa. Fundamental to BMS is a partitioning of the real-space model into interior and boundary components, and subsequent reduction of the interior via truncated normal mode expansion. In this paper, two enhancements are presented for the BMS method that reduce both computation time and error in band-structure calculations. The first enhancement improves the accuracy of the interior reduction by approximating the participation of the residual modes rather than simply truncating them. The original formulation of BMS includes a modal reduction of the boundary that must be recomputed for every wave vector. This limits computational benefits and prevents the reduced-order model from being useful for the inverse band-structure problem (i.e., the $\mathbf{k}(\omega)$ calculation). The second enhancement is a local boundary reduction that is independent of wave vector and thus does not suffer from the aforementioned limitations.

© 2017 Elsevier Inc. All rights reserved.

1. Introduction

Band structures fundamentally describe the characteristics of waves in materials and can be used to understand most if not all material properties. These waves may be electronic, electromagnetic, or mechanical in nature, with each class giving rise to its own set of material properties. In the current work, we focus on mechanical waves in periodic materials. This encompasses lattice vibrations at the atomic scale up to seismic waves at the macroscale. Regardless of the length scale, a material with spatial periodicity that influences the transfer of vibrational energy is referred to as a *phononic crystal* [1]. The phononic band-structure is a relationship between the wave number (or wave vector) and frequency for all vibrational waves admitted in the material. Band-structure diagrams highlight many interesting wave propagation phenomena including band gaps – frequency ranges within which waves cannot propagate. Band gaps are particularly well studied in electronic materials, where their presence (or absence) determines whether a material is an insulator, a semiconductor, or a conductor [2]. Acoustic/elastic metamaterials are a fascinating class of materials that has emerged in recent years. Named for their ability to produce dynamical properties outside the range of naturally occurring materials (such as negative refraction), these materials typically make use of resonators that can interact with wavelengths much larger than the unit-cell length scale.

* Corresponding author.

E-mail address: mih@colorado.edu (M.I. Hussein).

Using Bloch's theorem [3], a periodic material can be exactly modeled by considering a single unit cell. The solution field is approximated over the unit-cell domain using a set of basis functions, thus providing a numerical discretization. Numerous discretization methods exist and have been adapted to band-structure calculation problems. We focus on models obtained using the finite element (FE) method. This method has been studied extensively in the field of structural dynamics. It achieves good convergence, allows for local refinement to handle sharp gradients in the solution or material distribution, and produces sparse matrices which reduce the memory and computational requirements. The band structure is obtained by either looping through all of the wave vectors of interest and solving each resulting eigenvalue problem for the frequencies ($\omega(\mathbf{k})$ method), or by looping through all of the frequencies of interest and solving each resulting eigenvalue problem for the corresponding wave vectors ($\mathbf{k}(\omega)$ method).

Depending on the complexity of the material unit cell, and the level of accuracy required, the system of equations can be very large. Many algorithms exist for accurately and efficiently obtaining eigen solutions to systems of equations, but even so, the computations are very time consuming for large systems. This is exacerbated by the necessity to loop through and solve at multiple frequencies or wave vectors to obtain the band structure.

Due to the high computational cost, there have been several research efforts dedicated to more efficient band-structure calculation. Axemann and Kuchment noted that direct eigenvalue solution is very costly and does not take advantage of the sparsity of the system matrices [4]. They proposed a simultaneous coordinate over-relaxation (SICOR) [5] approach to obtain the extreme-valued eigenvalues iteratively, so that the sparsity of the system is exploited. Although SICOR has not seen widespread use, other eigenvalue algorithms that use an iterative approach and take advantage of the system's sparsity have gained considerable traction. Perhaps the most popular of these iterative eigenvalue solvers is the implicitly restarted Arnoldi method (implemented as `eigs` in MATLAB [6]). Dobson developed a Fourier preconditioner that takes advantage of the material periodicity, and can speed up convergence of iterative solvers [7]. He also examined the idea of using eigenvectors obtained at a previous \mathbf{k} point as a guess for the next \mathbf{k} point, thus further improving convergence of the iterative solver. Chern et al. used a multigrid solution method embedded in an inverse eigenvalue solver to speed up convergence of iterative eigenvalue solvers [8].

The previously discussed methods are primarily focused on speeding up convergence of iterative solvers. Another approach to speeding up band-structure calculation is model-order reduction whereby the system of equations is reduced in size so it may be solved more quickly. Depending on the reduction, the solution may be exact or approximate. For example, Guyan reduction [9,10], also known as static condensation, splits the problem into slave and master degrees of freedom (DOFs) and assumes zero inertial contribution from the slave DOFs. This allows the entire displacement field to be expressed in terms of just the master DOFs, thus reducing the problem size. This reduction is exact for static problems (i.e. $\omega = 0$), but is increasingly inaccurate the higher the frequency. The dynamic reduction works very similarly, but does not assume negligible inertia. Rather, it performs the reduction about some test frequency. The reduced model is then exact for that test frequency and decreases in accuracy as we move away from that frequency. This form of reduction is often used when solving the $\mathbf{k}(\omega)$ problem because a frequency must be selected before solving each eigenvalue problem anyway.

Modal reduction is another common form of reduced-order modeling that can be adapted from structural mechanics and applied to periodic materials. The mode shapes of interest for the full model are computed at a few key points throughout the Brillouin zone. These mode shapes are then collected and used as a reduced basis for other points in the Brillouin zone [11]. The key points may be selected using high-symmetry points following the reduced Bloch mode expansion (RBME) method [12], or can be obtained using an adaptive process [13]. Another form of modal reduction extends component mode synthesis (CMS) techniques to band-structure problems [14–17]. Compactly referred to as Bloch mode synthesis (BMS) [14], this approach decomposes the domain into interior and boundary DOFs, and then describes the interior with a truncated modal basis while leaving the boundary conditions in their original “real-space” form. This allows boundary conditions, used to enforce Bloch's theorem, to be applied as they normally would. The reduced BMS model is useful for both $\omega(\mathbf{k})$ and $\mathbf{k}(\omega)$ band-structure calculations. The latter was recently demonstrated by Palermo and Marzani [17]. The most important benefit of BMS is that the same reduced model can be used for any wave vector (or frequency), because application of Bloch boundary conditions occurs *after* the model reduction.

After a reduction of the interior DOFs, a majority of the remaining DOFs are typically contained in the boundary set. The characteristic constraint (CC) mode reduction was developed in order to perform a reduction of the interface size for substructured models [18]. This reduction is performed on the coupled interface so it is referred to here as system-level CC (S-CC). The S-CC reduction was applied to the unit-cell problem in reference [14]. Since the boundary has to be self coupled (by Bloch boundary conditions) before modal reduction, the reduction must be applied independently for every \mathbf{k} point. This limits efficiency gains from the reduction, and also precludes the reduced model from being solved via the $\mathbf{k}(\omega)$ method. In this paper, we augment the BMS method by making use of a local-level CC (L-CC) interface reduction. The L-CC reduction is designed in the context of CMS to reduce substructure interfaces locally (i.e., before assembling the substructures to form the system). When applied to unit-cell problems, L-CC allows for the boundary to be reduced prior to enforcing Bloch boundary conditions. Local interface reduction provides significant computational benefits, and also allows the reduced model to be solved via the $\mathbf{k}(\omega)$ method as well as the $\omega(\mathbf{k})$ method.

A second enhancement is used to improve the accuracy of the BMS reduced-order model. The reduction of the unit-cell interior is possible because the majority of the dynamically important information can be represented with a small number of normal modes. The higher frequency modes that are discarded are called residual modes. Including more residual modes in the model drives down the error, but increases the reduced-order model size and thus computation time. Rather than

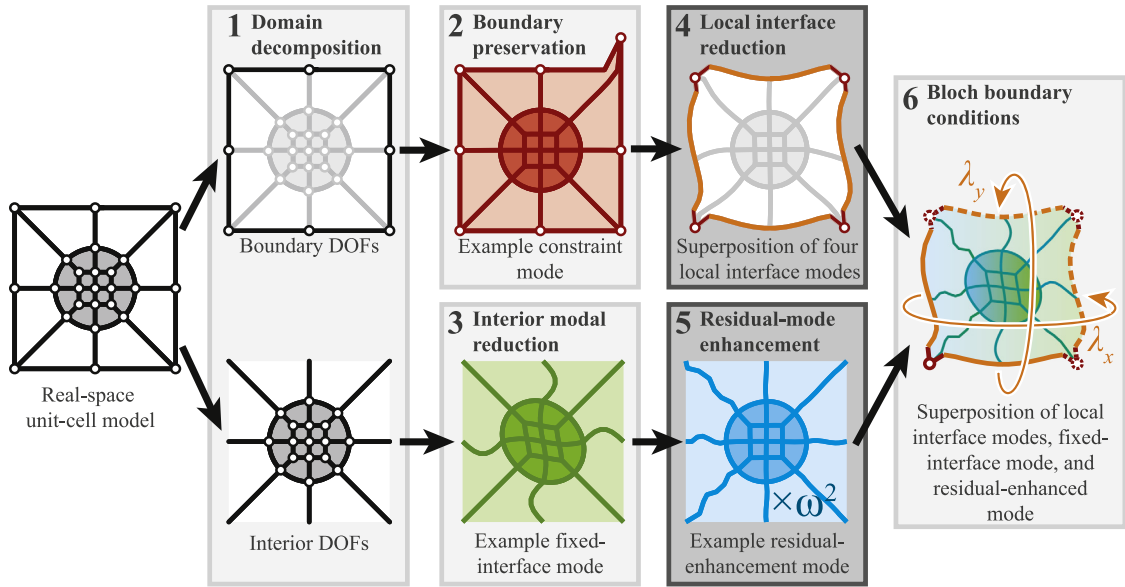


Fig. 1. Flowchart depicting the various steps in the generalized BMS reduction. The full model is first decomposed into interior and boundary partitions (1). The boundary DOF set is preserved using a static condensation of the interior onto the boundary via a set of constraint modes (2). Then, dynamic information is introduced into the interior using a truncated set of fixed-interface modes (3). A local interface reduction is then used to reduce the boundary of the model while maintaining displacement compatibility between opposite unit cell faces (4). Next the contribution of the truncated fixed-interface modes is approximated using residual correction modes that are scaled by ω^2 (5). Finally, Bloch boundary conditions are enforced in order to compute the band-structure diagram. The darkened boxes (steps 4 and 5) highlight the new contributions presented in this paper. The previous steps were proposed earlier in Ref. [14].

including more modes, it is possible to approximate the contribution of the residual modes in order to drive down the error. This process, termed residual-mode enhancement, does not significantly increase computation time, does not increase reduced-order model size, and is capable of reducing error in frequencies by one to two orders of magnitude.

Fig. 1 illustrates the main steps taken in the generalized BMS reduction, the focus of this paper. The free unit-cell model (with no boundary conditions applied) is decomposed into interior and boundary sets. Then modal reductions of the interior and the boundary are performed. Finally the interior and boundary sets are combined as a reduced basis for the model, and Bloch boundary conditions are applied.

The remaining sections are laid out as follows. The next section discusses the calculation of elastic band structure without any model reduction. The section that follows discusses the implementation of the generalized BMS model reduction, including the local interface reduction and the residual-mode enhancement. The final section shows results obtained with the generalized reduced-order model and evaluates the performance of the overall technique.

2. Band-structure computation

The wave propagation characteristics of a periodic medium depend on the configuration of the unit cell. According to Bloch's theorem, waves in crystals are separable into a plane-wave component and a periodic amplitude [3]. The plane-wave component depends on the wave vector, which in turn is related to the frequency via the dispersion relation which we seek to solve for. This decomposition allows the analysis of an infinitely repeating array of unit cells by consideration of just a single unit cell. This section discusses how the unit cell is modeled, how boundary conditions must be applied in order to enforce a Bloch solution, and finally how the band-structure diagram is obtained. For simplicity, the formulation in this paper is described using a 2D model with 2 directions of periodicity (2D2P). The concepts, however, are easily adapted to 1D, 2D, or 3D models and 1, 2, or 3 directions of periodicity.

