

High-order discretized ACMS method for the simulation of finite-size two-dimensional photonic crystals[Figure presented]

Giammatteo, Elena; Heinlein, Alexander; Lederer, Philip L.; Schlottbom, Matthias; Twente, University

DOI

[10.1016/j.camwa.2025.08.013](https://doi.org/10.1016/j.camwa.2025.08.013)

Publication date

2025

Document Version

Final published version

Published in

Computers & Mathematics with Applications

Citation (APA)

Giammatteo, E., Heinlein, A., Lederer, P. L., Schlottbom, M., & Twente, U. (2025). High-order discretized ACMS method for the simulation of finite-size two-dimensional photonic crystals[Figure presented]. *Computers & Mathematics with Applications*, 196, 376-394. <https://doi.org/10.1016/j.camwa.2025.08.013>

Important note


To cite this publication, please use the final published version (if applicable).
Please check the document version above.

Copyright

Other than for strictly personal use, it is not permitted to download, forward or distribute the text or part of it, without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license such as Creative Commons.

Takedown policy

Please contact us and provide details if you believe this document breaches copyrights.
We will remove access to the work immediately and investigate your claim.



High-order discretized ACMS method for the simulation of finite-size two-dimensional photonic crystals

Elena Giammatteo^{a,1}, Alexander Heinlein^{b, }, Philip L. Lederer^{c, },
Matthias Schlottbom^{a, },*

^a Department of Applied Mathematics, University of Twente, P.O. Box 217, Enschede, 7500 AE, the Netherlands

^b Delft Institute of Applied Mathematics, Delft University of Technology, Mekelweg 4, Delft, 2628 CD, the Netherlands

^c Department of Mathematics, University of Hamburg, Bundesstraße 55, Hamburg, 20146, Germany

ARTICLE INFO

Keywords:

Approximate component mode synthesis (ACMS)
Multiscale FEM
Helmholtz equation
hp Finite elements
Finite photonic crystals

ABSTRACT

The computational complexity and efficiency of the approximate mode component synthesis (ACMS) method is investigated for the two-dimensional heterogeneous Helmholtz equations, aiming at the simulation of large but finite-size photonic crystals. The ACMS method is a Galerkin method that relies on a non-overlapping domain decomposition and special basis functions defined based on the domain decomposition. While, in previous works, the ACMS method was realized using first-order finite elements, we use an underlying *hp*-finite element method. We study the accuracy of the ACMS method for different wavenumbers, domain decompositions, and discretization parameters. Moreover, the computational complexity of the method is investigated theoretically and compared with computing times for an implementation based on the open source software package NGSolve. The numerical results indicate that, for relevant wavenumber regimes, the size of the resulting linear systems for the ACMS method remains moderate, such that sparse direct solvers are a reasonable choice. Moreover, the ACMS method exhibits only a weak dependence on the selected domain decomposition, allowing for greater flexibility in its choice. Additionally, the numerical results show that the error of the ACMS method achieves the predicted convergence rate for increasing wavenumbers. Finally, to display the versatility of the implementation, the results of simulations of large but finite-size photonic crystals with defects are presented.

1. Introduction

Photonic crystals are nanostructures consisting of periodically repeating building blocks, the so-called unit cells, which usually consist of different materials, modeled by discontinuous material parameters [1,2]. Photonic crystals find many applications, such as enhanced absorption in solar cells and thin films [3–5], or the control of spontaneous emission of quantum emitters [6–9]. A mathematical treatment of photonic crystals can be found in [10,11].

* Corresponding author.

E-mail addresses: egiammatteo@utwente.nl (E. Giammatteo), a.heinlein@tudelft.nl (A. Heinlein), philip.lederer@uni-hamburg.de (P.L. Lederer), m.schlottbom@utwente.nl (M. Schlottbom).

¹ The research of this author was supported by the Dutch Research Council (NWO) via grant OCENW.GROOT.2019.071.

<https://doi.org/10.1016/j.camwa.2025.08.013>

Received 8 January 2025; Received in revised form 25 July 2025; Accepted 7 August 2025

The governing equations for light propagation in photonic crystals are Maxwell's equations, which may be reduced, under certain conditions, to the heterogeneous Helmholtz equations for transverse-magnetic and transverse-electric polarizations [2]. In view of the complex behavior of electromagnetic waves in photonic crystals, numerical approaches for the simulation of their optical properties are required, which we briefly discuss next.

Finite difference methods [12–16] are primarily used for structured grids, which allow for computationally highly efficient implementations. Using the regular sparsity structure of the resulting linear systems, problems up to tens of millions to one hundred million degrees of freedom can still be solved using sparse direct solvers [17,18]. Second order finite difference methods may suffer from a so-called pollution effect, i.e., large dispersion errors, which require much finer meshes than expected for a given wavenumber. Higher order finite difference methods weaken such strong assumption on the mesh, see, e.g., [19] for a discussion. However, the required smoothness of the fields may not be satisfied in the aforementioned application, e.g., due to discontinuous material parameters. Moreover, stair casing may occur for curved geometries, which requires non-trivial modifications for a robust and accurate implementation [20].

Finite element methods (FEM) are particularly well-suited for handling discontinuous material parameters due to their high geometric flexibility, but they also suffer from a pollution effect, which, in the FEM context, shows up in the form of wavenumber-dependent constants in the error analysis [21,22]. As a result, linear finite elements generally require excessive mesh resolution for growing wavenumbers, whereas high-order methods allow for overcoming the pollution effect; see, e.g., [23–27]. Let us also refer to interior penalty (high-order) discontinuous Galerkin methods [28,29] for the homogeneous Helmholtz equation. Additionally, boundary integral equation methods have been developed in [30] to solve plane wave scattering problems by periodic arrays of two-dimensional obstacles.

For large nanostructures, finite element methods lead to large linear systems, because of the required partitioning of the unit cells into sufficiently small elements. Therefore, solving these systems using a direct solver may become infeasible. Additionally, solving the resulting systems using iterative methods is also challenging for large wavenumbers: cf., e.g., [31,32]. To reduce the size of the linear systems, approaches that employ problem-adapted basis functions instead of piecewise polynomials have been employed. Trefftz discontinuous Galerkin methods [33–35] use solutions of the differential equation (plane waves) for the local approximation of the fields. For piecewise constant materials, these are often known, but for non-constant materials, the computation of such functions becomes more involved. In [36], polynomial discrete plane wave functions have been constructed numerically from a piecewise polynomial discretization. From all the degrees of freedom per element when considering a standard discontinuous Galerkin method, some are removed, while others are agglomerated into (polynomial) Trefftz ansatz functions. A detailed comparison for the latter is also given in [37].

For infinite periodically repeating nanostructures, the solutions are given in terms of Bloch modes [10,38]. In [39], a conforming finite element method that uses Bloch modes has been developed for 2D photonic crystals that are infinite in one direction and surrounded by an external medium in the other. The conformity requirement is satisfied by modulating the Bloch modes with polynomial functions, which makes the assembly of the corresponding matrices involved, but independent of the size of the structure. In [40, Chapter 5], photonic crystals which are finite in two-dimensions, are discretized using a discontinuous Galerkin method that treats the crystal as one element and that uses Bloch modes as basis functions within the crystal. Dropping the conformity requirement and exploiting properties of the Bloch modes greatly simplifies the assembly, yielding only a weak dependence on the number of unit cells; the dependence is linear instead of quadratic for crystals consisting of the same amount of unit cells in each direction.

When the periodic arrangement of the unit cells is broken, the aforementioned methods [39,40,30] are, however, not directly applicable. Such cases occur, e.g., in waveguides [2] or naturally in truly manufactured nanostructures due to unavoidable manufacturing flaws. For direct numerical simulations of real nanostructures and comparisons to the design structure using a discontinuous Galerkin method, see [41].

Multiscale methods [42–46] employ (local) problem adapted basis functions. As such, they aim to provide accurate approximation with only a few degrees of freedom, and at the same time, they can handle non-constant material properties and complex geometries. In practice, such methods usually also rely on an underlying discretization for their numerical realization.

This work focuses on the computational aspects and implementation of the approximate mode component synthesis (ACMS) method to solve two-dimensional heterogeneous Helmholtz equations. We refer to [43] for a theoretical analysis and preliminary numerical studies. The method has originally been introduced in [47] and further theoretically analyzed in [48] for heterogeneous diffusion problems. The ACMS method relies on a combination of a non-overlapping domain decomposition approach, which allows for straightforward parallelization over the number of unit cells during assembly, and the construction of problem-adapted basis functions. We mention that the ACMS method does not rely on (local) periodicity assumptions, allowing it to easily handle waveguides, defects, and other perturbations in the crystal structure. This flexibility broadens the method's applicability, including its use in the numerical simulation of quasicrystals; cf. [49]. The computational experiments in [43] considered relatively small examples and employed linear finite elements, which, as mentioned above, are susceptible to the pollution effect. Consequently, also the discrete realization of the ACMS method can potentially suffer from the pollution effect. Moreover, the efficient implementation of the ACMS method and its computational complexity has not been investigated in [43]; however, an efficient parallel implementation of the ACMS method employing linear finite elements for two-dimensional heterogeneous diffusion problems has been presented in [50].

Before going into a more detailed description of the ACMS method, we outline our main contributions, aimed at addressing the research gaps mentioned above. Our contributions include a detailed computational complexity analysis and an efficient implementation of the ACMS method using an underlying hp -finite element method. The efficiency of our implementation, i.e., assembly of linear systems and their solution using sparse direct solvers, is due to structural similarities to usual hp -finite element methods. In particular, the usual local degree p polynomials are replaced by problem-adapted basis functions; and when fixing the domain

decomposition, the ACMS method allows for increasing the number of local basis functions for each ACMS element (subdomain), similar to increasing the polynomial degree in hp -finite elements. The rather small number of basis functions required to achieve engineering relevant tolerances in practical wavenumber regimes yields linear systems of moderate size, allowing the use of sparse direct solvers. Additionally, since our implementation is based on underlying high-order finite elements, we can numerically investigate the wavenumber dependence of the method. This is particularly important for the Helmholtz equations, because Galerkin methods like the ACMS, require sufficient resolution in order to be well defined; see, e.g., [51]. Moreover, we study the influence on the computational complexity as well as on the accuracy of the choice of the domain decomposition and the number of modes per subdomain. Finally, we show that, by suitable choices of the discretization parameters, we can simulate large two-dimensional nanostructures. Our implementation employs the open source package NGSolve [52] and the code is publicly available at [53].