2.1. Unit-cell model

The fundamental equation governing linear waves in an elastic medium expressed in tensorial notation is

$$\nabla \cdot \mathbf{C} : \boldsymbol{\epsilon} = \rho \ddot{\mathbf{u}}, \tag{1}$$

where \mathbf{C} is the elasticity tensor, ρ is the density, \mathbf{u} is the displacement field, and the strain, $\boldsymbol{\epsilon}$, is given by

$$\boldsymbol{\epsilon} = \frac{1}{2} (\nabla \mathbf{u} + (\nabla \mathbf{u})^T). \tag{2}$$

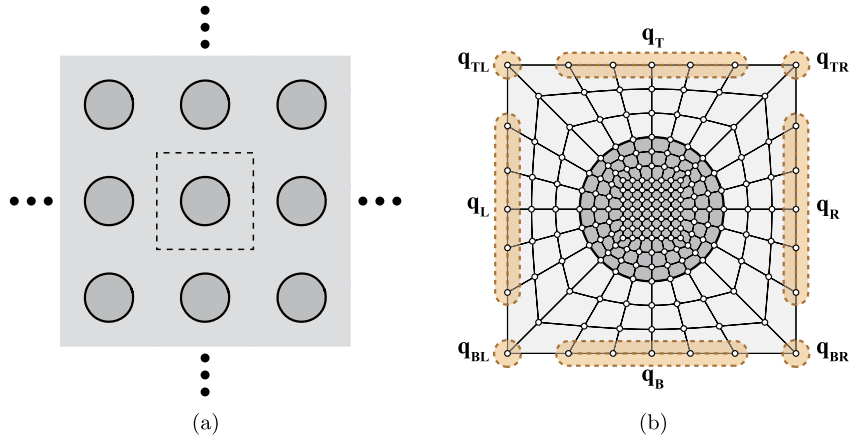


Fig. 2. (a) Periodic 2D material, and (b) isolated unit-cell FE mesh with boundary sets labeled.

The material parameters \mathbf{C} and ρ are known functions of space. Equation (1) thus represents a partial differential equation in the vector field \mathbf{u} . A spatial discretization is used to approximate the spatially varying quantities, yielding a system of ordinary differential equations (ODEs) that can be solved numerically. The three most common discretizations used for unit-cell analysis are the plane-wave expansion (PWE) method, the finite difference (FD) method, and the FE method. The BMS method relies on a partitioning of the system of equations into internal and boundary DOFs. Thus, a real-space discretization must be used where each DOF corresponds to a physical displacement of a node in the model. Both FE and FD methods produce real-space models, however the PWE method does not. For the remainder of this paper we focus on models that have been discretized via the FE method. An example of a 2D infinitely periodic material is shown in Fig. 2a with a single unit cell outlined. To model wave propagation in this unit cell, we begin by cutting the unit cell out of the periodic material so it is an isolated structure with free boundaries. An FE mesh for this free structure is shown in Fig. 2b. Neglecting damping, the resulting system of ODEs is expressed as follows in matrix notation,

$$\mathbf{K}\mathbf{q} + \mathbf{M}\ddot{\mathbf{q}} = \mathbf{0}, \quad (3)$$

where \mathbf{K} and \mathbf{M} are respectively the free stiffness and free mass matrices, and \mathbf{q} is the free displacement vector. Assuming time-harmonic displacements, the equation becomes,

$$(\mathbf{K} - \omega^2\mathbf{M})\mathbf{q} = \mathbf{0}, \quad (4)$$

where ω is the frequency of oscillation. An eigenvalue solution of Eq. (4) would yield the natural frequencies of the free structure, however we would like the isolated cell to behave as though it was still surrounded by an infinite periodic array of identical cells. To obtain this behavior within the isolated cell, we use Bloch boundary conditions.

2.2. Bloch boundary conditions

As mentioned earlier, Bloch's theorem governs waves traveling in periodic materials. For elastic waves this is expressed as follows,

$$\mathbf{u}(\mathbf{x}, \mathbf{k}) = \underbrace{\tilde{\mathbf{u}}(\mathbf{x}, \mathbf{k})}_{\text{periodic function}} \underbrace{e^{i\mathbf{k}^T\mathbf{x}}}_{\text{plane wave}}, \quad (5)$$

where $\mathbf{x} = [x \ y]^T$ is the position vector, and $\mathbf{k} = [k_x \ k_y]^T$ is the wave vector. The periodic function $\tilde{\mathbf{u}}$ must further satisfy periodicity constraints,

$$\tilde{\mathbf{u}}(\mathbf{x}, \mathbf{k}) = \tilde{\mathbf{u}}(\mathbf{x} + \mathbf{r}, \mathbf{k}), \quad \mathbf{r} = a_1\mathbf{r}_1 + a_2\mathbf{r}_2, \quad (6)$$

where \mathbf{r} is a lattice vector of the periodic material, \mathbf{r}_1 and \mathbf{r}_2 are primitive translation vectors of the lattice, and a_1 and a_2 are integers. Combining Eqs. (5) and (6), Bloch's theorem is expressed as a relationship between the boundaries of the unit cell,

$$\mathbf{u}(\mathbf{x} + \mathbf{r}, \mathbf{k}) = \mathbf{u}(\mathbf{x}, \mathbf{k})e^{i\mathbf{k}^T\mathbf{r}}. \quad (7)$$

For a discrete model like the one in Fig. 2b, Eq. (7) becomes a set of equations relating the DOFs on the boundary of the cell. The DOFs on the right boundary are related to the DOFs on the left boundary, and the DOFs on the top boundary are related to the DOFs on the bottom boundary,

$$\mathbf{q}_R = \mathbf{q}_L \lambda_1, \quad \mathbf{q}_T = \mathbf{q}_B \lambda_2, \tag{8}$$

where $\lambda_1 = e^{ik^T r_1}$, and $\lambda_2 = e^{ik^T r_2}$. A similar set of equations is used to link the corner DOFs,

$$\mathbf{q}_{BR} = \mathbf{q}_{BL} \lambda_1, \quad \mathbf{q}_{TR} = \mathbf{q}_{BL} \lambda_1 \lambda_2, \quad \mathbf{q}_{TL} = \mathbf{q}_{BL} \lambda_2. \tag{9}$$

These Bloch boundary equations relate the free DOF vector, \mathbf{q} , to the periodic DOF vector, $\tilde{\mathbf{q}}$,

$$\begin{pmatrix} \mathbf{q}_I \\ \mathbf{q}_L \\ \mathbf{q}_R \\ \mathbf{q}_B \\ \mathbf{q}_T \\ \mathbf{q}_{BL} \\ \mathbf{q}_{BR} \\ \mathbf{q}_{TR} \\ \mathbf{q}_{TL} \end{pmatrix} = \begin{pmatrix} \mathbf{I} & \mathbf{0} & \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \mathbf{I} & \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \lambda_1 \mathbf{I} & \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & \mathbf{I} & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & \lambda_2 \mathbf{I} & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{I} \\ \mathbf{0} & \mathbf{0} & \mathbf{0} & \lambda_1 \mathbf{I} \\ \mathbf{0} & \mathbf{0} & \mathbf{0} & \lambda_1 \lambda_2 \mathbf{I} \\ \mathbf{0} & \mathbf{0} & \mathbf{0} & \lambda_2 \mathbf{I} \end{pmatrix} \begin{pmatrix} \mathbf{q}_I \\ \mathbf{q}_L \\ \mathbf{q}_B \\ \mathbf{q}_{BL} \end{pmatrix}. \tag{10}$$

The Bloch periodicity matrix, $\mathbf{P}(\mathbf{k})$ is then decomposed into terms multiplying λ_1 and λ_2 ,

$$\mathbf{P}(\mathbf{k}) = \mathbf{P}_0 + \lambda_1 \mathbf{P}_1 + \lambda_2 \mathbf{P}_2 + \lambda_1 \lambda_2 \mathbf{P}_{12}. \tag{11}$$

The Bloch-periodicity sub-matrices, \mathbf{P}_0 , \mathbf{P}_1 , \mathbf{P}_2 , and \mathbf{P}_{12} , are formed by placing ones where entries in the free DOF vector should link together with entries in the periodic DOF vector. Note that the DOF vectors are not usually sorted as shown in Eq. (10). In order for the Bloch periodicity matrix to be compatible with the mass and stiffness matrices, they must have the same DOF sorting. This can be accomplished by re-sorting the rows and columns of the mass and stiffness matrices or by re-sorting the rows of the periodicity transformation. To simplify notation, the explicit dependence of \mathbf{P} on the wave vector, \mathbf{k} , is omitted from here on.

2.3. $\omega(\mathbf{k})$ solution

The $\omega(\mathbf{k})$ formulation involves stepping through the Brillouin zone and solving for the frequencies of wave propagation at each wave vector of interest. The steps involved in doing so are very simple. First we form \mathbf{P} for the current wave vector using Eq. (11). Next we substitute Eq. (10) into (4) and premultiply by the hermitian transpose of the Bloch periodicity matrix, \mathbf{P}^\dagger , to obtain the Bloch equations of motion.

$$\mathbf{P}^\dagger (\mathbf{K} - \omega^2 \mathbf{M}) \mathbf{P} \tilde{\mathbf{q}} = \mathbf{0}. \tag{12}$$

We now solve this eigenvalue problem for the frequencies and mode shapes and repeat for every wave vector until we have computed the full band structure.

2.3.1. Eigenvalue algorithm selection

Since the focus of this paper is computational efficiency, it is important to carefully select the eigenvalue algorithms used in obtaining the dispersion. Eigenvalue solution algorithms can be categorized into two classes: direct solvers and iterative solvers. Direct solvers involve processing an entire matrix to diagonalize it, generating all of the eigenvalues and sometimes eigenvectors of the system. The leading-order complexity of direct solvers is typically $\mathcal{O}(n^3)$. For relatively small matrices (less than about 1000×1000 entries), direct methods are very fast, but as n grows, the cubic-order computational complexity makes these algorithms impractical.

Iterative solvers start with a small set of random vectors and iteratively refine the vectors until they converge to a set of eigenvectors. The computationally limiting step in every iteration is a set of matrix–vector products and/or system solutions. This results in iterative methods having a leading-order complexity of $\mathcal{O}(n^2)$ for full matrices and $\mathcal{O}(n)$ for sparse matrices. As mentioned earlier, both FE and FD methods produce sparse matrices. Although iterative methods only give a few eigenvalue/eigenvector pairs, most applications only utilize the lowest-frequency branches for the band structure. As we move to large matrices (more than $10,000 \times 10,000$ entries) it is imperative both to use an iterative solver and to take advantage of the system sparsity. Previously reported results in Ref. [14], fail to take advantage of the sparsity. This gives timing results that unfairly favor BMS because the full-model calculations benefit more dramatically from sparse matrix calculations. In this paper, the full-model calculations are implemented using an iterative solver.

2.4. $\mathbf{k}(\omega)$ solution

In some situations, it may be desirable to solve the wave propagation problem for the wave vector in terms of the frequency. This $\mathbf{k}(\omega)$ problem involves stepping through the frequencies of interest and solving for the wave vectors. The $\mathbf{k}(\omega)$ method produces both real and complex wave-vector solutions. Complex wave-vector solutions are necessary for certain types of analysis because they contain information about the attenuation of waves in space.

To solve the $\mathbf{k}(\omega)$ problem, we begin by setting $\mathbf{D}(\omega) = \mathbf{K} - \omega^2 \mathbf{M}$ and expanding Eq. (12) in terms of the Bloch-periodicity sub matrices:

$$\left[\tilde{\mathbf{D}}_1 + \lambda_1 \tilde{\mathbf{D}}_2 + \lambda_1^* \tilde{\mathbf{D}}_2^T + \lambda_2 \tilde{\mathbf{D}}_3 + \lambda_2^* \tilde{\mathbf{D}}_3^T + \lambda_1 \lambda_2^* \tilde{\mathbf{D}}_4 + \lambda_1^* \lambda_2 \tilde{\mathbf{D}}_4^T + \lambda_1 \lambda_2 \tilde{\mathbf{D}}_5 + \lambda_1^* \lambda_2^* \tilde{\mathbf{D}}_5^T \right] \tilde{\mathbf{q}} = \mathbf{0}, \quad (13)$$

where

$$\tilde{\mathbf{D}}_1 = \mathbf{P}_0^T \mathbf{D}(\omega) \mathbf{P}_0 + \mathbf{P}_1^T \mathbf{D}(\omega) \mathbf{P}_1 + \mathbf{P}_2^T \mathbf{D}(\omega) \mathbf{P}_2 + \mathbf{P}_{12}^T \mathbf{D}(\omega) \mathbf{P}_{12}, \quad (14a)$$

$$\tilde{\mathbf{D}}_2 = \mathbf{P}_0^T \mathbf{D}(\omega) \mathbf{P}_1 + \mathbf{P}_2^T \mathbf{D}(\omega) \mathbf{P}_{12}, \quad (14b)$$

$$\tilde{\mathbf{D}}_3 = \mathbf{P}_0^T \mathbf{D}(\omega) \mathbf{P}_2 + \mathbf{P}_1^T \mathbf{D}(\omega) \mathbf{P}_{12}, \quad (14c)$$

$$\tilde{\mathbf{D}}_4 = \mathbf{P}_1^T \mathbf{D}(\omega) \mathbf{P}_2, \quad (14d)$$

$$\tilde{\mathbf{D}}_5 = \mathbf{P}_0^T \mathbf{D}(\omega) \mathbf{P}_{12}. \quad (14e)$$