The remainder of the paper is structured as follows. We first introduce the Helmholtz equation and the essential notation in Section 2. Then, in Section 3, we give a summary of the continuous ACMS method and recall relevant analytical results. In Section 4, we describe the numerical realization and computational cost of the ACMS method when employing an underlying hp -FEM discretization. Finally, in Section 5, we present numerical results to highlight the method's flexibility in the choice of the domain decomposition and the benefit of choosing a high-order approximation of the ACMS basis functions. We conclude with final remarks in Section 6.

2. Notation and model equation

In the coming section, we present the basic notation along with the Helmholtz equation.

Notation. Given a connected polygonal domain $\Omega \subset \mathbb{R}^2$ with piecewise C^2 boundary and strictly convex angles, we denote with $L^2(\Omega)$ the Lebesgue space of square-integrable functions $u, v : \Omega \rightarrow \mathbb{C}$ with inner product

$$(u, v)_{\Omega} = \int_{\Omega} u \bar{v} \, d\mathbf{x},$$

and with $H^1(\Omega)$ the usual Sobolev space of functions in $L^2(\Omega)$ with square-integrable weak derivatives. Additionally, let $H^{1/2}(\partial\Omega)$ be the trace space of functions in $H^1(\Omega)$. Moreover, $H_0^1(\Omega)$ denotes functions in $H^1(\Omega)$ with vanishing trace on $\partial\Omega$ and $L^\infty(\Omega)$ denotes the space of essentially bounded measurable functions. Corresponding notation is also used for other measurable sets besides Ω . We indicate with $\|\cdot\|_V$ the corresponding norm on a function space V .

Problem formulation. We consider the heterogeneous Helmholtz problem:

$$\begin{aligned} -\operatorname{div}(a \nabla u) - \kappa^2 u &= f \quad \text{in } \Omega, \\ a \partial_n u - i \omega \beta u &= g \quad \text{on } \partial\Omega, \end{aligned} \tag{1}$$

where i is the imaginary unit. The coefficient functions $a, c \in L^\infty(\Omega)$ describe material properties of the background medium occupying Ω and are such that $a \geq a_0 > 0$ and $c \geq c_0 > 0$ for some constants $a_0, c_0 \in \mathbb{R}$. Together with the positive angular frequency denoted by ω , we then obtain the wavenumber $\kappa = \omega/c$. The real-valued function $\beta \in L^\infty(\partial\Omega)$ is related to the transmission and reflection of the unknown function u on the boundary $\partial\Omega$, and we assume that either $\beta \geq \beta_0 > 0$ or $\beta \leq \beta_0 < 0$, for a given $\beta_0 \in \mathbb{R}$. Impedance boundary conditions are described by the function $g \in H^{1/2}(\partial\Omega)$, and any interior sources may be modeled by the function $f \in L^2(\Omega)$. In the current work, we focus on the case of a homogeneous interior source, i.e., $f \equiv 0$, and we refer to Remark 1 below for a discussion of the more general case.

To obtain a weak formulation of the Helmholtz problem, let us introduce the sesquilinear form $C : H^1(\Omega) \times H^1(\Omega) \rightarrow \mathbb{C}$ defined by

$$C(u, v) = (a \nabla u, \nabla v)_{\Omega} - (\kappa^2 u, v)_{\Omega} - i(\omega \beta u, v)_{\partial\Omega},$$

and the antilinear functional $G : H^1(\Omega) \rightarrow \mathbb{C}$ defined by

$$G(v) = (g, v)_{\partial\Omega}. \tag{2}$$

The weak form of the Helmholtz problem (1) then reads

$$\text{Find } u \in H^1(\Omega) : \quad C(u, v) = G(v), \quad \text{for all } v \in H^1(\Omega). \tag{3}$$

With the aforementioned regularity assumptions on the coefficients and the functions, the Helmholtz problem is well-posed and the weak formulation (3) admits a unique solution; see [51, Theorem 2.4].

3. The continuous ACMS method

In the following section, we recap the ACMS method and recall some theoretical convergence results obtained in [43].

Domain decomposition. We introduce by $D = \{\Omega_j, j = 1, \dots, J\}$ a conforming decomposition of Ω into J non-overlapping open subdomains with piecewise smooth boundaries, and we define the interface Γ of the domain decomposition as

$$\Gamma = \bigcup_{j=1}^J \partial\Omega_j,$$

which also includes the boundaries of the domain Ω . Moreover, we introduce the set \mathcal{E} of edges and the set \mathcal{V} of vertices of the domain decomposition:

$$\mathcal{E} = \{e \subset \Gamma : \bar{e} = \partial\Omega \cap \partial\Omega_j \text{ or } \bar{e} = \partial\Omega_j \cap \partial\Omega_i, \text{ for some } i, j = 1, \dots, J\},$$

$$\mathcal{V} = \{q \in \Gamma : \exists e \in \mathcal{E} : q \in \partial e\}.$$

ACMS spaces. The proper formulation of the ACMS method relies on the solvability of local Helmholtz problems with Dirichlet boundary conditions [43]. To that end, fix $j \in \{1, \dots, J\}$, and let us consider the following eigenvalue problem: Find $b \in H_0^1(\Omega_j) \setminus \{0\}$ and $\lambda \in \mathbb{R}$ such that

$$(a \nabla b, \nabla v)_{\Omega_j} = \lambda(\kappa^2 b, v)_{\Omega_j} \quad \text{for all } v \in H_0^1(\Omega_j). \quad (4)$$

The local solvability condition can then be conveniently characterized by requiring that $\lambda \neq 1$, for all $j \in \{1, \dots, J\}$, which we assume in the following. To discuss how the condition $\lambda \neq 1$ is related to ω , which is important for investigating the wavenumber dependency of the method, let us denote by λ_i^j , $i \in \mathbb{N}$, the non-decreasing sequence of eigenvalues of (4) for the subdomains Ω_j . Since $\kappa = \omega/c$, we observe that $\lambda_i^j = \tilde{\lambda}_i^j / \omega^2$ with $\tilde{\lambda}_i^j$ independent of ω . The condition $\lambda_i^j \neq 1$ is then equivalent to $\tilde{\lambda}_i^j \neq \omega^2$. Therefore, to ensure that $|\lambda_i^j - 1|$ is uniformly positive for changing ω , we additionally assume in the following that there is $\gamma > 0$ such that

$$\omega \in (0, \infty) \setminus \bigcup_{i,j} \left[(1 - \gamma) \sqrt{\tilde{\lambda}_i^j}, (1 + \gamma) \sqrt{\tilde{\lambda}_i^j} \right]. \quad (5)$$

We note that in other works, such as [44,54,55], the condition (5) is replaced by a stronger local coercivity assumption, which can be expressed as $\omega \leq (1 - \gamma) \sqrt{\tilde{\lambda}_1^j}$. The validity of the local coercivity condition implies that for increasing frequencies ω , the smallest eigenvalue $\tilde{\lambda}_1^j$ is required to grow. Since $\tilde{\lambda}_1^j$ increases if the area of Ω_j decreases properly, this can be related to a refinement of the domain decomposition, i.e., a sufficient resolution condition.

Assuming (5) and recalling that we assumed $f = 0$, the solution u to (3) is determined by its trace on Γ , see Remark 1 below for the case $f \neq 0$. The ACMS method now relies on the approximation of $u|_\Gamma$ using functions that are locally supported on Γ . We will introduce the local functions associated with vertices and edges as follows. For all $q \in \mathcal{V}$, let $\varphi^q : \Gamma \rightarrow \mathbb{R}$ be edgewise harmonic functions defined on Γ , i.e., $\varphi^q(r) = \delta_{q,r}$ for $q, r \in \mathcal{V}$, where δ is the Kronecker delta, and

$$\int_e \partial_e \varphi^q \partial_e \eta \, ds = 0 \quad \text{for all } \eta \in H_0^1(e),$$

where ∂_e is the tangential derivative along the edge e . The vertex-based space is then defined as the linear combination of the functions φ^q :

$$V_{\mathcal{V}} = \text{span}\{\varphi^q : q \in \mathcal{V}\}.$$

Additionally, let us define the edge modes as solutions to the following weak formulation of the edge-Laplace eigenvalue problems: for each $e \in \mathcal{E}$, find $(\tau_i^e, \lambda_i^e) \in H_0^1(e) \times \mathbb{R}$, $i \in \mathbb{N}$, such that

$$(\partial_e \tau_i^e, \partial_e \eta)_e = \lambda_i^e (\tau_i^e, \eta)_e, \quad \text{for all } \eta \in H_0^1(e). \quad (6)$$

According to [56, p. 415], the eigenvalues depend quadratically on the index of the mode, that is

$$\lambda_i^e \sim \left(\frac{i\pi}{|e|} \right)^2. \quad (7)$$

In the following, we assume that the eigenvalues λ_i^e are ordered non-decreasingly. With a slight abuse of notation, we continue to use the symbol τ_i^e to denote the extension by zero of the edge basis function to the whole interface Γ . The edge-based space is then defined as the linear combination of the eigenmodes τ_i^e for all edges

$$V_{\mathcal{E}} = \text{span}\{\tau_i^e : e \in \mathcal{E}, i \in \mathbb{N}\}.$$

It follows from [43, Lemma 3.6], that $V_{\mathcal{V}} + V_{\mathcal{E}}$ is dense in the trace space

$$H^{1/2}(\Gamma) = \{v : \Gamma \rightarrow \mathbb{C} : \forall j = 1, \dots, J, \exists u_j \in H^1(\Omega_j) \text{ s.t. } u_j|_{\partial\Omega_j} = v|_{\partial\Omega_j}\}.$$

In order to approximate (3), we next extend functions in $H^{1/2}(\Gamma)$ to functions in $H^1(\Omega)$. To do this, we first introduce the extensions E^j associated with the subdomains Ω_j , and then the extension E^Γ associated with the full interface Γ . For a given $\Omega_j \in \mathcal{D}$ and $\tau \in H^{1/2}(\partial\Omega_j)$, let us define the local Helmholtz-harmonic extension operator $E^j : H^{1/2}(\partial\Omega_j) \rightarrow H^1(\Omega_j)$ such that

$$\begin{aligned} (a \nabla E^j \tau, \nabla \eta)_{\Omega_j} - (\kappa^2 E^j \tau, \eta)_{\Omega_j} &= 0, \quad \text{for all } \eta \in H_0^1(\Omega_j), \\ (E^j \tau)|_{\partial\Omega_j} &= \tau. \end{aligned} \quad (8)$$

The Helmholtz-harmonic extension operator is well-defined and bounded independently of ω if (5) holds, which follows from the above assumption on the solvability of local Helmholtz problems [43]. Let us also introduce $E^\Gamma : H^{1/2}(\Gamma) \rightarrow H^1(\Omega)$, which first restricts a given function $\tau \in H^{1/2}(\Gamma)$ to the boundaries of the subdomains Ω_j , for all $j = 1, \dots, J$, and then applies the local Helmholtz-harmonic extension operator, i.e., $(E^\Gamma \tau)|_{\Omega_j} = E^j(\tau|_{\partial\Omega_j})$. We note that a vertex basis function φ^q is supported on all edges that share the vertex q while its extension $E^\Gamma \varphi^q$ is supported on all subdomains Ω_j that share the vertex q . In contrast, the extension of an edge basis function $E^\Gamma \tau_i^e$ is supported on the two neighboring subdomains that share the edge e . Moreover, we can write the solution u of (3) in terms of its trace $u|_\Gamma \in H^{1/2}(\Gamma)$ that is extended to $H^1(\Omega)$ with the operator that was just introduced: $u = E^\Gamma(u|_\Gamma)$.