The direction of wave propagation is given by the ratio of propagation constants $\mu_1 = \mathbf{k}^T \mathbf{r}_1$ and $\mu_2 = \mathbf{k}^T \mathbf{r}_2$. If the ratio between μ_1 and μ_2 is rational, then we can set $\mu_1 = \sigma b$, and $\mu_2 = \sigma c$, where b and c are integers. Setting $\gamma = e^{i\sigma}$ gives,

$$\lambda_1 = \gamma^b, \quad \lambda_1^* = \gamma^{-b}, \quad \lambda_2 = \gamma^c, \quad \lambda_2^* = \gamma^{-c}. \quad (15)$$

Substituting Eq. (15) into (13) and multiplying by γ^{b+c} gives the polynomial eigenvalue equation of order $2b + 2c$,

$$\left[\gamma^{b+c} \tilde{\mathbf{D}}_1 + \gamma^{2b+c} \tilde{\mathbf{D}}_2 + \gamma^c \tilde{\mathbf{D}}_2^T + \gamma^{2c+b} \tilde{\mathbf{D}}_3 + \gamma^b \tilde{\mathbf{D}}_3^T + \gamma^{2b} \tilde{\mathbf{D}}_4 + \gamma^{2c} \tilde{\mathbf{D}}_4^T + \gamma^{2b+2c} \tilde{\mathbf{D}}_5 + \tilde{\mathbf{D}}_5^T \right] \tilde{\mathbf{q}} = \mathbf{0}. \quad (16)$$

Given an m th order polynomial eigenvalue problem of the form

$$\left[\sum_{j=0}^m \mathbf{A}_j \gamma^j \right] \mathbf{q} = \mathbf{0}, \quad (17)$$

a linearization process reformulates the problem as a general linear eigenvalue problem,

$$\left(\left[\begin{array}{c|c} \mathbf{A}_0 & \mathbf{I} \\ \hline & \ddots \\ & \mathbf{I} \end{array} \right] - \gamma \left[\begin{array}{cccc} -\mathbf{A}_1 & \cdots & -\mathbf{A}_{m-1} & -\mathbf{A}_m \\ \mathbf{I} & & & \\ & \ddots & & \\ & & \mathbf{I} & \mathbf{0} \end{array} \right] \right) \left\{ \begin{array}{c} \mathbf{q} \\ \gamma \mathbf{q} \\ \vdots \\ \gamma^{m-1} \mathbf{q} \end{array} \right\} = \mathbf{0}. \quad (18)$$

The solution of Eq. (18) is difficult because matrix sizes in the linearized eigenproblem are $(nm \times nm)$, thus severely limiting the size of the unit-cell models that can be considered.

If the ratio of propagation constants is not rational, then we can still proceed to Eq. (13), but b and c are no longer integers, so it is not possible to use a polynomial eigenvalue approach. Rather, a root finding algorithm such as Newton's method must be used to solve for γ . This type of approach is typically very time consuming, especially for large models, and requires a reasonably good guess for the solutions in order for the method to converge.

Since the computational cost of performing a direct solution of Eq. (18) can be very large, it is common to perform a dynamic condensation of $\mathbf{D}(\omega)$ before forming the polynomial eigenvalue problem. The dynamic condensation removes all interior DOFs from the system without any loss in accuracy. We first partition Eq. (4) into interior and boundary equations,

$$\overbrace{\left[\begin{array}{cc} \mathbf{D}_{II} & \mathbf{D}_{IA} \\ \mathbf{D}_{AI} & \mathbf{D}_{AA} \end{array} \right]}^{\mathbf{D}=\mathbf{K}-\omega^2 \mathbf{M}} \overbrace{\left\{ \begin{array}{c} \mathbf{q}_I \\ \mathbf{q}_A \end{array} \right\}}^{\mathbf{q}} = \left\{ \begin{array}{c} \mathbf{0} \\ \mathbf{0} \end{array} \right\}, \quad (19)$$

where the subscript I refers to interior DOFs and the subscript A refers to interface DOFs. Thus \mathbf{q}_A is a collection of all of the boundary DOF sets shown in Fig. 2b. We solve the upper half of this equation to obtain an expression for \mathbf{q}_I ,

$$\mathbf{q}_I = -\mathbf{D}_{II}^{-1} \mathbf{D}_{IA} \mathbf{q}_A. \quad (20)$$

This allows us to describe the entire free DOF vector in terms of just the boundary DOFs,

$$\begin{Bmatrix} \mathbf{q}_I \\ \mathbf{q}_A \end{Bmatrix} = \begin{matrix} \mathbf{w}(\omega) \\ \left[\begin{array}{c} -\mathbf{D}_{II}^{-1} \mathbf{D}_{IA} \\ \mathbf{I} \end{array} \right] \end{matrix} \mathbf{q}_A. \tag{21}$$

We now dynamically reduce $\mathbf{D}(\omega)$ using $\mathbf{W}(\omega)$,

$$\mathcal{D}(\omega) = \mathbf{W}(\omega)^T \mathbf{D}(\omega) \mathbf{W}(\omega). \tag{22}$$

Returning to Eq. (14), we use $\mathcal{D}(\omega)$ instead of $\mathbf{D}(\omega)$. Note that since we have removed all interior DOFs from the system, \mathbf{P} must be modified slightly. Otherwise we simply proceed as before. The obvious benefit of using a dynamic reduction is that the size of the eigenvalue problem to be solved is greatly reduced without any sacrifice in accuracy. The main cost of the dynamic reduction is the system solution of Eq. (20) which must be computed at every frequency. Efficient algorithms for solving large sparse systems with multiple right-hand sides exist, but this still represents a large part of the computational cost.

For completeness, we also highlight an alternative approach to obtaining the $\mathbf{k}(\omega)$ band structure. The FE mass and stiffness matrices may be formed using a Bloch-operator approach wherein the element formulation explicitly considers a Bloch solution [12]. When formed in this way, the stiffness matrix has the following form:

$$\tilde{\mathbf{K}} = \tilde{\mathbf{K}}_0 + \tilde{\mathbf{K}}_x k_x + \tilde{\mathbf{K}}_y k_y + \tilde{\mathbf{K}}_{xx} k_x^2 + \tilde{\mathbf{K}}_{xy} k_x k_y + \tilde{\mathbf{K}}_{yy} k_y^2, \tag{23}$$

and the mass matrix has no dependence on wave vector. This form of the system matrices leads to a quadratic eigenvalue problem for any direction of wave propagation [19]. This is a significant advantage over the previously discussed $\mathbf{k}(\omega)$ method based on Bloch boundary conditions that gives higher order polynomial eigenvalue problems depending on the direction of wave propagation. Unfortunately the FE matrices obtained with the Bloch-operator approach do not always converge as quickly as traditional FE matrices. This means that finer meshes are necessary to obtain equivalent accuracy. Nevertheless, the ability to easily consider any direction of wave propagation makes this approach superior for obtaining $\mathbf{k}(\omega)$ band-structure surfaces over the entire Brillouin zone. We will not discuss the Bloch operator FE method further here because the BMS formulation cannot be extended to matrices obtained in this manner.

3. Generalized Bloch mode synthesis model reduction

BMS is a model reduction technique introduced in Ref. [14]. It resembles CMS, a substructuring technique widely used for the vibration analysis of structures. The main idea of BMS is to split the model into interior and boundary DOFs. A small set of modes is then used to replace the large set of interior DOFs. The power of this approach is that the boundary DOFs are not affected by the reduction so Bloch boundary conditions are still applied exactly as before. This means that a reduced model can be formed *a-priori*, and then used throughout the Brillouin zone.

Reference [14] showed that a reduction of the boundary DOFs is possible using the S-CC reduction. The effectiveness of this boundary reduction is limited because it must be performed *after* applying Bloch boundary conditions. This means that a new S-CC reduction must be computed and applied for every point in the Brillouin zone. We show here that a local interface reduction [20] can be used to reduce the boundary in such a way that it must only be performed once, independently of the Brillouin zone.

3.1. Interior modal reduction

We begin by highlighting the main steps involved in the Hurty/Craig–Bampton (HCB) reduction methodology [21,22]. This is a CMS substructure representation that is very popular for its good performance and simplicity. Another advantage of the HCB method is that it lends itself very effectively to a secondary reduction of the boundary DOFs. The method starts by partitioning Eq. (4) into interior and boundary equations,

$$\left(\begin{bmatrix} \mathbf{K}_{II} & \mathbf{K}_{IA} \\ \mathbf{K}_{AI} & \mathbf{K}_{AA} \end{bmatrix} - \omega^2 \begin{bmatrix} \mathbf{M}_{II} & \mathbf{M}_{IA} \\ \mathbf{M}_{AI} & \mathbf{M}_{AA} \end{bmatrix} \right) \begin{Bmatrix} \mathbf{q}_I \\ \mathbf{q}_A \end{Bmatrix} = \mathbf{0}. \tag{24}$$

Next we compute a set of normal mode shapes, Φ_I , to describe the interior,

$$\left(\mathbf{K}_{II} - \omega_1^2 \mathbf{M}_{II} \right) \phi_I = \mathbf{0}, \tag{25a}$$

$$\Phi_{FI} = \left[\{ \phi_I \}_1 \quad \{ \phi_I \}_2 \quad \cdots \quad \{ \phi_I \}_{n_1^\phi} \right], \tag{25b}$$

where the number of interior mode shapes n_1^ϕ is much smaller than the original interior dimension, n_I . The interior mode shapes are referred to as fixed-interface mode shapes because they are the structural mode shapes that would result from

clamping the boundaries of the free unit-cell structure. Fig. A.1 shows a selection of fixed-interface modes for a simple unit cell with a circular inclusion.

Next, we compute a set of static constraint modes to describe the boundary DOFs,

$$\Psi = -\mathbf{K}_{\text{II}}^{-1} \mathbf{K}_{\text{IA}}. \quad (26)$$

Fig. A.2 shows a selection of constraint modes for the circular-inclusion material. We now form the HCB transformation matrix by expressing the original DOF vector as a linear combination of fixed-interface modes and constraint modes,

$$\begin{Bmatrix} \mathbf{q}_I \\ \mathbf{q}_A \end{Bmatrix} = \overbrace{\begin{bmatrix} \Phi_{\text{FI}} & \Psi \\ \mathbf{0} & \mathbf{I} \end{bmatrix}}^{\mathbf{T}_{\text{HCB}}} \begin{Bmatrix} \eta_I \\ \mathbf{q}_A \end{Bmatrix}, \quad (27)$$

where η_I is a set of modal DOFs corresponding to the set of fixed-interface modes. The constraint modes, Ψ , capture the influence of interface motion on the interior. Note that the interface DOF set remains unchanged in this transformation.

We now introduce the generalized BMS transformation, \mathbf{B} , which denotes a transformation from the free equations of motion to a reduced set of equations (i.e. the free BMS equations of motion) that retains the ability to accept periodic boundary conditions. In other words, the generalized BMS transformation is independent of Brillouin-zone location. If we form the free BMS mass and stiffness matrices,

$$\mathcal{M} = \mathbf{B}^T \mathbf{M} \mathbf{B}, \quad \mathcal{K} = \mathbf{B}^T \mathbf{K} \mathbf{B}. \quad (28)$$

we can still apply Bloch boundary conditions. Then to obtain the band structure we simply make the substitutions $\mathbf{M} \rightarrow \mathcal{M}$ and $\mathbf{K} \rightarrow \mathcal{K}$ in Eq. (12) or (13), and update \mathbf{P} to match the reduced system DOFs.

If we want to perform a BMS reduction using the HCB transformation then we simply set $\mathbf{B} = \mathbf{T}_{\text{HCB}}$. This is done in Ref. [14] to reduce the system of equations and quickly step through all wave vectors of interest. Note that this reference also included a boundary reduction based on S-CC modes, but this reduction would not fit into our definition of a generalized BMS transformation because the boundary reduction is dependent on Brillouin-zone location.

The HCB transformation is excellent for demonstration purposes, but we can use any CMS transformation for \mathbf{B} . For example, another popular technique referred to in the literature as the Craig–Chang method [23] uses free-interface modes instead of fixed-interface modes, and attachment modes instead of constraint modes. Different substructure representations may change how Bloch boundary conditions are applied.

In this paper we do not consider damping, however CMS methods have been extended to include certain forms of damping [24,25]. The extension of BMS to include damping should thus follow the same procedure. Depending on the form of damping however, there may not be a straightforward way to apply BMS.