Using the above extensions, we can now approximate the solution space using a combination of local spaces associated with vertex functions and edge modes:

$$V(\mathcal{D}) := E^\Gamma V_\mathcal{V} + E^\Gamma V_\mathcal{E} \subseteq H^1(\Omega).$$

We are now able to define the finite-dimensional ACMS space $V_A(\mathcal{D}) \subset V(\mathcal{D})$. For every edge, we select the eigenmodes associated with the $I_\mathcal{E} \in \mathbb{N}$ smallest eigenvalues, assuming that they are non-decreasingly ordered thus obtaining

$$V_A(\mathcal{D}) := E^\Gamma V_\mathcal{V} + E^\Gamma \mathcal{V}_\mathcal{E}^{I_\mathcal{E}}, \quad \text{with } \mathcal{V}_\mathcal{E}^{I_\mathcal{E}} = \text{span}\{\tau_i^e : e \in \mathcal{E}, 1 \leq i \leq I_\mathcal{E}\}. \quad (9)$$

Remark 1. The methodology presented here can be extended to more general f by replacing the antilinear functional G defined in (2) with

$$(f, v)_\Omega + (g, v)_{\partial\Omega}.$$

In this case, the technique illustrated here only approximates $E^\Gamma(u|_\Gamma)$ and not the full solution u . In fact, we have the decomposition $u = u_B + E^\Gamma(u|_\Gamma)$. The term u_B , which is not discussed here, satisfies $u_B|_{\Omega_j} \in H_0^1(\Omega_j)$ for each subdomain Ω_j and can, for instance, be approximated using the local eigenfunctions specified in (4). Hence, the numerical approximation of u_B can be done independently for each subdomain Ω_j and in parallel. Moreover, the computation of u_B is independent of the approximation of $E^\Gamma(u|_\Gamma)$. We refer the reader to [43] for more details.

ACMS approximation of the Helmholtz equation. The weak formulation (3) of the Helmholtz equation can be approximated in the ACMS space: find $u_A \in V_A(\mathcal{D})$ such that

$$C(u_A, v_A) = G(v_A), \quad \text{for all } v_A \in V_A(\mathcal{D}). \quad (10)$$

The well-posedness of the approximation problem (10) relies on the smallness of the adjoint approximability constant σ^* ; see, for instance, [42,44,51]. To introduce σ^* in our context, denote $T^* : L^2(\Omega) \rightarrow V(\mathcal{D})$ the solution operator of the dual problem to (3) restricted to $V(\mathcal{D})$, i.e., for $\chi \in L^2(\Omega)$ we define $T^* \chi = z$, where $z \in V(\mathcal{D})$ is the unique solution to

$$C(v, z) = (v, \chi), \quad \text{for all } v \in V(\mathcal{D});$$

see [43] for details and [51] for the well-posedness of the dual problem. The adjoint approximability constant specifies how well the solution to the dual problem can be approximated by functions in $V_A(\mathcal{D})$, i.e.,

$$\sigma^* = \sup_{\varphi \in L^2(\Omega) \setminus \{0\}} \frac{\inf_{v_A \in V_A(\mathcal{D})} \|T^*(\kappa^2 \varphi) - v_A\|_B}{\|\kappa \varphi\|_{L^2(\Omega)}},$$

where the B -norm is defined by

$$\|v\|_B^2 = (a \nabla v, \nabla v)_\Omega + (\kappa^2 v, v)_\Omega, \quad v \in H^1(\Omega).$$

In [43, Theorem 4.6], it has been shown that $\sigma^* = \mathcal{O}(C_\kappa \|\kappa\|_\infty^2 / I_\mathcal{E})$ if the coefficients a and β are Lipschitz continuous in a neighborhood of the interface Γ . Here, C_κ is the stability constant of the Helmholtz problem, i.e., a bound on the solution operator associated with (3), mapping $f \in L^2(\Omega)$ to $u \in H^1(\Omega)$, for $g = 0$. We note that C_κ in general depends on κ , and we refer to [51] for a discussion. Hence, σ^* can be made small by using sufficiently many edge modes. Moreover, under these regularity assumptions on a and β , it has been shown in [43, Theorem 4.7] that

$$\|u - u_A\|_{L^2(\Omega)} \leq C \frac{\max_{e \in \mathcal{E}} |e|^3 C_\kappa \|\kappa\|_\infty^2}{I_\mathcal{E}^3} \sum_{e \in \mathcal{E}} \|u\|_{H^3(e)}, \quad (11)$$

where C is a constant independent of κ . These results indicate that the wavenumber κ influences the number of edge modes required to obtain a given L^2 -error of the ACMS approximation against the exact solution. Since other factors enter the estimate, the relation between edge modes and wavenumber in this bound will be further investigated numerically; see Sections 5.1 and 5.2.

We close this section with some remarks on the error bounds. In addition to the stated L^2 -error bound (11), the smallness of σ^* implies a corresponding H^1 -error bound in terms of $\|\kappa\|_\infty \sum_{e \in \mathcal{E}} |e|^2 \|u\|_{H^3(e)} / I_{\mathcal{E}}^2$; see [43, Lemma 4.4]. Moreover, we note that the regularity $u|_e \in H^3(e)$ is often met in practice. For example, by elliptic regularity, if the boundary data are smooth and the material parameters are constant in a neighborhood of Γ , which is the case in our numerical examples below. In case of discontinuous parameters the solution u may be less smooth. The ACMS method is then still applicable, but the convergence rates are lower and σ^* decays sub-linearly in $1/I_{\mathcal{E}}$. We refer to [57] for a detailed error analysis.

4. The discrete ACMS method

The approximation space (9) is spanned by extensions of the vertex basis functions and of the edge modes defined in (6). In practice, though, we need to numerically approximate the eigenproblems and harmonic extension operators. We now discuss the numerical realization of the ACMS method when employing an hp -FEM discretization.

Triangulation. For the sake of simplicity, we suppose that all edges are line segments, although the following discussion would hold similarly in the case of curved edges. Let $\mathcal{T}_h = \{T_1, \dots, T_N\}$ be a triangulation of the domain Ω without hanging nodes, with mesh size

$$h = \max_{T \in \mathcal{T}} h_T,$$

where h_T is the diameter of T ; see, e.g., [58]. If not explicitly mentioned differently, we use a quasi-uniform triangulation. Additionally, we want the mesh to be conforming with the domain decomposition, namely, we require every open triangle $T \in \mathcal{T}_h$ to be contained exactly in one subdomain $T \subset \Omega_j$.

FEM approximation of the Helmholtz equation. For the sake of the discussion, let us present a finite element approximation of the Helmholtz equation.

Let \mathcal{P}_p be the set of all polynomials in two variables of degree less or equal to $p \geq 1$. Then, the usual finite element space $V^{h,p}(\Omega)$ defined on the triangulation of the domain Ω consisting of piecewise polynomials of degree p is given by

$$V^{h,p}(\Omega) = \{v \in C^0(\overline{\Omega}) : v|_T \in \mathcal{P}_p, \forall T \in \mathcal{T}_h\};$$

see, e.g., [58]. The finite element approximation of the weak formulation (3) reads as: find $u^{h,p} \in V^{h,p}(\Omega)$ such that

$$C(u^{h,p}, v^{h,p}) = G(v^{h,p}), \quad \text{for all } v^{h,p} \in V^{h,p}(\Omega). \quad (12)$$

By choosing a basis for $V^{h,p}(\Omega)$, equation (12) reads

$$\mathbb{S}_F \mathbf{u}_F = \mathbf{g}_F,$$

where $\mathbf{u}_F \in \mathbb{R}^{N_F}$ is the solution vector, $\mathbb{S}_F \in \mathbb{R}^{N_F \times N_F}$ is the system matrix and $\mathbf{g}_F \in \mathbb{R}^{N_F}$ is the right-hand side vector, and $N_F = \dim V^{h,p}(\Omega)$.