3.2. Local interface modal reduction

Reference [14] reduces the boundary DOFs via the S-CC reduction [18,26]. This reduction is designed to reduce the interface DOFs after they have been coupled to other interface DOFs. In other words, this means that a new reduction must be performed for every different wave vector, which results in high computational costs. Furthermore, this approach prevents the reduction from being applicable to the $\mathbf{k}(\omega)$ problem. We would like to perform just a single reduction of the boundary that is independent of the Brillouin zone. In order to do so, we need to be able to reduce the boundary DOFs before coupling the interface DOFs together. This is possible using L-CC interface reduction [20].

We begin by partitioning the reduced equations of motion,

$$\left(\begin{bmatrix} \mathcal{K}_{\text{II}} & \mathcal{K}_{\text{IA}} \\ \mathcal{K}_{\text{AI}} & \mathcal{K}_{\text{AA}} \end{bmatrix} - \omega^2 \begin{bmatrix} \mathcal{M}_{\text{II}} & \mathcal{M}_{\text{IA}} \\ \mathcal{M}_{\text{AI}} & \mathcal{M}_{\text{AA}} \end{bmatrix} \right) \begin{Bmatrix} \eta_I \\ \mathbf{q}_A \end{Bmatrix} = \mathbf{0}. \quad (29)$$

Next we compute a set of normal mode shapes, Φ_A , to describe the boundary motion,

$$\left(\mathcal{K}_{\text{AA}} - \omega_A^2 \mathcal{M}_{\text{AA}} \right) \phi_A = \mathbf{0}, \quad (30a)$$

$$\Phi_A = \left[\{\phi_A\}_1 \quad \{\phi_A\}_2 \quad \cdots \quad \{\phi_A\}_{n_A} \right]. \quad (30b)$$

An example of such a boundary mode shape is shown in Fig. 3a. We wish to perform a modal reduction using the boundary modes and then apply Bloch boundary conditions. We begin by partitioning the modes according to the boundary sets in Fig. 2b,

$$\Phi_A = \left\{ \Phi_{\text{L}}^T \quad \Phi_{\text{R}}^T \quad \Phi_{\text{B}}^T \quad \Phi_{\text{T}}^T \quad \Phi_{\text{BL}}^T \quad \Phi_{\text{BR}}^T \quad \Phi_{\text{TR}}^T \quad \Phi_{\text{TL}}^T \right\}^T. \quad (31)$$

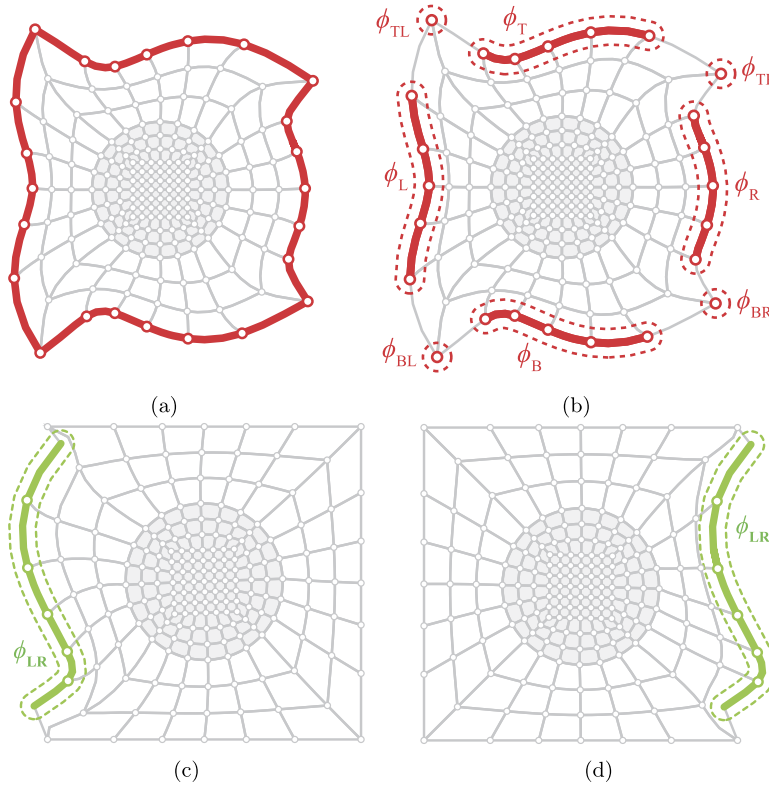


Fig. 3. (a) Full boundary mode, (b) partitioned boundary mode, and compatible boundary modes for (c) left and (d) right boundaries. Although boundary modes only contain boundary DOFs, they define a linear combination of constraint modes which can be used to compute deflection in the unit-cell interior. The modes shown here include this interior deflection in light gray.

Fig. 3b shows an example of a mode shape partitioned into boundary sets. Suppose we simply reduced each boundary segment by the corresponding mode partition. To illustrate the problems with this approach, we focus for the moment on just the right and left boundary sets. The modal reduction will take the following form,

$$\mathbf{q}_L = \Phi_L \boldsymbol{\eta}_L, \quad \mathbf{q}_R = \Phi_R \boldsymbol{\eta}_R. \tag{32}$$

We must now find a way to apply Bloch boundary conditions, so we substitute Eq. (32) into Eq. (8),

$$\Phi_L \boldsymbol{\eta}_L = \Phi_R \boldsymbol{\eta}_R \lambda_x. \tag{33}$$

Next we must solve for $\boldsymbol{\eta}_R$, but this is not possible because Φ_R generally does not span the same column space as Φ_L . To circumvent this issue, we use a single modal basis to reduce both the right and left boundary sets, pursuant to [20],

$$\mathbf{q}_L = \Phi_{LR} \boldsymbol{\eta}_L, \quad \mathbf{q}_R = \Phi_{LR} \boldsymbol{\eta}_R. \tag{34}$$

Now, Eq. (33) becomes,

$$\Phi_{LR} \boldsymbol{\eta}_R = \Phi_{LR} \boldsymbol{\eta}_L \lambda_x, \tag{35}$$

which allows us to express the dependent modal coordinates, $\boldsymbol{\eta}_R$, in terms of the independent modal coordinates, $\boldsymbol{\eta}_L$,

$$\boldsymbol{\eta}_R = \boldsymbol{\eta}_L \lambda_x. \tag{36}$$

The only remaining question is how to obtain the single modal basis, Φ_{LR} so that it contains the important dynamic information from both the left and right boundary mode sets, Φ_L , and Φ_R . This is accomplished by simply combining the original left and right mode sets, $\Phi_{LR} = [\Phi_L \ \Phi_R]$. To ensure that Φ_{LR} is not singular, we also orthogonalize the combined set of mode shapes either via QR decomposition or singular value decomposition. For the examples shown later, we choose the latter. Figs. 3c and 3d show an example mode shape selected from Φ_{LR} which is used to reduce the left and right boundaries. We see that the left and right boundaries are now compatible for connection by Bloch boundary conditions.

the model interior, n_I , and the number of DOFs in the model boundary, n_A . For the purpose of comparing computational complexity, we can relate the interior and boundary sizes to the full-model size, n , as follows,

$$\mathcal{O}(n_I) = \mathcal{O}(n), \tag{41}$$

$$\mathcal{O}(n_A) = \mathcal{O}(n^{1/2}) \text{ in 2D, } \quad \mathcal{O}(n_A) = \mathcal{O}(n^{2/3}) \text{ in 3D.} \tag{42}$$

Assuming that the interior partitions of the stiffness and mass matrices are sparse, and an iterative eigenvalue algorithm is used (e.g., `eigs` in MATLAB), then the fixed-interface mode calculation will be $\mathcal{O}(n)$. Since \mathbf{K}_{II} is positive definite, the system solution performed in Eq. (26) can be computed with the sparse supernodal cholesky algorithm [27] (automatically called by Matlab’s backslash operator). In our tests we observed this solver to have a computational order of approximately $\mathcal{O}(n)$ for a fixed number of right-hand-side vectors. The number of right-hand-side vectors varies however with model size as it is equivalent to the number of boundary DOFs. Thus, we can expect the constraint mode calculation to be $\mathcal{O}(n^{3/2})$ in 2D and $\mathcal{O}(n^{5/3})$ in 3D.

In comparison, the computational order of the full-model band-structure calculation will be $\mathcal{O}(n)$, because we are using an iterative eigenvalue algorithm on a sparse system. This shows that as we move to larger and larger systems, the performance of the BMS method based on the HCB reduction will deteriorate with respect to the full band-structure calculation.

Even though the BMS model has higher computational order than simply computing the full band structure, the BMS model is often faster because the computationally expensive steps are performed just once rather than at every \mathbf{k} point. Thus the computational benefits emerge when the number of \mathbf{k} points is very large. Still, as we move to larger and larger systems, the computational benefit will diminish. It may be possible to lower the computational order of the BMS reduction using automated multilevel substructuring (AMLS) [28,29]. The main idea in AMLS is to split the unit cell up into smaller sub-structures, perform separate reductions on all of the subcomponents, and then combine the subcomponents into a single reduced-order model. If the subcomponents are still too large to be handled efficiently, then the process can be repeated for each subcomponent. This forms a tree-like organization with multiple levels of substructures. In our experience, the AMLS algorithm was not able to reduce the computational cost of the BMS reduction, so the results are not included in this paper. The literature reports dramatic speedups over the HCB method however, therefore the problem may be in our algorithm implementation. This route warrants further research.

3.4. Approximating residual-mode contribution for enhanced accuracy

The HCB method is exact when all modes are kept in the transformation. Of course this is counterproductive because no reduction in model size is obtained. It has been shown, however, that an approximation of the residual-mode contributions can lead to improved accuracy [30]. A recent extension of this technique obtains this improved accuracy without any increase in the reduced-model size [31–33]. This approach was first developed in the context of HCB reductions, but was shown to be useful for AMLS models as well.

As mentioned earlier, the HCB representation is exact when all interior modes are kept. The interior mode set must be truncated however if the model size is to be reduced. Rather than simply truncating the residual-mode set, the contribution of the residual modes can be approximated by modifying the HCB transformation. We begin by adding the residual modes into Eq. (27) so that the transformation represents a change of coordinates but not a model reduction,

$$\begin{Bmatrix} \mathbf{q}_I \\ \mathbf{q}_A \end{Bmatrix} = \begin{bmatrix} \Phi_{FI}^d & \Phi_{FI}^r & \Psi \\ \mathbf{0} & \mathbf{0} & \mathbf{I} \end{bmatrix} \begin{Bmatrix} \eta_I^d \\ \eta_I^r \\ \mathbf{q}_A \end{Bmatrix}, \tag{43}$$

where the superscript *d* denotes dominant modes that are kept in the original HCB representation, and the superscript *r* denotes residual modes that are truncated in the original HCB representation. Adding in the local interface reduction gives,

$$\begin{Bmatrix} \mathbf{q}_I \\ \mathbf{q}_A \end{Bmatrix} = \begin{bmatrix} \eta_I^d & \eta_I^r & \Psi \mathbf{L} \\ \mathbf{0} & \mathbf{0} & \mathbf{L} \end{bmatrix} \begin{Bmatrix} \eta_I^d \\ \eta_I^r \\ \eta_A \end{Bmatrix}. \tag{44}$$

Applying this transformation to Eq. (24) gives,

$$\left(\begin{bmatrix} \Lambda^d & \mathbf{0} & \mathbf{0} \\ & \Lambda^r & \mathbf{0} \\ \text{sym} & & \mathbf{L}^T \mathcal{K}_{AA} \mathbf{L} \end{bmatrix} - \omega^2 \begin{bmatrix} \mathbf{I} & \mathbf{0} & \begin{Bmatrix} \Phi_{FI}^d \\ \Phi_{FI}^r \end{Bmatrix}^T \mathbf{M}_c \\ & \mathbf{I} & \begin{Bmatrix} \Phi_{FI}^d \\ \Phi_{FI}^r \end{Bmatrix}^T \mathbf{M}_c \\ \text{sym} & & \mathbf{L}^T \mathcal{M}_{AA} \mathbf{L} \end{bmatrix} \right) \begin{Bmatrix} \eta_I^d \\ \eta_I^r \\ \eta_A \end{Bmatrix} = \begin{Bmatrix} \mathbf{0} \\ \mathbf{0} \\ \mathbf{0} \end{Bmatrix}, \tag{45}$$

where

$$\mathbf{M}_c = (\mathbf{M}_{IA} + \mathbf{M}_{II} \Psi) \mathbf{L}, \tag{46}$$

and Λ^d and Λ^r respectively contain the dominant and residual eigenvalues on their diagonals. Considering just the second row of Eq. (45), we solve for η_1^r ,

$$\eta_1^r = \left(\Lambda^r - \omega^2 \mathbf{I} \right)^{-1} \left\{ \Phi_{\text{FI}}^r \right\}^T \mathbf{M}_c \eta_A. \quad (47)$$

Substituting this back into the first row of Eq. (44) gives,

$$\mathbf{q}_1 = \Phi_{\text{FI}}^d \eta_1^d + \left(\Psi \mathbf{L} + \omega^2 \mathbf{F}_r \mathbf{M}_c \right) \eta_A, \quad (48)$$

with

$$\mathbf{F}_r = \Phi_{\text{FI}}^r \left(\Lambda^r - \omega^2 \mathbf{I} \right)^{-1} \left\{ \Phi_{\text{FI}}^r \right\}^T, \quad (49)$$

where \mathbf{F}_r is the residual-flexibility matrix. Assuming the residual frequencies are much higher than the highest frequency of interest, ω , we can approximate the residual flexibility as,

$$\mathbf{F}_r \approx \underbrace{\Phi_{\text{FI}}^r \left\{ \Lambda^r \right\}^{-1} \left\{ \Phi_{\text{FI}}^r \right\}^T}_{\mathbf{F}_{rs}} + \omega^2 \underbrace{\Phi_{\text{FI}}^r \left\{ \Lambda^r \right\}^{-2} \left\{ \Phi_{\text{FI}}^r \right\}^T}_{\mathbf{F}_{rm}}. \quad (50)$$

Substituting back into Eq. (48) and keeping only first order terms in ω^2 gives,

$$\mathbf{q}_1 = \Phi_{\text{FI}}^d \eta_1^d + \left(\Psi \mathbf{L} + \omega^2 \mathbf{F}_{rs} \mathbf{M}_c \right) \eta_A. \quad (51)$$

It is not practical to compute the static residual flexibility, \mathbf{F}_{rs} using residual modes, because these modes are not computed in the original HCB representation. Rather, the static residual flexibility can be expressed in terms of the dominant modes as follows,

$$\mathbf{F}_{rs} = \mathbf{K}_{\text{II}}^{-1} - \Phi_{\text{FI}}^d \left\{ \Lambda^d \right\}^{-1} \left\{ \Phi_{\text{FI}}^d \right\}^T. \quad (52)$$

At this point, we rewrite Eq. (44) in terms of η_1^d and η_A ,

$$\underbrace{\begin{Bmatrix} \mathbf{q} \\ \mathbf{q}_1 \\ \mathbf{q}_A \end{Bmatrix}}_{\mathbf{q}} = \underbrace{\left(\begin{array}{c} \mathbf{B} \\ \left[\begin{array}{cc} \Phi_{\text{FI}}^d & \Psi \mathbf{L} \\ \mathbf{0} & \mathbf{L} \end{array} \right] + \omega^2 \left[\begin{array}{cc} \mathbf{0} & \mathbf{F}_{rs} \mathbf{M}_c \\ \mathbf{0} & \mathbf{0} \end{array} \right] \end{array} \right)}_{\mathbf{B}_e} \underbrace{\begin{Bmatrix} \eta \\ \eta_1^d \\ \eta_A \end{Bmatrix}}_{\eta}. \quad (53)$$

The term $\mathbf{F}_{rs} \mathbf{M}_c$ can be thought of as a set of residual-correction modes that are scaled by the frequency to correct for the truncated modes. Fig. A.4 shows a selection of correction modes for the circular-inclusion material. If we wish to calculate the $\mathbf{k}(\omega)$ band structure, then forming \mathbf{B}_e is straightforward because ω is known. For the $\omega(\mathbf{k})$ formulation we need to approximate ω^2 . Since this ω^2 will change as a function of the wave vector, we need to use a slightly different transformation for each wave vector. This goes against our earlier definition of a generalized BMS transformation, however the modification performed at every wave vector is very computationally inexpensive so it does not affect the performance. First, we replace \mathbf{B} with \mathbf{B}_e in Eq. (28),

$$\mathcal{M}_e = \mathbf{B}_e^T \mathbf{M} \mathbf{B}_e, \quad \mathcal{K}_e = \mathbf{B}_e^T \mathbf{K} \mathbf{B}_e. \quad (54)$$

Expanding this product gives,

$$\mathcal{K}_e = \underbrace{\mathbf{B}^T \mathbf{K} \mathbf{B}}_{\mathcal{K}_0} + \omega^2 \left(\underbrace{\mathbf{B}^T \mathbf{K} \mathbf{B}_r}_{\mathcal{K}_2} + \underbrace{\mathbf{B}_r^T \mathbf{K} \mathbf{B}}_{\mathcal{K}_2^T} \right) + \omega^4 \underbrace{\mathbf{B}_r^T \mathbf{K} \mathbf{B}_r}_{\mathcal{K}_4}, \quad (55)$$

$$\mathcal{M}_e = \underbrace{\mathbf{B}^T \mathbf{M} \mathbf{B}}_{\mathcal{M}_0} + \omega^2 \left(\underbrace{\mathbf{B}^T \mathbf{M} \mathbf{B}_r}_{\mathcal{M}_2} + \underbrace{\mathbf{B}_r^T \mathbf{M} \mathbf{B}}_{\mathcal{M}_2^T} \right) + \omega^4 \underbrace{\mathbf{B}_r^T \mathbf{M} \mathbf{B}_r}_{\mathcal{M}_4}. \quad (56)$$

We now pre- and postmultiply \mathcal{K}_e and \mathcal{M}_e by \mathbf{P} to enforce Bloch boundary conditions,

$$\tilde{\mathcal{K}}_e = \tilde{\mathcal{K}}_0 + \omega^2 \left(\tilde{\mathcal{K}}_2 + \tilde{\mathcal{K}}_2^T \right) + \omega^4 \tilde{\mathcal{K}}_4, \quad (57)$$

$$\tilde{\mathcal{M}}_e = \tilde{\mathcal{M}}_0 + \omega^2 (\tilde{\mathcal{M}}_2 + \tilde{\mathcal{M}}_2^T) + \omega^4 \tilde{\mathcal{M}}_4, \tag{58}$$

where the tilde accent denotes premultiplication by \mathbf{P}^\dagger and postmultiplication by \mathbf{P} (e.g., $\tilde{\mathcal{M}} = \mathbf{P}^\dagger \mathcal{M} \mathbf{P}$).

At this point a $\mathbf{k}(\omega)$ solution can be carried out because the ω is a known quantity at the outset of every calculation. In order to perform a $\omega(\mathbf{k})$ calculation however, the dependence of $\tilde{\mathcal{K}}_e$ and $\tilde{\mathcal{M}}_e$ on ω^2 must be removed. Consider briefly the reduced-order characteristic equation with no enhancement terms, after enforcement of periodic boundary conditions,

$$(\tilde{\mathcal{K}}_0 - \omega^2 \tilde{\mathcal{M}}_0) \tilde{\eta} = \mathbf{0}. \tag{59}$$

This can be rearranged and solved to give,

$$\omega^2 \tilde{\eta} = \underbrace{(\tilde{\mathcal{M}}_0)^{-1} \tilde{\mathcal{K}}_0}_{\mathbf{Y}} \tilde{\eta}. \tag{60}$$

This substitution is referred to in Refs. [31,32] as O’Callahan’s approach and further details may be found there. The substitution shown in Eq. (60) does not directly provide an approximation of ω^2 . Clearly, the matrix \mathbf{Y} cannot be viewed as an approximation of the scalar ω^2 . Rather the effect that ω^2 has on the modal DOF vector, η , is being approximated by \mathbf{Y} . Also note that \mathbf{Y} is formed after applying periodic boundary conditions, so this must be carried out at every \mathbf{k} point. Nevertheless, obtaining \mathbf{Y} is computationally inexpensive because $\tilde{\mathcal{K}}_0$ and $\tilde{\mathcal{M}}_0$ are small relative to the full model. Substituting the previous result and the Bloch periodicity equation into Eq. (53) gives,

$$\mathbf{q} = (\mathbf{B} + \omega^2 \mathbf{B}_r) \mathbf{P} \tilde{\eta} = \underbrace{(\mathbf{B} \mathbf{P} + \mathbf{B}_r \mathbf{P} \mathbf{Y})}_{\mathbf{H}} \tilde{\eta}. \tag{61}$$

Finally, the residual-mode enhanced matrices are obtained by pre- and postmultiplying the free mass and stiffness matrices by the composite transformation \mathbf{H} ,

$$\tilde{\mathcal{K}}_e = \mathbf{H}^T \mathbf{K} \mathbf{H} = \tilde{\mathcal{K}}_0 + (\tilde{\mathcal{K}}_2 \mathbf{Y} + \mathbf{Y}^T \tilde{\mathcal{K}}_2^T) + \mathbf{Y}^T \tilde{\mathcal{K}}_4 \mathbf{Y}, \tag{62}$$

$$\tilde{\mathcal{M}}_e = \mathbf{H}^T \mathbf{M} \mathbf{H} = \tilde{\mathcal{M}}_0 + (\tilde{\mathcal{M}}_2 \mathbf{Y} + \mathbf{Y}^T \tilde{\mathcal{M}}_2^T) + \mathbf{Y}^T \tilde{\mathcal{M}}_4 \mathbf{Y}. \tag{63}$$

Finally, we obtain the band-structure frequencies by solving the residual-enhanced BMS eigenvalue problem,

$$(\tilde{\mathcal{K}}_e - \omega^2 \tilde{\mathcal{M}}_e) \tilde{\eta} = \mathbf{0}. \tag{64}$$

The approximation of the residual-mode contribution allows for significant improvements in accuracy. This allows band-structure error to be decreased with just a small trade-off in additional computation time. Note that the residual-mode correction provides an enhancement to the interior modes but not to the interface modes. A similar enhancement for local interface modal reduction may be possible but has not yet been developed. In the present work, when the residual-enhanced technique is used to improve the accuracy of the HCB reduction of the interior, then a larger number of boundary modes most likely needs to be used so that the error in the boundary-mode reduction does not dominate the overall error.

3.5. Mode selection

Up until this point we have not addressed how the modes used in the generalized BMS reduction should be selected. The simplest strategy is to use the lowest frequency modes. Although the lowest-frequency normal modes are not guaranteed to be the optimal modes for representing the low frequency system modes, this heuristic approach to mode selection is widely used because it provides good approximations for most problems, and is very simple to implement because the lowest-frequency normal modes are easily obtained using an iterative eigenvalue solver. The number of lowest-frequency modes kept can be varied to reach the desired model size and/or solution accuracy.

An alternative to frequency-based mode selection is to rank modes by their effective interface mass [34]. This metric quantifies the forces generated at the boundary due to each normal mode, and selects modes that show large forces because these will have the largest impact on the system.

In the numerical results presented in this paper, we use frequency-based mode selection both for the interior and boundary modes. We specify a number of modes for the interior, n_1^ϕ , and also the number of modes for the boundary, n_A^ϕ . For the performance plots, the number of boundary modes used is related to the number of interior modes by a constant multiplier, p :

$$n_A^\phi = p \cdot n_1^\phi. \tag{65}$$

For a given error tolerance, we can find an appropriate number of modes to use by successively increasing n_1^ϕ and recomputing a small section of the reduced-model band structure until the change between subsequent solutions is small.

Table 1
Abbreviations for different generalized BMS schemes.

Abbreviation	Interior reduction type	Residual-mode enhancement	Local interface reduction
BMS _{HCB}	HCB	✗	✗
BMS _{HCB+}	HCB	✓	✗
BMS _{HCB} ^{I-CC}	HCB	✗	✓
BMS _{HCB+} ^{I-CC}	HCB	✓	✓

In practice, it is convenient to perform this evaluation in reverse because modes can be removed from a BMS model by simply deleting the corresponding rows and columns of the mass and stiffness matrices. Once we have n_1^ϕ , we repeat this same process with n_A^ϕ . A few successive iterations between the two sets may be necessary to find the smallest reduced-order model that meets the target error tolerance.

4. Numerical experiments

In this section, we demonstrate some results using two example materials. The first is a 2D material with an 8-lobed inclusion in every unit cell. The second is a 3D locally resonant elastic metamaterial [35] consisting of a heavy spherical resonator coated in a soft silicon rubber embedded in an epoxy matrix. In our numerical experiments, we are concerned with both the speedups obtained relative to the full model and the error in the reduced-model results with respect to the full-model results.