4.1. Setup of the discrete ACMS method

In order to define the discrete counterparts of the elements in $V_A(D)$, we introduce auxiliary spaces on a generic set $\mathcal{M} \subseteq \overline{\Omega}$ in the following way:

$$\begin{aligned} V^{h,p}(\mathcal{M}) &= \{v|_{\mathcal{M}} : v \in V^{h,p}(\Omega)\}, \\ V_0^{h,p}(\mathcal{M}) &= \{v|_{\mathcal{M}} : v \in V^{h,p}(\Omega), v|_{\partial\mathcal{M}} = 0\}. \end{aligned}$$

The discrete vertex basis function of a generic vertex $q \in \mathcal{V}$ is referred to as $\varphi_h^q \in V^{h,p}(\Gamma)$ and is such that $\varphi_h^q(r) = \delta_{q,r}$ for all $q, r \in \mathcal{V}$. Moreover, the function is piecewise harmonic on the edges in the discrete sense, meaning that it satisfies

$$(\partial_e \varphi_h^q, \partial_e \eta^{h,p})_e = 0, \quad \text{for all } \eta^{h,p} \in V_0^{h,p}(e), \text{ for all } e \in \mathcal{E}. \quad (13)$$

The discrete edge modes are computed by numerically solving the eigenvalue problem (6): for $e \in \mathcal{E}$ and $i = 1, \dots, I_{\mathcal{E}}$, find $(\tau_{e,i}^{h,p}, \lambda_i^e) \in V_0^{h,p}(e) \times \mathbb{R}$ such that

$$(\partial_e \tau_{e,i}^{h,p}, \partial_e \eta^{h,p})_e = \lambda_i^e (\tau_{e,i}^{h,p}, \eta^{h,p})_e, \quad \text{for all } \eta^{h,p} \in V_0^{h,p}(e). \quad (14)$$

Again, with a slight abuse of notation we continue to use the symbol $\tau_{e,i}^{h,p}$ to denote the extension by zero of the edge basis function to the whole interface Γ . As in (8), we introduce the discrete Helmholtz-harmonic extension $E_h^j : V^{h,p}(\partial\Omega_j) \rightarrow V^{h,p}(\Omega_j)$ such that:

$$\begin{aligned} (a \nabla E_h^j \tau^{h,p}, \nabla v^{h,p}) - (\kappa^2 E_h^j \tau^{h,p}, v^{h,p})_{\Omega_j} &= 0, \quad \text{for all } v^{h,p} \in V_0^{h,p}(\Omega_j), \\ (E_h^j \tau^{h,p})|_{\partial\Omega_j} &= \tau^{h,p}. \end{aligned} \quad (15)$$

Let us introduce the discrete extension $E_h^\Gamma : V^{h,p}(\Gamma) \rightarrow V^{h,p}(\Omega)$, which first restricts functions to the boundaries of the subdomains Ω_j and then applies E_h^j .

The discrete ACMS space, indicated with the subscript A , is then defined on the domain decomposition \mathcal{D} as the span of the extended discrete edge modes and discrete vertex basis functions:

$$V_A^{h,p}(\mathcal{D}) = \text{span}\{E_h^\Gamma \tau_{e,i}^{h,p}, \forall e \in \mathcal{E}, i = 1, \dots, I_{\mathcal{E}}\} \cup \text{span}\{E_h^\Gamma \phi_h^q, \forall q \in \mathcal{V}\}.$$

The weak formulation (10) can be rewritten in the ACMS space: find $u_A^{h,p} \in V_A^{h,p}(\mathcal{D})$ such that

$$C(u_A^{h,p}, v_A^{h,p}) = G(v_A^{h,p}), \quad \text{for all } v_A^{h,p} \in V_A^{h,p}(\mathcal{D}). \quad (16)$$

Finally, the weak formulation in the ACMS space (16) turns into

$$\mathbb{S}_A \mathbf{u}_A = \mathbf{g}_A, \quad (17)$$

where $\mathbf{u}_A \in \mathbb{R}^{N_A}$ is the solution vector, $\mathbb{S}_A \in \mathbb{R}^{N_A \times N_A}$ is the system matrix and $\mathbf{g}_A \in \mathbb{R}^{N_A}$ is the right-hand side vector. Here, $N_A = \dim V_A^{h,p}(\mathcal{D})$ is the space dimension.

4.2. Computational costs

In the following, we give some insights on the expected computational costs, in terms of FLOPS, when using the discrete ACMS method. For more details, we refer to [50], where an efficient and parallel implementation for $p = 1$ is discussed. We are interested in the dependence of the cost on the number of modes $I_{\mathcal{E}}$, the number of subdomains J and related quantities, such as $|\mathcal{V}|$, $|\mathcal{E}|$, and the corresponding dimension of the ACMS space $N_A = |\mathcal{V}| + |\mathcal{E}| I_{\mathcal{E}}$.

To ease the presentation, we make the following assumptions:

- all subdomains $\Omega_j \in \mathcal{D}$ of the domain decomposition are of comparable size, and correspondingly also the dimensions $N_j = \dim V^{h,p}(\Omega_j)$ are similar,
- all edges $e \in \mathcal{E}$ are of comparable size, and correspondingly also the dimensions $n_e = \dim V^{h,p}(e)$ are similar.

For a given mesh size h and polynomial order p of the underlying finite element space, we can introduce the following two costs

C_{ext} : cost of solving the extension problem (15) on a subdomain $\Omega_j \in \mathcal{D}$,

$C_{eig}(I_{\mathcal{E}})$: cost of solving the eigenvalue problem (14) on an edge $e \in \mathcal{E}$.

Edge basis computation. The cost for the computation of the edge modes is given by solving the eigenvalue problem (14) for all edges $e \in \mathcal{E}$ and by the application of the extension operator to the corresponding neighboring subdomains for each mode (and each edge) separately; see Remark 3. Thus, it is given by the respective contributions

$$\mathcal{O}(|\mathcal{E}| C_{eig}(I_{\mathcal{E}})) + \mathcal{O}(|\mathcal{E}| I_{\mathcal{E}} C_{ext}). \quad (18)$$

Remark 2. When the eigenvalue computation has to be done for each edge separately, the first term in (18) might dominate when $I_{\mathcal{E}}$ grows, although, in practice, we expect that the problem structure can be exploited such that the eigenvalue problem only needs to be solved $\mathcal{O}(1)$ times; see Section 4.3.

Vertex basis computation. The costs of the vertex basis functions, where we only consider the application of the harmonic extensions, are given by

$$\mathcal{O}(|\mathcal{V}| C_{ext}). \quad (19)$$

Note that, for the above cost, we assumed that the restrictions on the edges of functions ϕ_h^q which fulfill property (13) are either given or very cheap to compute. This assumption is motivated by the fact that, in case the edges of the domain decomposition are straight lines, the restriction of ϕ_h^q on an edge $e \in \mathcal{E}$ is given by a linear function.

Assembly. From an implementation perspective, the ACMS method exhibits some similarities to a high-order finite element method, which can be utilized during setup and assembly; see also Remarks 6 and 7. More precisely, as already discussed in [47,50], the assembly is done in the usual way, via a loop over the subdomains $\Omega_j \in \mathcal{D}$, see also, e.g., [59]. On each cell Ω_j , one then first computes the local ACMS basis functions (as discussed in the previous paragraph) and then calculates the local ACMS stiffness matrix




Fig. 1. Left: Sparsity pattern of \mathbb{S}_A for $I_E = 32$ and $J = 36$ domains for Ω as in Section 4.3. The bandwidth is approximately $b_A \sim 3\sqrt{J}I_E$, hence $b_A \sim 384$. Right: dof numbering of a crystal with $J = 4$ and $I_E = 2$.

\mathbb{S}_A by means of the stiffness matrix \mathbb{S}_F of the underlying finite element space $V^{h,p}(\Omega_j)$. With a slight abuse of notation let us write the local contributions on $\Omega_j \in \mathcal{D}$ as

$$\mathbb{S}_A|_{\Omega_j} = \mathbb{B}_j^T \mathbb{S}_F|_{\Omega_j} \mathbb{B}_j, \quad (20)$$

where the restriction is understood as taking only into account degrees of freedom of basis functions (of the corresponding finite element spaces) which have a non-zero support on Ω_j . Likewise, we indicate with \mathbb{B}_j the matrix where each column corresponds to a restriction of the ACMS-basis functions supported on Ω_j . Consequently, the column dimension of \mathbb{B}_j scales with $\mathcal{O}(I_E)$. Since $\mathbb{S}_F|_{\Omega_j}$ is sparse with $\mathcal{O}(N_j)$ bandwidth, the costs for (20) are $\mathcal{O}(N_j I_E^2)$. Since this has to be done on each subdomain, the total cost for the assembly is given by

$$\mathcal{O}(J N_j I_E^2).$$

Solving the system. Finally, the cost for solving the system (17) is $\mathcal{O}(N_A^3)$ in case the matrix \mathbb{S}_A is dense. However, since the matrix \mathbb{S}_A shows a sparsity pattern, due to the aforementioned local support of the ACMS basis functions, as depicted in Fig. 1, these costs can be further reduced to

$$\mathcal{O}(N_A b_A^2), \quad (21)$$

where b_A is the bandwidth of \mathbb{S}_A ; see [60, Ch. 2.1.1]. We note that on structured grids this estimate might be pessimistic; see [61]. Indeed, we observe a better scaling in the following section.

Remark 3. The cost for the extension C_{ext} on a given domain $\Omega_j \in \mathcal{D}$, corresponds to solving a sparse system associated with problem (15) and is of order $\mathcal{O}(N_j n_e^2)$; see [60, Ch. 2.1.1]. In the implementation of the extension operators, we only compute the factorization once (using a sparse Cholesky solver), and subsequently apply the forward/backward substitution for each basis function independently. Therefore, the contribution of C_{ext} is split into factorization and application costs. The factorization is reused for both edge and vertex basis functions and, therefore, is carried out once per subdomain, accounting for a total cost of $\mathcal{O}(J n_e^4)$. The application is carried out I_E times per subdomain in the edge basis computation (18), meaning $\mathcal{O}(|\mathcal{E}| I_E n_e^3)$, and only once per subdomain in the vertex basis computation (19), which is $\mathcal{O}(|\mathcal{E}| n_e^3)$.

Remark 4. We emphasize that $C_{eig}(I_E)$ clearly depends on the choice of the number of edge modes I_E . In fact, the cost for solving the eigenvalue problem on an edge $e \in \mathcal{E}$ is of order $\mathcal{O}(n_e^2 I_E)$, using, e.g., an Arnoldi GMRES, see [60, Ch. 2.6.2].

Remark 5. We remind that, since it always holds that $I_E \leq n_e$, the computation of a high number of edge modes will require enriching the underlying finite element space $V_0^{h,p}(e)$. That can be obtained by either refining the mesh size or increasing the polynomial degree of approximation.

Remark 6. If the ACMS method is interpreted as a high-order finite element method, increasing the maximum number of modes I_E is similar to having a p -FEM method, while increasing the number of domains J can be related to an h -FEM method.

Remark 7. So far, we have not considered the advantages of potential parallel implementation. Nevertheless, we want to emphasize that, as with standard finite element methods, there is considerable room for improvement. For example, the assembly process and the independent application of the extensions (see Remark 3) can be parallelized.




Fig. 2. Timings and reference lines for varying number of edge modes I_E (left) and number of subdomains J (right).