Frequency error, e_f , is defined as the deviation from the nominal frequency divided by the nominal frequency. This definition does not work well when the nominal frequency is zero, so zero-frequency points are ignored in frequency error discussions. Wave-vector error is evaluated by normalizing the complex magnitude of the wave-vector deviation by π/L_x . We define the error between two mode shape vectors, ϕ_A and ϕ_B as,

$$e_\phi = 1 - \frac{|\langle \phi_A | \phi_B \rangle|}{\|\phi_A\| \|\phi_B\|}, \quad (66)$$

where $|\cdot|$ denotes the complex modulus, $\langle \cdot | \cdot \rangle$ denotes the dot-product operator, and $\|\cdot\|$ denotes the 2-norm of a vector. Note that the eigenvectors computed with reduced models are in a reduced modal space and thus do not represent physical coordinates. To compare the full-model mode shape with the reduced-model mode shape, both modes must first be transformed to physical space. This is accomplished using the BMS transformation matrix and the Bloch boundary condition matrix,

$$\phi_{\text{BMS}}^{\text{phys}} = \mathbf{BP}\tilde{\phi}_{\text{BMS}}, \quad (67)$$

$$\phi_{\text{full}}^{\text{phys}} = \mathbf{P}\tilde{\phi}_{\text{full}}. \quad (68)$$

Note that \mathbf{P} should be the periodicity transformation specific to the model that created the mode shape. Thus, \mathbf{P} is not the same between Eqs. (67) and (68).

The performance of the algorithm is estimated by timing both the reduced- and full-model band-structure calculations. The FE modeling time is not included, so the full-model computation has no up-front calculation cost, just a computation time per point (per \mathbf{k} point for $\omega(\mathbf{k})$ calculations or per frequency point for $\mathbf{k}(\omega)$ calculations). When multiplied by the number of points used in the computation, we obtain the total computation time for the full model,

$$t_{\text{full}} = n_p \cdot t_{\text{full}}^p. \quad (69)$$

In addition to the per-point computation time, reduced-model computations have an up-front computation time associated with performing the reduction itself,

$$t_{\text{BMS}} = n_p \cdot t_{\text{BMS}}^p + t_{\text{BMS}}^u. \quad (70)$$

The BMS reduction becomes more worthwhile the larger n_p is. When n_p is very small, it may be faster to simply perform the full-model band-structure calculations. As n_p is increased, a break-even point will be reached after which the BMS reduction becomes increasingly beneficial.

The examples in this section explore different combinations of the interface and interior reduction strategies described earlier. To make it simpler to discuss the various reduction strategies, we use the abbreviations shown in Table 1.

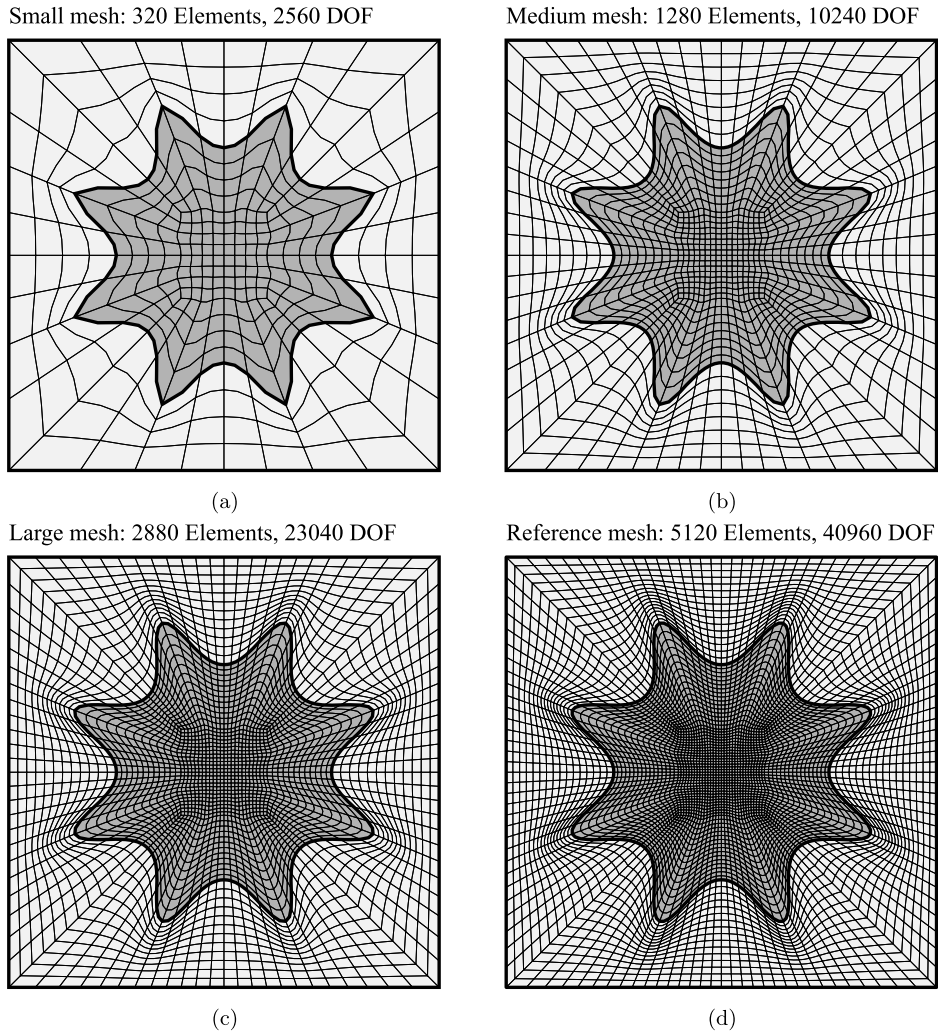


Fig. 4. Various FE meshes of the unit cell of a 2D material with 8-lobed inclusions (resembling the Phononics conference helicoid–catenoid logo [36]).

4.1. Eight-petal inclusion model

The first material to be considered has a 2D unit cell with an 8-lobed inclusion. The inclusion is shaped to resemble the Helicoid–Catenoid logo of the Phononics conference [36]. The unit cell is modeled under the assumption of plane strain using four different meshes as shown in Fig. 4. The elements are 9-node quadrilaterals with Lagrange-polynomial shape functions. We use the varying mesh densities to analyze the impact of model size on the effectiveness of the reduction. The computations for this material are performed on a laptop computer with 2.8 GHz Intel Core i7 processor and 16 GB of memory.

4.1.1. Discretization error

Before going into an in-depth analysis of the error introduced by the model reduction, we should take into account discretization error. Discretization error arises due to the inability of a discrete model to capture the behavior (and sometimes geometry) of a continuous structure. Mesh refinement can be used to decrease the discretization error and get closer to the continuous-model solution. Reduction error, on the other hand, occurs due to the inability of a reduced modal space to capture the exact dynamics of the full mesh. Reduction error decreases as we use more modes in the BMS model.

The overall error in a BMS-reduced model consists of discretization error and reduction error. To explore the importance of discretization error relative to reduction error, we compute the $\omega(\mathbf{k})$ band-structure of the eight-petal inclusion model using each mesh in Fig. 4. The maximum frequency error, e_f^{\max} , in the first ten branches of the band structure is computed for each mesh with respect to the reference mesh and plotted as thick, faint, solid lines in Fig. 5. This approximates the discretization error in each mesh. The solid and dotted lines show the maximum frequency error versus number of interior modes for BMS_{HCB} and BMS_{HCB+} models, respectively. At this stage, no boundary reduction is performed. The number of




Fig. 5. Combined FE discretization and BMS model-reduction error for small, medium, and large meshes of the 8-lobed material with respect to the reference mesh. The BMS models are formed using the HCB, and residual-enhanced HCB reductions of the interior. No boundary reduction is performed in this example.

interior modes is varied for the BMS models by increasing the cutoff frequency from 1 to 10 times the maximum frequency of interest. Note that typically we evaluate a reduced model by comparing to its respective full model, but here we compare each reduced model with the reference mesh so that we can see both discretization and reduction error.

The discretization errors seen in Fig. 5 are relatively large. Even for the large mesh with 23,040 DOFs, the error is barely below 0.1%. Better results can be obtained with more sophisticated (higher order) FEs, but our focus here is not on FE convergence. Rather, we simply stress that there is not much benefit in creating a very accurate reduced-order model if the discretization error is comparatively large. We also observe that the error due to the BMS reductions drops quickly at first, but the benefit of adding more modes to the reduced models diminishes rapidly. Finally, if we compare the errors in the BMS_{HCB} and BMS_{HCB+} models, we see that the residual-enhanced technique is able to converge much more quickly. Although there is slightly more overhead associated with the residual-enhanced technique, it is typically worthwhile because the reduced-order model can be made extremely small without compromising accuracy.

4.1.2. Band-structure calculation using $\omega(\mathbf{k})$ method

Using the $\omega(\mathbf{k})$ approach, we compute the first 10 dispersion branches of the 23040 DOF large mesh for 97 \mathbf{k} points. We then create a BMS_{HCB+}^{L-CC} reduced-order model with 80 DOFs (30 interior modes and 50 boundary modes), and again we compute the dispersion. Fig. 6a shows both the full- and reduced-model dispersion diagrams, and indicates very good visual agreement in the dispersion frequencies. Indeed the maximum frequency error over the first 10 branches is just 0.03% relative to the full model. This ensures that the overall error is dominated by the discretization and not by the reduction, but still produces a very compact reduced-order model. The reduced-model band-structure calculation time is just 4.1 seconds which is over 40 times faster than the full-model calculation time of 170.1 seconds. Of these 4.1 seconds, 3.9 are spent on the up-front calculations, and less than a quarter of a second is needed to visit all of the \mathbf{k} points. The performance details are summarized in Table 2.

Fig. 6b shows a mode-shape comparison between the full and reduced models for the point highlighted on the dispersion plot. The error between the full and reduced mode shapes is $e_\phi = 1.71 \times 10^{-6}$. The maximum mode error for the band structure computed in Fig. 6a is $e_\phi = 7.15 \times 10^{-5}$.

4.1.3. Complex band-structure calculation using $\mathbf{k}(\omega)$ method

As a second example using the same model, we compute both the propagating and evanescent band structure using the $\mathbf{k}(\omega)$ approach. Fig. 7 shows a comparison of the full-model and a BMS-reduced-order-model complex band structures computed for the Γ -X direction. For plotting and error comparison, 200 frequency steps are evaluated, and for each, only the 8 wave-vector solutions nearest the gamma point are considered.

The BMS_{HCB+}^{L-CC} model used to compute the complex band structure is the same 80 DOF model used in the previous section. The reduced-model band-structure calculation time is just 5.6 seconds which is over 1200 times faster than the full-model calculation time of 6858 seconds (approximately 2 hours). Less than 3 seconds are needed to visit all of the 200 ω points. The maximum error in the BMS wave vectors is 6.5% with respect to the full model.

The full model $\mathbf{k}(\omega)$ calculation takes significantly longer than the full model $\omega(\mathbf{k})$ calculation because a direct eigenvalue solver is used rather than an iterative solver. The eigenvalues are of the form $\lambda = e^{i\mathbf{k}^T \mathbf{r}}$, so finding the smallest wave-vector

Table 2
Comparison of full and BMS reduced-model calculation time and model-size for 8-lobed inclusion material band-structure calculations.

	Full model	BMS ^{L-CC} _{HCB+} model
Model details		
Model size (DOF)	23040	80 (30 interior, 50 boundary)
Up-front time (s)	–	1.9 (Interior mode) 1.1 (Constraint mode) + 0.9 (Other)
<hr/>		
$\omega(\mathbf{k})$ Calc.		
Per-k-point time (s)	1.75	0.0028
Total time, 97 k pts. (s)	170.1	4.1
Max. frequency error (%)	–	0.034
Max. mode error	–	7.15×10^{-5}
<hr/>		
$\mathbf{k}(\omega)$ Calc.		
Per- ω -pt. time (s)	51.5	0.0070
Total time, 200 ω pts. (s)	6858	5.6
Max. wave-vector error (%)	–	6.5

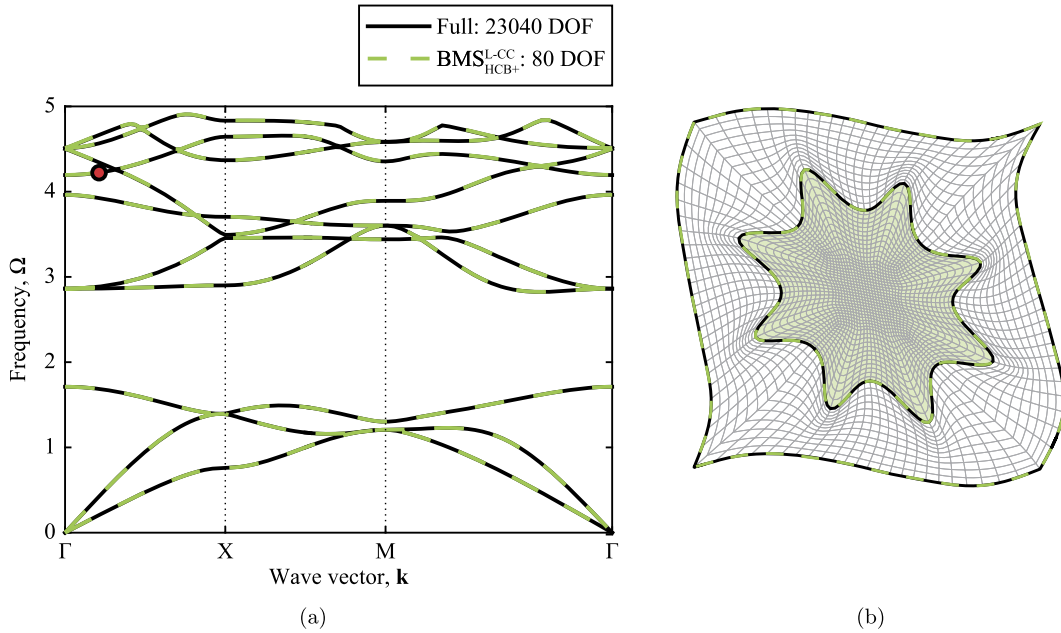


Fig. 6. (a) Dispersion-diagram comparison for 8-lobed inclusion material where $\Omega = \omega L_x \sqrt{\rho/E}$ is the non-dimensional frequency, and (b) selected mode-shape, computed by BMS, comparison for the mode highlighted in the dispersion diagram.