4.3. Comparison of costs and timings for a prototype crystal

In this section, we want to discuss the above derived computational costs and compare them to the timings of our implementation. We choose $a = \beta = \kappa = 1$. As the computational domain, we consider a square-shaped prototype photonic crystal structure $\Omega = [0, \sqrt{J}] \times [0, \sqrt{J}]$ decomposed into J unit square subdomains that we refer to as unit cells. For a given mesh size h and order p , considering a quasi-uniform and shape regular triangulation of Ω , we have $N_j = n_e^2 \sim (p/h)^2$ for every subdomain. It is clear that, for this setting, the assumptions from Section 4.2 are fulfilled. Additionally, we consider a mesh such that vertical and horizontal edges have the same (one-dimensional) mesh and thus the eigenvalue problem (15) has to be solved only once (and the cost is ignored); see Remark 2. Since $|\mathcal{E}| \sim J$ and $|\mathcal{V}| \sim J$, and hence $N_A = |\mathcal{E}|I_E + |\mathcal{V}| \sim JI_E$, the total computational costs are given by

$$\underbrace{\mathcal{O}(JI_E C_{ext})}_{\text{basis computation}} + \underbrace{\mathcal{O}(JN_j I_E^2)}_{\text{assembly}} + \underbrace{\mathcal{O}(J^2 I_E^3)}_{\text{solve}}, \quad (22)$$

where we used the fact that the bandwidth of our structured prototype crystal is $b_A \sim 3\sqrt{J}I_E$. The latter is derived by considering the following degrees of freedom (dof) numbering of the ACMS basis: we start with the lower left vertex and number all dofs on vertices and edges increasingly along the x-axis, i.e., we consider only horizontal edges. Next, all dofs on vertical edges are numbered increasingly along the y-axis of the lowest row of the domain. This pattern is repeated for each row. Then the aforementioned bandwidth follows since we have \sqrt{J} unit cells in each direction and I_E basis functions per edge, see also Fig. 1 for an illustration of the sparsity pattern and the dof numbering. For the total costs (22), we observe that, similarly to standard finite element methods, solving the system could potentially be the dominant factor for very large J and I_E . However, for the discrete ACMS method, we expect that N_A is small enough that the assembly and the basis computation are the limiting factors in the computations. This is motivated by the limit case when $I_E = n_e = \sqrt{N_j}$, then, taking into account Remarks 3 and 4, the full cost (22) becomes

$$\underbrace{\mathcal{O}(JI_E^4)}_{\text{basis computation}} + \underbrace{\mathcal{O}(JI_E^4)}_{\text{assembly}} + \underbrace{\mathcal{O}(J^2 I_E^3)}_{\text{solve}}.$$

When doing hp -refinements while keeping the domain decomposition fixed, the number of modes I_E becomes much larger than the number of subdomains J and the cost is dominated by the basis computation and assembly.

Scaling with respect to I_E . In Fig. 2 (left), we present the timings for the computation of the basis functions t_{bas} , the assembly t_{ass} , the time to solve the system t_{sol} , and the sum of these timings t_{tot} . For our computation, we choose $h = 0.01$, $p = 6$, $J = 4$, and we set $I_E = 2^l$ with $l = 1, \dots, 9$. According to (22), we expect a linear scaling $\mathcal{O}(I_E)$ for the basis computation, a quadratic scaling $\mathcal{O}(I_E^2)$ for the assembly and a cubic scaling $\mathcal{O}(I_E^3)$ for solving the system. While the times of the basis computations and the assembly given in Fig. 2 (left) are in accordance with the theory, we see a better scaling for the solving times. Note, that (21) might be pessimistic; see [61]. In our computations we used SciPy to solve the system, which uses the sparse UMFPACK solver; see [62,63]. In practice, considering such a high number of modes might not be necessary due to the fast error reduction; see (11).

Scaling with respect to J . In Fig. 2 (right), we again present the timings for the computation of the basis functions t_{bas} , the assembly t_{ass} , the time to solve the system t_{sol} , and the sum of these timings t_{tot} . For our computation, we choose $h = 0.1$, $p = 4$, $I_E = 4$,




Fig. 3. In a, the domain Ω and the domain decomposition D . In b, the corresponding coarse finite element mesh with $h = 0.1$. The vertices in \mathcal{V} are marked with dots.

and we set $J = l^2$ with $l = 1, \dots, 40$. According to (22) we expect a linear scaling $\mathcal{O}(J)$ for the basis computation and the assembly, and a quadratic scaling $\mathcal{O}(J^2)$ for solving the system. Again, while the timings for the basis computations and the assembly are in accordance to the theory, the solver provided by SciPy now scales with the optimal rate $\mathcal{O}(J^{3/2})$; see [61].

5. Numerical results

The main goal of this section is to exemplify how to choose the different parameters of the discrete ACMS method in order to solve large scale problems. Namely, we will present a study on suitable choices of the mesh size of the (underlying) finite element space, the polynomial order, and the number of edge modes needed to achieve a certain accuracy, depending on the material parameters and the chosen wavenumber. The computations were done using the finite element library Netgen/NGSolve, see [52]. Further, for reproduction purposes, all data can be found in the repository [53].

We consider three different numerical examples. First, in Section 5.1, we investigate the effect of using an underlying higher order FEM for the ACMS basis computation on the method's accuracy in relation to the wavenumber. Then, in Section 5.2, we test a configuration with discontinuous coefficients and, given the flexibility of the ACMS method, we examine how the chosen domain decomposition impacts the accuracy and efficiency of the method. Finally, in Section 5.3, we show the applicability and flexibility of the method on a large but finite-size structure that has defects.

Note that, in the following, we focus exclusively on the L^2 -norm errors. We do not present the corresponding tests in the H^1 -norm, as they align with the theoretical convergence rates mentioned after (11).

5.1. Example 1: circular domain

We consider (1) on the unit disk $\Omega = B_1(0)$ and set the constant functions $a = 1$, $\beta = 1$, $c = 1$, and $f = 0$. We define the boundary source term g such that the problem admits the plane wave $u = e^{-i\kappa(0.6x+0.8y)}$ as its analytical solution, with constant wavenumber $\kappa = \omega$. We mention that this setting is as in [43, Section 5.1]. We choose a domain decomposition D as shown in Fig. 3a, with $J = 8$ subdomains, 12 edges in \mathcal{E} , and 5 vertices in \mathcal{V} . In Fig. 3b, we show a coarse triangulation that is uniformly refined in order to compute the approximated ACMS solution. Below, we will compare the relative L^2 -error, i.e., $\|u - u_A^{h,p}\|_{L^2(\Omega)} / \|u\|_{L^2(\Omega)}$ of the approximated ACMS solution against the analytical solution; this will verify the correct implementation of the high-order ACMS method. In this example, we examine the effects of increasing the maximum number of computable edge modes by either refining the mesh size or driving up the polynomial order of approximation; in both cases, n_e is increased, as commented in Remark 5.

First, we investigate how the degree of approximation of the underlying finite element space, used to approximate the computed ACMS basis functions, affects the accuracy of the ACMS method. In Fig. 4, we show the L^2 -relative errors for different discretization parameters: $h = 0.025$ and $I_{\mathcal{E}} = 32$ (left), $h = 0.0125$ and $I_{\mathcal{E}} = 64$ (middle), and $h = 0.00625$ and $I_{\mathcal{E}} = 128$ (right). We consider an increasing degree of approximation $p = 1, 2, 3, 4, 5$, and repeat the tests for different wavenumbers $\kappa = 16, 32, 64, 128$. On the one hand, we observe that the approximation gets worse for higher wavenumbers. In particular, there is hardly any error reduction for $p = 1$ and $\kappa = 128$, which might be explained by a pollution effect. On the other hand, the error decreases approximately by an order of magnitude if $I_{\mathcal{E}}$ is doubled for each $p \geq 3$ fixed. The observed saturation effect in Fig. 4 occurs, due to the fixed number of edge modes used and the decreasing value of h .

To overcome this saturation effect, it is necessary to employ more edge modes while keeping h fixed. In Fig. 5, we adjust the number of modes based on the degree of approximation, since, as noted in Remark 5, the maximum number of possible edge modes $I_{\mathcal{E}}$ increases for a higher polynomial order. We choose mesh size $h = 0.025$ and the modes considered are $I_{\mathcal{E}} = 32$ for degree $p = 1$, $I_{\mathcal{E}} = 64$ for degree $p = 2, 3$ and $I_{\mathcal{E}} = 128$ for degree $p = 4, 5$. This choice was made because these numbers of modes are the largest values we can compute for each degree of approximation while still being powers of two, ensuring consistency with the previous tests. As expected, we can now see convergence without saturation and correspondingly the full benefit of choosing a high-order




Fig. 4. Example in Section 5.1, on a circular domain with decomposition D depicted in Fig. 3a. The L^2 -relative error $\|u - u_A^{h,p}\|_{L^2(\Omega)} / \|u\|_{L^2(\Omega)}$ is computed against the exact solution for a fixed mesh size and number of edge modes. We choose $h = 0.025$ and $I_E = 32$ (left), $h = 0.0125$ and $I_E = 64$ (middle), $h = 0.00625$ and $I_E = 128$ (right). Test with increasing degree of approximation $p = 1, 2, 3, 4, 5$ (horizontal axis), for different wavenumber values $\kappa = 16, 32, 64, 128$. In the blue region, the values are highlighted for a better comparison with Fig. 5.




Fig. 5. Example in Section 5.1, on a circular domain with decomposition D depicted in Fig. 3a. The L^2 -relative error $\|u - u_A^{h,p}\|_{L^2(\Omega)} / \|u\|_{L^2(\Omega)}$ is computed against the exact solution for mesh size $h = 0.025$ and for different wavenumber values $\kappa = 16, 32, 64, 128$. The modes considered are $I_E = 32$ for degree $p = 1$, $I_E = 64$ for degree $p = 2, 3$ and $I_E = 128$ for degree $p = 4, 5$. In the blue region, the values are highlighted for a better comparison with Fig. 4.

approximation. Edge modes are better approximated because the smooth eigenfunctions, which underlie the ACMS basis functions, are effectively approximated using a p -method. Additionally, the number of computable edge modes increases, which enriches the approximation space of the ACMS method and further reduces the approximation error. To further highlight this observation, we point out that we obtain the same error accuracy if we compare the error in Fig. 4 (right) (where $h = 0.00625$) and in Fig. 5 (where $h = 0.025$), for example, for degree $p = 5$ and $I_E = 128$ (errors within the blue ellipses). This means that, for the coarsest mesh, we obtain already the most accurate solution if we use high polynomial order p and employ a suitable number of modes.

Before proceeding to the next example, we want to verify numerically the dependence between the number of edge modes and the wavenumber that we observed in our theoretical result (11). In Fig. 6, we plot in logarithmic scale the relative L^2 -error for varying number of edge modes $I_E = 1, \dots, 128$ for different wavenumber $\kappa = 16, 32, 64, 128$ and for fixed mesh size $h = 0.025$ and degree of approximation $p = 5$. Qualitatively, the curves show the same behavior. First the error stagnates, and then there is a significant drop after which the asymptotic convergence rate is achieved. Stagnation of error might be explained by requiring a minimal resolution condition, compare this to the situation of piecewise linear approximation, where $h\kappa \approx 1/10$ is required to be able to interpolate a corresponding wave accurately on the given grid. The transition from stagnating errors to monotonically decreasing errors is marked with a black circle. We observe that the required number of modes to enter this monotone behavior almost doubles if the wavenumber is doubled. This might be related to a minimal resolution condition. Moreover, for this example, we may state that the adjoint approximability constant σ^* behaves better than the scaling $\mathcal{O}(C_\kappa \|\kappa\|_\infty^2 / I_E)$ predicted by theory, see the discussion following (10). We conclude that, in this example, when doubling the wavenumber, the errors achieve the predicted convergence rates and enter the asymptotic regime when the number of edge modes is also doubled.

5.2. Example 2: square crystal with circular pores

We consider (1) on a square domain $\Omega = [0, 16] \times [0, 16]$ that is meant to model a finite-size photonic crystal made of silicon with periodically arranged pores (circular holes of radius 0.25 carved in a periodic fashion), as depicted in Fig. 7. The difference in the




Fig. 6. Example in Section 5.1, on a circular domain with decomposition \mathcal{D} depicted in Fig. 3a. The L^2 -relative error $\|u - u_A^{h,p}\|_{L^2(\Omega)} / \|u\|_{L^2(\Omega)}$, shown in logarithmic scale, is computed against the exact solution for a fixed order of approximation $p = 5$ and with mesh size $h = 0.025$. We consider an increasing number of edge modes per edge $I_E = 1, \dots, 128$ for different wavenumber values $\kappa = 16, 32, 64, 128$. The small black circles indicate the index I_E after which the error decays monotonically.




Fig. 7. $\Omega = (0, 16) \times (0, 16)$ with 16×16 square unit cells. In this case the decomposition $\mathcal{D}_{16 \times 16}$ contains only one subdomain Ω .

domain material, silicon for the crystal and air for the pore, is encoded in the discontinuous diffusion coefficient $a = 1/\varepsilon$, where ε is the permittivity with $\varepsilon = 12.1$ in silicon (crystal) and $\varepsilon = 1$ in air (pores) [2]. We note that the units are chosen such that the speed of light is $c = 1$; see [40,41] for similar example configurations. We set $\beta = -1$, $f = 0$, and define the boundary source term as an incoming plane wave that propagates in the x -direction, namely $g = \kappa(1 - (1, 0)^T \cdot \hat{n})e^{-i\kappa x}$, with \hat{n} being the unit outward normal vector, and constant wavenumber $\kappa = \omega$. Since it is of interest to model how light propagates and interacts with the pores within a photonic crystal, we take the wave frequency as a multiple of the length of a unit cell, specifically $\kappa = 1$ in this section.

With no analytical solution available for this example, we show the relative L^2 -error decay $\|u^{h,10} - u_A^{h,p}\|_{L^2(\Omega)} / \|u^{h,10}\|_{L^2(\Omega)}$ of the ACMS approximation against a high-order FEM solution $u^{h,10}$ with degree of approximation $p = 10$, which we refer to as the ground truth solution.

Denoting a 1×1 square that has a circular pore in its center a unit cell, a natural choice for a domain decomposition is to assign to each unit cell a subdomain Ω_j for $j = 1, \dots, 256$, see Fig. 8a.

The ACMS method is, however, applicable for more general domain decompositions, and we next investigate the accuracy and efficiency of the corresponding ACMS approximation. Therefore, we repeat the numerical tests for different subdivisions of the domain Ω as depicted in Fig. 8, which correspond to domain decompositions into $J = 1, 4, 16, 64, 256$ subdomains with number of edges $|\mathcal{E}| = 4, 12, 40, 144, 544$, and number of vertices $|\mathcal{V}| = 4, 9, 25, 81, 289$. We refer to each domain decomposition as $\mathcal{D}_{j \times j}$, where j^2 indicates the number of unit cells per subdomain. As j increases, the number of subdomains decreases.

For each choice of such a domain decomposition, the material coefficients are constant in a suitable neighborhood of the corresponding interface. Therefore, the asymptotic error bound (11) holds in each case. The length of the edges, which influences the

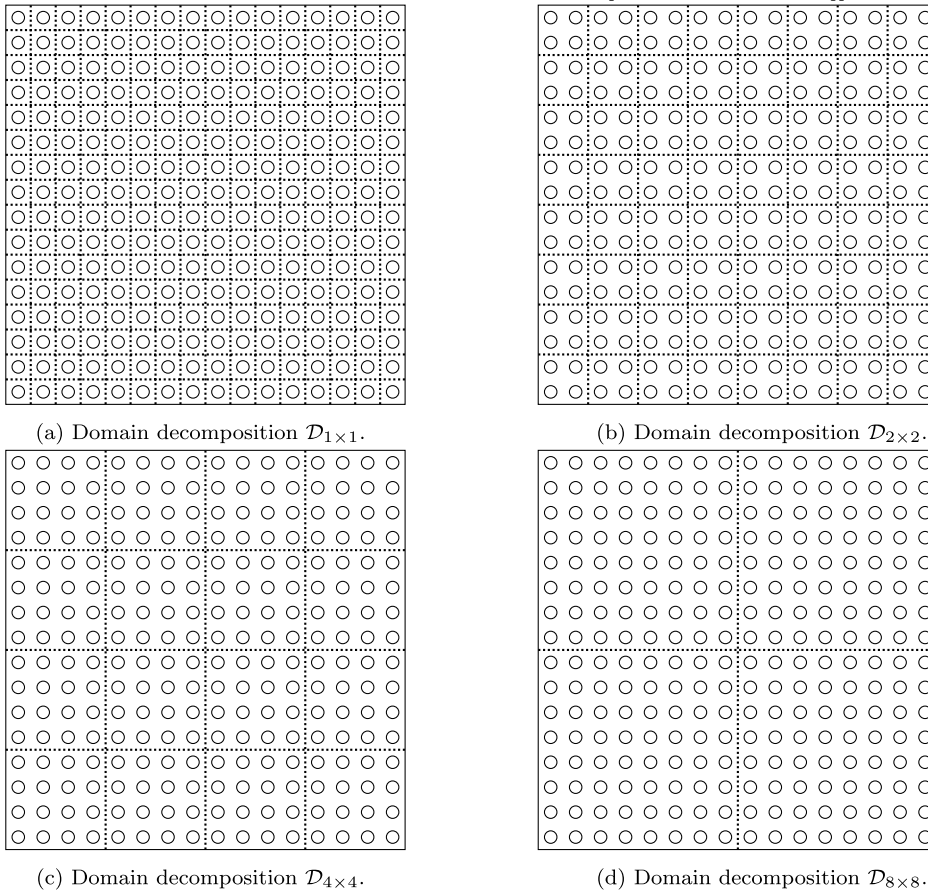


Fig. 8. Ω is a 16×16 square, decomposed into 256 (a), 64 (b), 16 (c) and 4 (d) subdomains.

error estimate (11), varies for each domain decomposition. Hence, in order to counterbalance the effect of the increased edge length on the error, we have to consider twice as many modes in our numerical tests when going from $\mathcal{D}_{j \times j}$ to $\mathcal{D}_{2j \times 2j}$.

In Table 1, we present the approximation error for the different domain decompositions for a fixed mesh size $h = 0.05$ and for an increasing polynomial order $p = 1, 2, 3, 4, 5, 6$. The number of edge modes $I_{\mathcal{E}}$, indicated in brackets, increases with the order p and doubles if the subdomain size doubles. From the data shown in the table, we observe that all domain decompositions $\mathcal{D}_{j \times j}$, $j = 1, 2, 4, 8, 16$ give similar results. Hence, we may conclude that the constants in the error estimate (11) depend only weakly on the chosen domain decomposition. The slight increase in the error for $p = 5$ might be explained by the fact that the respective ACMS spaces for $p = 4, 5, 6$, which employ the same $I_{\mathcal{E}}$ if $\mathcal{D}_{j \times j}$ is fixed, are not nested.

We next discuss the influence on the computational costs when going from $\mathcal{D}_{j \times j}$ to $\mathcal{D}_{2j \times 2j}$ and at the same time doubling the number of modes for accuracy reasons, see (11). We first observe that the corresponding number of subdomains J decreases by 4 and the number of degrees of freedom on an edge n_e roughly doubles, and $N_j = n_e^2$ approximately increases by four. We recall (22)

$$\underbrace{\mathcal{O}(JI_{\mathcal{E}}C_{ext})}_{\text{basis computation}} + \underbrace{\mathcal{O}(JN_jI_{\mathcal{E}}^2)}_{\text{assembly}} + \underbrace{\mathcal{O}(J^2I_{\mathcal{E}}^3)}_{\text{solve}}.$$

The basis computation cost depends on the cost C_{ext} for the extension. Assuming $C_{ext} = \mathcal{O}(J(n_e^4 + I_{\mathcal{E}}n_e^3))$, cf. Remark 3, the resulting basis computation cost of $\mathcal{O}(J^2(I_{\mathcal{E}}n_e^4 + I_{\mathcal{E}}^2n_e^3))$ favors larger J , i.e., smaller j . Similarly, the assembly cost favors smaller $I_{\mathcal{E}}$, i.e., larger J , since JN_j is constant. The cost for solving of $\mathcal{O}(J^2I_{\mathcal{E}}^3)$ favors smaller J , but if it scales only like $\mathcal{O}(J^{3/2}I_{\mathcal{E}}^3)$, cf. Fig. 2, it becomes independent of the domain decomposition, because $J^{3/2}I_{\mathcal{E}}^3$ stays constant. From these considerations, we may conclude that choosing j as small as possible is beneficial. However, parallelization allows to reduce the influence of the basis computation and assembly costs. Moreover, if subdomains have moderate number N_j of degrees of freedom, then cache effects and dense linear algebra may reduce the influence of C_{ext} on the runtime. Since the applicability of sparse direct solvers may be limited by the size of the corresponding linear system, keeping the overall system size small is favorable. Since the system size corresponds to the dimension of the ACMS space $N_A = \mathcal{O}(JI_{\mathcal{E}})$, it is favorable to keep J moderate, i.e., to employ larger subdomains. For these reasons, we continue with the domain decomposition $\mathcal{D}_{2 \times 2}$.