solutions is not as simple as computing the smallest eigenvalues. Rather, one must compute *all* of the eigenvalues (using a direct solver), convert them to wave-vectors, and then select those with smallest magnitude. The full-model $\mathbf{k}(\omega)$ calculation is significantly accelerated by making use of dynamic reduction, but even so the computation times are quite long due to the system solution of Eq. (20) required by the dynamic reduction at every frequency. The system solution can make use of a very fast sparse algorithm based on the supernodal Cholesky algorithm as long as the interior portion of the dynamic matrix, $\mathbf{D}_{II} = \mathbf{K}_{II} - \omega^2 \mathbf{M}_{II}$, is positive definite. When ω is large enough, \mathbf{D}_{II} becomes indefinite and a more general and much slower algorithm based on the multifrontal method must be used. The BMS model is not impacted very much by the need to use a direct solver because the reduced model is very small and fully populated. This leads to extreme computational gains over the full model.

It is also interesting to note that the BMS wave-vector errors are much larger than the frequency errors from the previous section. This occurs because reduced-order substructure representations, like the HCB method used here, produce models that are too stiff (meaning that frequencies are over-predicted). Effectively, every branch in the reduced-order model is shifted upwards slightly in frequency with respect to the full-model branches. Even when this frequency shift is very tiny,

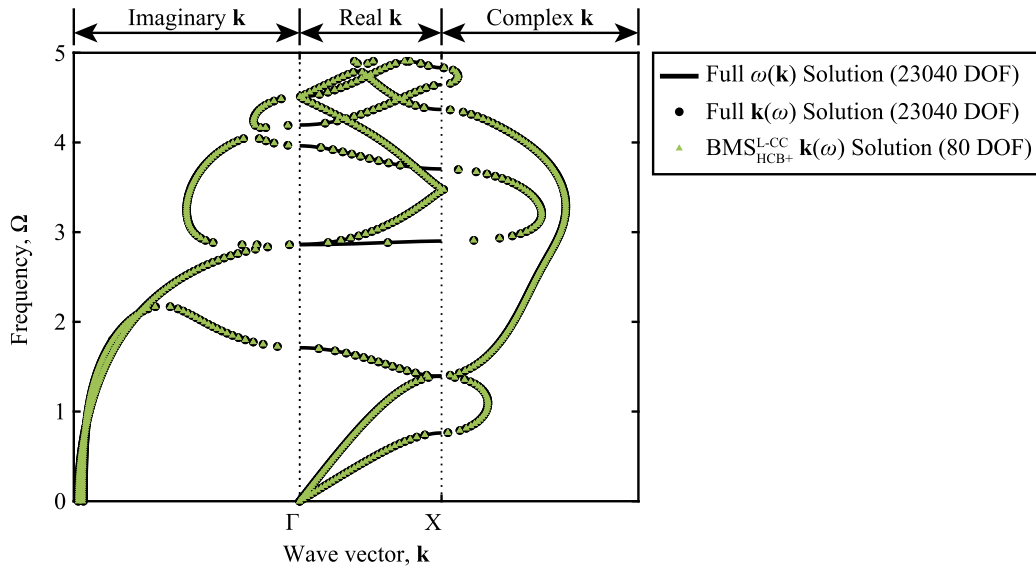


Fig. 7. Complex dispersion-diagram comparison for the 8-lobed inclusion material where $\Omega = \omega L_x \sqrt{\rho/E}$ is the non-dimensional frequency.




Fig. 8. Performance plot of maximum frequency error versus computation time, generated by computing $\omega(\mathbf{k})$ band structures with the large mesh of the 8-lobe inclusion material.

the resulting wave-vector errors can be relatively large when the branches are very flat. For the same reason, the $\mathbf{k}(\omega)$ error tends to be very sensitive to the frequency discretization as well. In light of these difficulties in evaluating $\mathbf{k}(\omega)$ error, we advise readers not to attach too much importance to this result.

4.1.4. Performance evaluation

We study the performance of the BMS method by computing the $\omega(\mathbf{k})$ band structure for various BMS reductions of the large mesh, and plotting the maximum frequency error versus computation time in Fig. 8. The maximum frequency error is evaluated for the first 10 branches of the band structure across 97 \mathbf{k} points. The curves in Fig. 8 are generated by varying the number of fixed-interface modes, n_A^ϕ , and performing either interface reduction using $p = 0.25$, or no interface reduction at all.

All of the methods show a sharp initial drop in error with little increase in computational time as the cutoff frequency is raised. Further increase in the cutoff frequencies produce diminishing returns as the curves level off. The residual-mode enhancement is able to reduce error levels by over two orders of magnitude with little additional computational cost. Interface reduction is able to reduce computation time without significantly increasing error levels.




Fig. 9. FE mesh for elastic metamaterial unit cell consisting of a soft rubber (white) and heavy lead (dark grey) spherical resonator embedded in an epoxy (light grey) matrix.

Table 3
Timing and model-size details for acoustic metamaterial.

	Full model	BMS _{HCB+} ^{L-CC} model
Model details		
Model size (DOF)	67656	628 (250 interior, 378 boundary)
Up-front calculation time (s)	–	1190 (Interior mode) 950 (Constraint mode) + 310 (Other)
		2450
$\omega(\mathbf{k})$ Calc.		
Per- \mathbf{k} -point time (s)	1398	0.61
Total time, 65 \mathbf{k} pts. (s)	90870	2490
Max. frequency error (%)	–	1.89
Max. mode error	–	0.0039

4.2. Elastic metamaterial model

The second material to be considered has a 3D cubic unit cell composed of resin with a resonator inside. The resonator consists of a lead sphere coated in silicon rubber. The mesh is plotted in Fig. 9 with a section removed to show the spherical resonator. The FEM model consists of 352 quintic Lagrangian brick elements and contains 67,656 DOFs (after enforcing Bloch boundary conditions). The computations for this material are performed on a 64-core server with 1.4 GHz AMD processors and 512 GB of memory.

4.2.1. Band-structure calculation using $\omega(\mathbf{k})$ method

We form a BMS_{HCB+}^{L-CC} reduced-order model containing 250 interior modes and 378 boundary modes for a total of 628 DOFs. The first 150 curves of the band structure are computed using the full model and the BMS model, and both are plotted in Fig. 10a. We first note that there are many flat bands whose frequencies do not change throughout the Brillouin zone. These represent resonator modes that couple very weakly or not at all with waves moving through the material. Although they do not carry any energy through the system, these modes are physically meaningful. The flat bands are grayed out in the band-structure diagram so that the energy carrying branches can be seen more easily.

We see that the BMS model produces band-structure curves that match the full model very well. The maximum error in band-structure frequencies is: 1.9%. Fig. 10b also shows a reduced-model mode shape corresponding to the point highlighted in the band-structure diagram. The maximum mode error over the first 148 branches is 0.0039. The 149th and 150th branches are ignored because they contain degeneracies with higher uncomputed branches. The error and timing results are summarized in Table 3.

4.2.2. Performance evaluation

We study the performance of the BMS method by computing the $\omega(\mathbf{k})$ band structure for various BMS reductions of the large mesh, and plotting the maximum frequency error versus computation time in Fig. 11. The maximum frequency

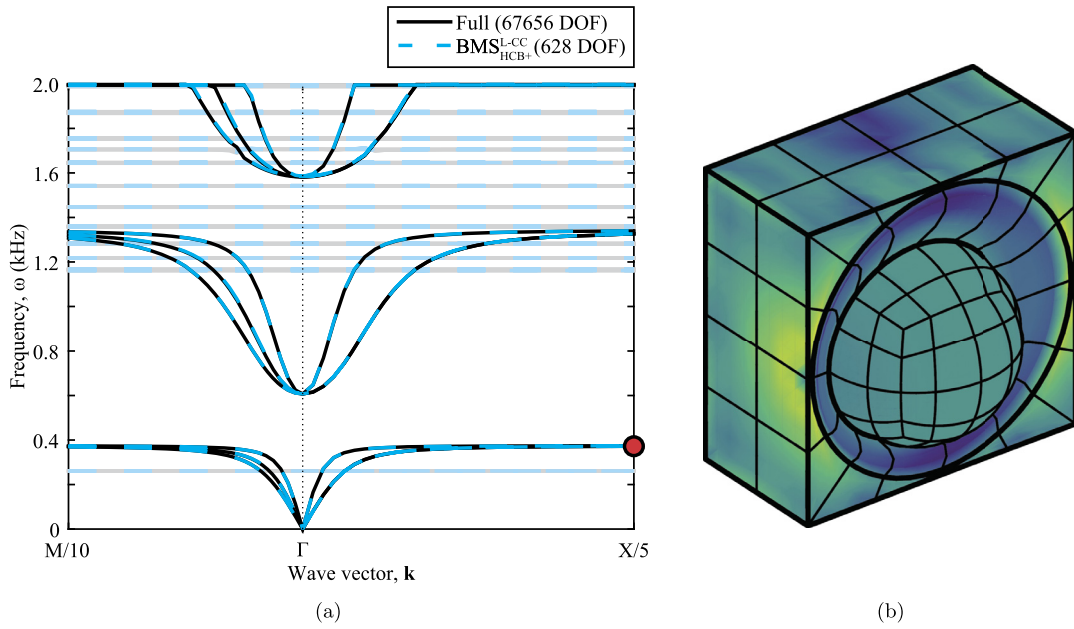


Fig. 10. (a) Dispersion-diagram comparison for elastic metamaterial, and (b) selected mode shape computed with BMS.

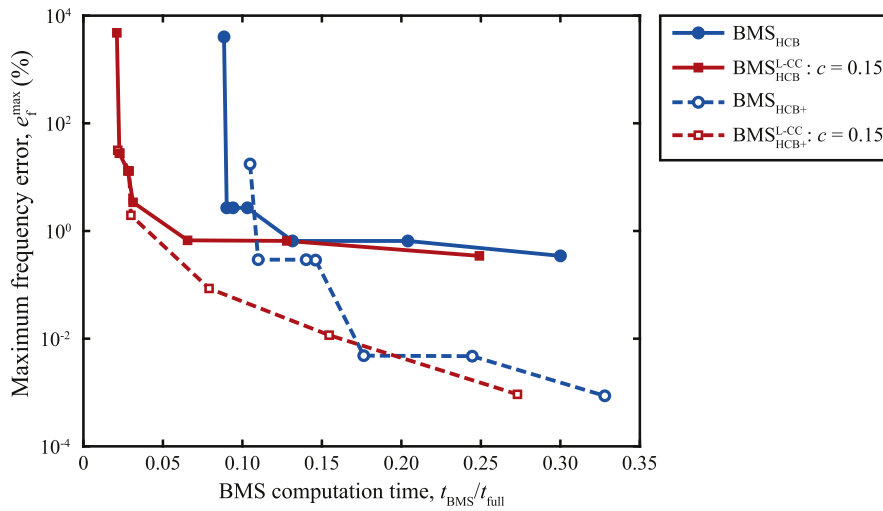


Fig. 11. Performance plot of maximum frequency error versus computation time, generated by computing $\omega(\mathbf{k})$ band structures for the resonator model.

error is evaluated for the first 150 branches of the band structure across 65 \mathbf{k} points. The curves in Fig. 11 are generated by varying the number of fixed-interface modes, n_A^ϕ , and performing either interface reduction using $p = 0.15$, or no interface reduction at all.

As with the previous model, all of the performance curves show a sharp initial drop in error with little increase in computational time as the cutoff frequency is raised. The error levels do not get quite as low as those in the previous example. A possible explanation for this is the presence of the flat branches that do not carry energy. Each of these flat branches corresponds to a fixed-interface mode. The flat branches effectively “use up” a fixed-interface mode that will not contribute to the solution at the elastic branches.