Table 1

Example in Section 5.2, on the square domain Ω with different decompositions depicted in Figs. 7 and 8: $D_{1 \times 1}$, $D_{2 \times 2}$, $D_{4 \times 4}$, $D_{8 \times 8}$, $D_{16 \times 16}$. The L^2 -relative error $\|u^{h,p} - u_A^{h,p}\|_{L^2(\Omega)} / \|u^{h,10}\|_{L^2(\Omega)}$ is computed against the ground truth solution. We choose mesh size $h = 0.05$, degree of approximation $p = 1, 2, 3, 4, 5, 6$, and number of edge modes I_E as indicated in brackets.

	p					
	1	2	3	4	5	6
$D_{1 \times 1}$	$6.9 \cdot 10^{-2}$ (16)	$2.4 \cdot 10^{-4}$ (32)	$6.5 \cdot 10^{-6}$ (32)	$4.1 \cdot 10^{-7}$ (64)	$9.1 \cdot 10^{-7}$ (64)	$7.2 \cdot 10^{-7}$ (64)
$D_{2 \times 2}$	$7.4 \cdot 10^{-2}$ (32)	$2.6 \cdot 10^{-4}$ (64)	$6.5 \cdot 10^{-6}$ (64)	$4.1 \cdot 10^{-7}$ (128)	$8.8 \cdot 10^{-7}$ (128)	$7.0 \cdot 10^{-7}$ (128)
$D_{4 \times 4}$	$7.7 \cdot 10^{-2}$ (64)	$2.8 \cdot 10^{-4}$ (128)	$6.5 \cdot 10^{-6}$ (128)	$4.1 \cdot 10^{-7}$ (256)	$8.7 \cdot 10^{-7}$ (256)	$7.0 \cdot 10^{-7}$ (256)
$D_{8 \times 8}$	$7.9 \cdot 10^{-2}$ (128)	$2.9 \cdot 10^{-4}$ (256)	$6.6 \cdot 10^{-6}$ (256)	$4.1 \cdot 10^{-7}$ (512)	$8.7 \cdot 10^{-7}$ (512)	$7.0 \cdot 10^{-7}$ (512)
$D_{16 \times 16}$	$8.0 \cdot 10^{-2}$ (256)	$2.9 \cdot 10^{-4}$ (512)	$6.6 \cdot 10^{-6}$ (512)	$4.1 \cdot 10^{-7}$ (1024)	$8.6 \cdot 10^{-7}$ (1024)	$6.9 \cdot 10^{-7}$ (1024)




Fig. 9. Example in Section 5.2 on the square domain Ω with domain decomposition $D_{2 \times 2}$, Fig. 8b. The L^2 -relative error of the ACMS approximation is computed against a ground truth solution, for $h = 0.05$ and $p = 6$. We show the convergence rate for an increasing number of edge modes per edge $I_E = 1, \dots, 128$ for different wavenumber values $\kappa = 0.5, 1, 2, 4, 8, 16$.

Finally, besides the dependence on the chosen decomposition, we further want to confirm the linear correlation between the number of modes and the wavenumber that was shown in Fig. 6, in case of a discontinuous diffusion coefficient and a different setting. In this test, we choose the wavenumbers based on an appropriate scaling of the example in [41]: wavelengths between 400 nm and 5000 nm correspond to ω values ranging from 0.628 to 7.85, hence we choose $\kappa \in [0.5, 8]$. We additionally consider $\kappa = 16$ to test the method's limits. In Fig. 9, we show the relative L^2 -error for the domain decomposition $D_{2 \times 2}$, order of approximation $p = 6$ and an increasing number of edge modes $I_E = 1, \dots, 128$. Again, we mark with a black circle the drop after which the error decreases with the predicted rate, and we observe that the number of edge modes required in order to have a monotonic error decay varies linearly with the wavenumber, i.e., when doubling κ we have to double I_E .

5.3. Example 3: large crystal with non periodic pores

In this section we want to show that the ACMS method can be employed to simulate large but finite-size crystals, which possibly contain defects. To do so, we compute the transmission of a “localized” wave that is placed in front of a waveguide-like crystal structure. We consider a rectangular shaped crystal $\Omega = [-J_x/2, J_x/2] \times [-J_y/2, J_y/2]$ consisting of $J_x \times J_y$ unit cells of size $[0, 1] \times [0, 1]$. We choose the same parameter as before, i.e., $f = 0$, $\beta = -1$, $c = 1$, $a = 1$ in the pores and $a = 1/12.1$ outside, and the right hand side

$$g = \begin{cases} i\kappa e^{-i\kappa x} e^{-y^2} & \text{on } [-J_x/2] \times [-J_y/2, J_y/2], \\ 0 & \text{else.} \end{cases}$$

Different to the previous examples, some unit cells consist of silicon only ($a = 1/12.1$), while others include pores with a radius $r = 0.25$ centered within the cell as before. We refer to Fig. 10 for an example of the geometry for $J_x = 10$, $J_y = 10$, where two vertical lines, $\gamma_{in} = [-J_x/2] \times [-3, 3]$ and $\gamma_{out} = [J_x/2] \times [-3, 3]$, are also added on the boundary since they are used for later computations. The




Fig. 10. Geometry of a crystal (green) with pores (red) and a waveguide like structure (depicted by the vertical line pattern) for $J_x = 10$ and $J_y = 10$. Further we plot two lines $\gamma_{in} = [-J_x/2] \times [-3, 3]$ and $\gamma_{out} = [J_x/2] \times [-3, 3]$ in cyan and magenta, respectively. The dotted lines depict the structure of the domain decomposition used for the ACMS computations.




Fig. 11. The norms E_{in} and E_{out} on a crystal structure with $J_x = J_y = 10$ and wavenumber $\kappa \in [0.5, 2]$.

dashed lines depict the $D_{2 \times 2}$ decomposition of the crystal that is considered for all computations. Furthermore, we choose polynomial order $p = 5$, mesh size $h = 0.1$ and number of edge modes per edge $I_E = 16$.

First, we are going to perform a frequency sweep over $\kappa \in [0.5, 2]$. To compare the results we measure the L^2 -norm on the left and right side of the crystal

$$E_{in} := \int_{-J_y/2}^{J_y/2} |u_A^{h,p}(-J_x/2, y)|^2 dy \quad \text{and} \quad E_{out} := \int_{-J_y/2}^{J_y/2} |u_A^{h,p}(J_x/2, y)|^2 dy.$$

Note that E_{out} is related to the time-averaged energy flow through γ_{out} via the complex Poynting vector [2], while E_{in} corresponds to the time-averaged energy flow through γ_{in} up to contributions that are due to the inhomogeneous boundary terms. In the upper plot of Fig. 11, where we have chosen $J_x = J_y = 10$, one can clearly see the wavenumber dependence of E_{in} and E_{out} showing that the transmission pattern varies with wavenumber; see also [41] for an analogous behavior. While some areas exhibit a transmission pattern, i.e., $E_{in} \approx E_{out}$, others suppress transmission, i.e., $E_{out} \ll E_{in}$, e.g., for $\kappa \approx 1$. In addition, we further plot in more detail the values of E_{in} and E_{out} for $\kappa \in [1.2, 1.32]$ and $\kappa \in [1.42, 1.54]$. First, we observe that the smallest amplitude of E_{out} occurs at $\kappa = 1.26$ (within the considered wavenumber range $[0.5, 2]$), indicating poor wave transmission. In contrast, at $\kappa = 1.48$, we have $E_{in} = E_{out}$, suggesting that, at this frequency, the incoming localized wave can pass through the crystal structure, or waveguide, with minimal dissipation. Indeed, in Fig. 12 the absolute value of $u_A^{h,p}$ is plotted and we can observe the expected behavior.

We note that the dimension of the ACMS space is $N_A = 516$ for $J_x = J_y = 10$, while the underlying FEM space has a dimension $N_F = 253926$, i.e., about $N_j \approx 2500$ degrees of freedom per unit cell. For inclusions with simple geometries and moderate wavenumbers such a resolution might not be necessary and finite element methods with lower resolution may still yield accurate solutions for the




Fig. 12. Absolute value $|u_A^{hp}|$ for $\kappa = 1.26$ (left) and $\kappa = 1.48$ (right) for a crystal structure with $J_x = J_y = 10$.

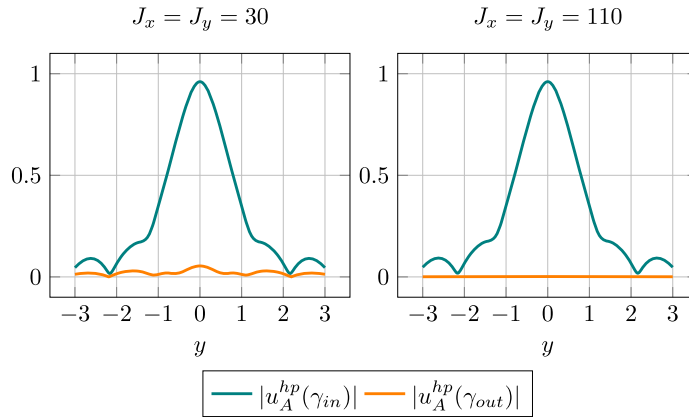


Fig. 13. Absolute value $|u_A^{hp}|$ evaluated on γ_{in} and γ_{out} for a crystal with $J_x = J_y = 30$ (left) and $J_x = J_y = 110$ (right) and a frequency $\kappa = 1.26$.

full crystal problem. However, high resolutions guarantee that the ACMS basis functions are realized accurately. Moreover, highly resolved FEM spaces do not cause computational issues for the ACMS method, because the local systems related to the extension operator remain small. Next we illustrate this observation by increasing the crystal size.