5. Conclusion

The original BMS formulation is an adaptation of CMS techniques to unit-cell problems aimed at greatly speeding up band-structure calculations based on FE or FD models. It does so by creating a reduced-order model independent of the Brillouin zone that is used to quickly solve for all of the band-structure frequencies.

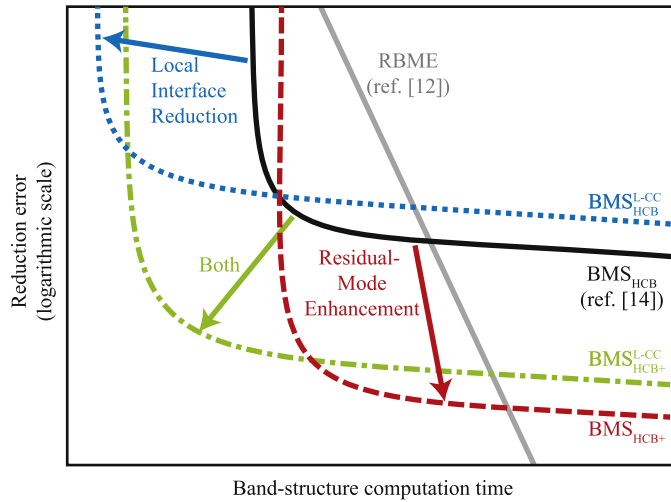


Fig. 12. Schematic qualitatively showing the effect of local interface reduction and residual-mode enhancement on BMS performance. Performance of the RBME method [12] is included for comparison.

In this paper we have presented two enhancements to the BMS method. First, local interface reduction reduces the model boundary in a way that is independent of Brillouin zone. Second, residual-mode enhancement increases the accuracy of the BMS reduction by approximating the contribution of mode shapes that are typically truncated in the BMS representation. The effect of these two improvements on the performance of the BMS method are summarized in Fig. 12. Generally, local interface reduction decreases computation time with very little increase in error. Residual-mode enhancement improves the accuracy but slightly increases the computation time. The combined effect is typically a reduced-order model that makes the band structure both faster to compute and more accurate than the original BMS reduction. The best performance for any type of BMS reduction is at the “corner” of each performance curve. Toward the tails, the error will either be very high or the computation time very long. To achieve good performance, it is important to choose an appropriate number of modes both for the interior reduction and for the boundary reduction.

Depending on the model being considered, computation times can be reduced by multiple orders of magnitude. We have shown in our examples that computation time can be reduced by over 40 times for $\omega(\mathbf{k})$ calculations, and by more than 1200 times for $\mathbf{k}(\omega)$ calculations. Furthermore, these speedups are realized without introducing error beyond the levels that are already present in our FE discretizations.

Acknowledgements

This work has been supported by the National Science Foundation grant No. 1538596.

Appendix A. Bloch mode synthesis basis functions

To visualize the different basis functions used in forming the BMS reduced-order model, an FE mesh of the circular-inclusion material unit cell is generated using linear elements. Fig. A.1 shows a few fixed-interface modes, Fig. A.2 shows a selection of constraint modes, Fig. A.3 shows local interface modes, and Fig. A.4 shows residual correction modes.




Fig. A.1. Selected fixed-interface modes for the circular-inclusion material unit cell. The fixed-interface modes are normal vibration modes of the unit-cell interior, and are computed by solving for the lowest frequency eigenvectors with all boundary DOFs fixed.

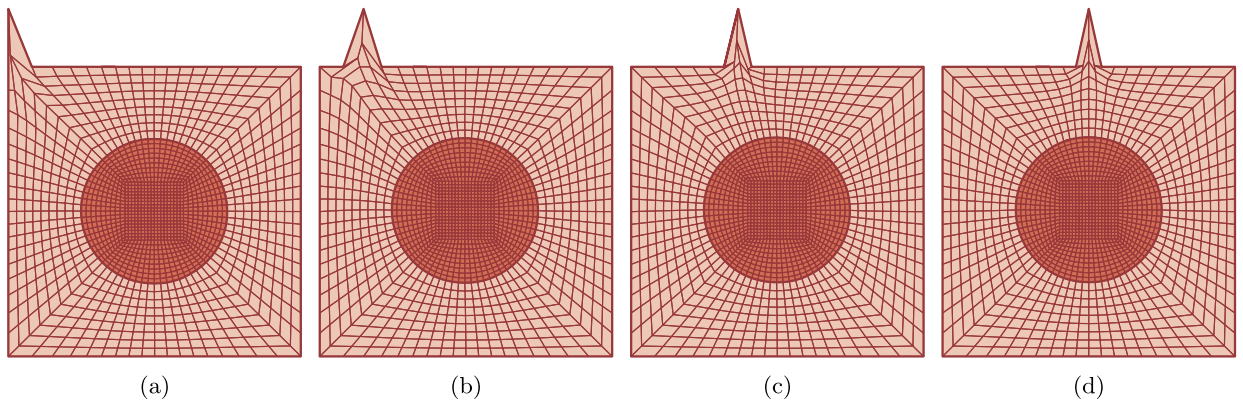


Fig. A.2. Selected constraint modes for the circular-inclusion material unit cell. Each constraint mode shows the static response of the interior to unit deflection of a single boundary DOF with all other boundary DOFs fixed in place.




Fig. A.3. Selected L-CC interface modes for the circular-inclusion material unit cell. The L-CC modes are formed by concatenating and then orthogonalizing boundary mode sets extracted from opposite unit-cell faces. Subplots (a), (b), and (c) show mode shapes with motion out to the plane of the interface, whereas the mode shape shown in (d) has motion entirely in the plane of the interface.




Fig. A.4. Selected residual correction modes for the circular-inclusion material unit cell. Each residual correction mode corresponds to an interface DOF. In this case, the interface DOFs are modal DOFs associated with the local interface modes. The correction modes shown in subplots (a–d) correspond to the local interface modes shown in subplots (a–d) of Fig. A.1 respectively. Looking closely, one can see a correlation between the deflection patterns of the local interface modes and their respective residual correction modes.

References

- [1] M.I. Hussein, M.J. Leamy, M. Ruzzene, Dynamics of phononic materials and structures: historical origins, recent progress, and future outlook, *Appl. Mech. Rev.* 66 (May 2014) 040802.
- [2] C. Kittel, *Introduction to Solid State Physics*, 8th ed., Wiley, Hoboken, NJ, 2005.
- [3] Felix Bloch, Über die Quantenmechanik der Elektronen in Kristallgittern, *Z. Phys.* (1928).
- [4] W. Axemann, P. Kuchment, An efficient finite element method for computing spectra of photonic and acoustic band-gap materials: I. scalar case, *J. Comput. Phys.* 150 (1999) 468–481.
- [5] R.C. Raffanetti, A simultaneous coordinate relaxation algorithm for large, sparse matrix eigenvalue problems, *J. Comput. Phys.* 32 (3) (1979) 403–419.

- [6] MATLAB Release 2017.
- [7] D.C. Dobson, J. Gopalakrishnan, J.E. Pasciak, An efficient method for band structure calculations in 3d photonic crystals, *J. Comput. Phys.* 161 (July 2000) 668–679.
- [8] R.L. Chern, C.C. Chang, C.C. Chang, R.R. Hwang, Large full band gaps for photonic crystals in two dimensions computed by an inverse method with multigrid acceleration, *Phys. Rev. E* 68 (Aug. 2003).
- [9] R.J. Guyan, Reduction of stiffness and mass matrices, *AIAA J.* 3 (2) (1965) 380.
- [10] R. Anderson, B. Irons, O. Zienkiewicz, Vibration and stability of plates using finite elements, *Int. J. Solids Struct.* 4 (1968) 1031–1055.
- [11] E.L. Shirley, Optimal basis sets for detailed Brillouin-zone integrations, *Phys. Rev. B* 54 (23) (1996) 16464.
- [12] M.I. Hussein, Reduced Bloch mode expansion for periodic media band structure calculations, *Proc. R. Soc. A, Math. Phys. Eng. Sci.* 465 (Sept. 2009) 2825–2848.
- [13] C. Scheiber, A. Schultschik, O. Biro, R. Dyczij-Edlinger, A model order reduction method for efficient band structure calculations of photonic crystals, *IEEE Trans. Magn.* 47 (May 2011) 1534–1537.
- [14] D. Krattiger, M.I. Hussein, Bloch mode synthesis: ultrafast methodology for elastic band-structure calculations, *Phys. Rev. E* 90 (Dec. 2014).
- [15] C.W. Zhou, J.P. Lainé, M.N. Ichchou, A.M. Zine, Multi-scale modelling for two-dimensional periodic structures using a combined mode/wave based approach, *Comput. Struct.* 154 (July 2015) 145–162.
- [16] C.W. Zhou, J.P. Lainé, M.N. Ichchou, A.M. Zine, Wave finite element method based on reduced model for one-dimensional periodic structures, *Int. J. Appl. Mech.* 07 (Apr. 2015) 1550018.
- [17] A. Palermo, A. Marzani, Extended Bloch mode synthesis: ultrafast method for the computation of complex band structures in phononic media, *Int. J. Solids Struct.* (2016).
- [18] M.P. Castanier, Y.-C. Tan, C. Pierre, Characteristic constraint modes for component mode synthesis, *AIAA J.* 39 (June 2001) 1182–1187.
- [19] M. Collet, M. Ouisse, M. Ruzzene, M.N. Ichchou, Floquet–Bloch decomposition for the computation of dispersion of two-dimensional periodic, damped mechanical systems, *Int. J. Solids Struct.* 48 (Oct. 2011) 2837–2848.
- [20] S.-K. Hong, B.I. Epureanu, M.P. Castanier, Next-generation parametric reduced-order models, *Mech. Syst. Signal Process.* 37 (May 2013) 403–421.
- [21] W.C. Hurty, Dynamic analysis of structural systems using component modes, *AIAA J.* 3 (Apr. 1965) 678–685.
- [22] M.C.C. Bampton, R.R. Craig, Coupling of substructures for dynamic analyses, *AIAA J.* 6 (7) (1968) 1313–1319.
- [23] R.R. Craig, C.J. Chang, On the Use of Attachment Modes in Substructure Coupling for Dynamic Analysis, American Institute of Aeronautics and Astronautics, San Diego, CA, Mar. 1977.
- [24] B. de Kraker, V. F. Werktuigkunde, Generalization of the Craig–Bampton CMS Procedure for General Damping, Report nr. WFW 93.023, Technische Universiteit Eindhoven, Eindhoven, February 1993.
- [25] R.R. Craig, Z. Ni, Component mode synthesis for model order reduction of nonclassically damped systems, *J. Guid. Control Dyn.* 12 (4) (1989) 577–584.
- [26] R.R. Craig, A.J. Kurdila, *Fundamentals of Structural Dynamics*, John Wiley & Sons, 2006.
- [27] Y. Chen, T.A. Davis, W.W. Hager, S. Rajamanickam, Algorithm 887: CHOLMOD, supernodal sparse Cholesky factorization and update/downdate, *ACM Trans. Math. Softw.* 35 (Oct. 2008) 22, 14 pp.
- [28] J.K. Bennighof, R.B. Lehoucq, An automated multilevel substructuring method for eigenspace computation in linear elastodynamics, *SIAM J. Sci. Comput.* 25 (Jan. 2004) 2084–2106.
- [29] M.F. Kaplan, Implementation of Automated Multilevel Substructuring for Frequency Response Analysis of Structures, PhD thesis, The University of Texas at Austin, Dec. 2001.
- [30] L. Suarez, M. Singh, Improved fixed interface method for modal synthesis, *AIAA J.* 30 (12) (1992) 2952–2958.
- [31] J.-G. Kim, P.-S. Lee, An enhanced Craig–Bampton method, *Int. J. Numer. Methods Eng.* 103 (July 2015) 79–93.
- [32] J.-G. Kim, S.-H. Boo, P.-S. Lee, An enhanced AMLS method and its performance, *Comput. Methods Appl. Mech. Eng.* 287 (Apr. 2015) 90–111.
- [33] S.-H. Boo, J.-G. Kim, P.-S. Lee, A simplified error estimator for the CB method and its application to error control, *Comput. Struct.* 164 (Feb. 2016) 53–62.
- [34] D.C. Kammer, M.J. Triller, Selection of component modes for Craig–Bampton substructure representations, *J. Vib. Acoust.* 118 (Apr. 1996) 264–270.
- [35] Z. Liu, Locally resonant sonic materials, *Science* 289 (Sept. 2000) 1734–1736.
- [36] The 4th International Conference on Phononic Crystals/Metamaterials, Phonon Transport/Coupling, Topological Phononics.