Since the previous computations were performed on a very small crystal, we conclude this section by extending the computations to larger crystal structures. For this purpose, we select two frequencies, $\kappa = 1.26$ and $\kappa = 1.48$, and compute the solution on two crystals, one with $J_x = J_y = 30$ and the other with $J_x = J_y = 110$, i.e., 900 and 12 100 unit cells, respectively. We note that the dimension of the approximation spaces is $N_A = 4096$ and $N_F = 2405776$ for $J_x = J_y = 30$; while $N_A = 52416$ and $N_F = 30636426$ for $J_x = J_y = 110$, which prohibits the solution of the FEM discretization with standard sparse direct solvers on our machines. In Figs. 13 and 14, we plot the absolute value of the solution along the two aforementioned lines, γ_{in} and γ_{out} . Motivated by the findings from the previous example, we expect that for $\kappa = 1.26$ the wave is less transmitted through the crystal, while for $\kappa = 1.48$ we expect that the wave is passing through the crystal with a minimal amount of absorption and possible reflection. This behavior is indeed observed in Fig. 13 and Fig. 14 for both cases $J_x = J_y = 30$ and $J_x = J_y = 110$, where in the latter case the amplitude on γ_{out} for $\kappa = 1.26$ has almost vanished. Additionally, it is clear that the amplitude on γ_{in} is significantly higher for $\kappa = 1.26$ compared to $\kappa = 1.48$, due to greater reflection. The same conclusion can also be drawn for the values E_{in} and E_{out} in Fig. 12 for $\kappa = 1.26$ and $\kappa = 1.48$.

6. Discussion and conclusions

We studied the computational and implementational aspects of the approximate component mode synthesis (ACMS) method. The numerical study has been carried out using the Python interface of the open source NGSolve [52], which allows access to state-of-the-art implementations of high-order finite element methods in a non-intrusive manner. We identified similarities of the ACMS method to usual hp -FEM, which permit efficient assembly and solution procedures, using sparse direct solvers. We gave a full computational complexity analysis, which we verified numerically. From this analysis, it also became clear that the system assembly can be parallelized straightforwardly.

We studied numerically the influence of the underlying hp -FEM approximation on the convergence of the ACMS method, focusing on varying wavenumbers. We found a strong improvement in accuracy when the degree of approximation is larger than one, which




Fig. 14. Absolute value $|u_A^{hp}|$ evaluated on γ_{in} and γ_{out} for a crystal with $J_x = J_y = 30$ (left) and $J_x = J_y = 110$ (right) and a frequency $\kappa = 1.48$.

is in line with the discussion about the pollution effect, see Section 1. We numerically confirmed the theoretical results of [43] for different wavenumbers, see (11), i.e., that the L^2 -error decays cubically in the number of edge modes. We observed that, for increasing wavenumber, the error starts to decrease monotonically for a number of modes that scales linearly in the wavenumber. This is better than expected by the theoretical results given in [43]. Since the error decay is initially very rapid, we may conclude that the ACMS method can achieve engineering relevant accuracy by employing a moderate number of basis functions. If higher accuracy is required, one may use exponentially convergent schemes [42,44,45]. We leave a detailed computational comparison of the different methods for future research.

We further studied the dependence of the error on the employed domain decomposition, and we find a weak dependence. Hence, we may conclude that the ACMS method offers flexibility in choosing the domain decomposition to minimize computational cost, e.g., to exploit fast dense linear algebra. This flexibility, in turn, allowed us to simulate a large but finite-size crystal that has a line defect.

Code availability statement

The code used to generate the findings of this study is openly available in Zenodo at <https://doi.org/10.5281/zenodo.13898694>.

References

- [1] M. Adhikary, R. Uppu, C.A.M. Harteveld, D.A. Grishina, W.L. Vos, Experimental probe of a complete 3D photonic band gap, *Opt. Express* 28 (3) (2020) 2683, <https://doi.org/10.1364/OE.28.002683>.
- [2] J.D. Joannopoulos, S.G. Johnson, J.N. Winn, R.D. Meade, *Molding the Flow of Light*, Princeton University Press, 2008.
- [3] D.M. Callahan, K.A.W. Horowitz, H.A. Atwater, Light trapping in ultrathin silicon photonic crystal superlattices with randomly-textured dielectric incouplers, *Opt. Express* 21 (25) (2013) 30315, <https://doi.org/10.1364/OE.21.030315>.
- [4] D. Zhou, R. Biswas, Photonic crystal enhanced light-trapping in thin film solar cells, *J. Appl. Phys.* 103 (9) (5 2008), <https://doi.org/10.1063/1.2908212>.
- [5] D. Sharma, S.B. Hasan, R. Saive, J.J.W. van der Vegt, W.L. Vos, Enhanced absorption in thin and ultrathin silicon films by 3D photonic band gap back reflectors, *Opt. Express* 29 (25) (2021) 41023, <https://doi.org/10.1364/OE.435412>.
- [6] A.F. Koenderink, L. Bechger, H.P. Schriemer, A. Lagendijk, W.L. Vos, Broadband fivefold reduction of vacuum fluctuations probed by dyes in photonic crystals, *Phys. Rev. Lett.* 88 (14) (2002) 143903, <https://doi.org/10.1103/PhysRevLett.88.143903>.
- [7] M.D. Leistikow, A.P. Mosk, E. Yeganegi, S.R. Huisman, A. Lagendijk, W.L. Vos, Inhibited spontaneous emission of quantum dots observed in a 3D photonic band gap, *Phys. Rev. Lett.* 107 (19) (2011) 193903, <https://doi.org/10.1103/PhysRevLett.107.193903>.
- [8] P. Lodahl, A. Floris van Driel, I.S. Nikolaev, A. Irman, K. Overgaag, D. Vanmaekelbergh, W.L. Vos, Controlling the dynamics of spontaneous emission from quantum dots by photonic crystals, *Nature* 430 (7000) (2004) 654–657, <https://doi.org/10.1038/nature02772>.
- [9] S. Ogawa, M. Imada, S. Yoshimoto, M. Okano, S. Noda, Control of light emission by 3D photonic crystals, *Science* 305 (5681) (2004) 227–229, <https://doi.org/10.1126/science.1097968>.
- [10] W. Dörfler, A. Lechleiter, M. Plum, G. Schneider, C. Wieners, *Photonic Crystals: Mathematical Analysis and Numerical Approximation*, Springer, Basel, Basel, 2011.
- [11] P. Kuchment, An overview of periodic elliptic operators, *Bull. Am. Math. Soc.* 53 (3) (2016) 343–414, <https://doi.org/10.1090/bull/1528>.
- [12] K. Busch, G. von Freymann, S. Linden, S. Mingaleev, L. Tskhelashvili, M. Wegener, Periodic nanostructures for photonics, *Phys. Rep.* 444 (3–6) (2007) 101–202, <https://doi.org/10.1016/j.physrep.2007.02.011>.
- [13] Y.D. Ho, P.S. Ivanov, E. Engin, M.F.J. Nicol, M.P.C. Taverne, Chengyong Hu, M.J. Cryan, I.J. Craddock, C.J. Railton, J.G. Rarity, FDTD simulation of inverse 3-D face-centered cubic photonic crystal cavities, *IEEE J. Quantum Electron.* 47 (12) (2011) 1480–1492, <https://doi.org/10.1109/JQE.2011.2170404>.
- [14] M. Minkov, V. Savona, Automated optimization of photonic crystal slab cavities, *Sci. Rep.* 4 (1) (2014) 5124, <https://doi.org/10.1038/srep05124>.
- [15] A. Taflov, S.C. Hagness, M. Picket-May, *Computational Electromagnetics: The Finite-Difference Time-Domain Method*, in: *The Electrical Engineering Handbook*, Elsevier, 2005, pp. 629–670.
- [16] K. Yee, Numerical solution of initial boundary value problems involving Maxwell's equations in isotropic media, *IEEE Trans. Antennas Propag.* 14 (3) (1966) 302–307, <https://doi.org/10.1109/TAP.1966.1138693>.
- [17] P. Amestoy, R. Brossier, A. Buttari, J.-Y. L'Excellent, T. Mary, L. Métivier, A. Miniussi, S. Operto, Fast 3D frequency-domain full-waveform inversion with a parallel block low-rank multifrontal direct solver: application to OBC data from the North Sea, *Geophysics* 81 (6) (2016) R363–R383, <https://doi.org/10.1190/geo2016-0052.1>.

- [18] S. Operto, P. Amestoy, H. Aghamiry, S. Beller, A. Buttari, L. Combe, V. Dolean, M. Gerest, G. Guo, P. Jolivet, J.-Y. L'Excellent, F. Mamfoumbi, T. Mary, C. Puglisi, A. Ribodetti, P.-H. Tournier, Is 3D frequency-domain FWI of full-azimuth/long-offset OBN data feasible? The Gorgon data FWI case study, *Lead. Edge* 42 (3) (2023) 173–183, <https://doi.org/10.1190/le42030173.1>.
- [19] C. Tong, X. Fang, M. Zhao, Fast high order algorithm for three-dimensional Helmholtz equation involving impedance boundary condition with large wave numbers, *Am. J. Comput. Math.* 13 (02) (2023) 211–229, <https://doi.org/10.4236/ajcm.2023.132011>.
- [20] R. Nicolaides, A method for complex geometries in finite-difference solutions of Maxwell's equations, *Comput. Math. Appl.* 48 (7–8) (2004) 1111–1119, <https://doi.org/10.1016/j.camwa.2004.10.009>.
- [21] F. Ihlenburg, I. Babuška, Finite element solution of the Helmholtz equation with high wave number Part I: the h-version of the FEM, *Comput. Math. Appl.* 30 (9) (1995) 9–37, [https://doi.org/10.1016/0898-1221\(95\)00144-N](https://doi.org/10.1016/0898-1221(95)00144-N).
- [22] I.M. Babuška, S.A. Sauter, Is the pollution effect of the FEM avoidable for the Helmholtz equation considering high wave numbers?, *SIAM J. Numer. Anal.* 34 (6) (1997) 2392–2423, <https://doi.org/10.1137/S0036142994269186>.
- [23] T. Chaumont-Frelet, S. Nicaise, Wavenumber explicit convergence analysis for finite element discretizations of general wave propagation problems, *IMA J. Numer. Anal.* 40 (2) (2020) 1503–1543, <https://doi.org/10.1093/imanum/drz020>.
- [24] H. Chen, P. Lu, X. Xu, A hybridizable discontinuous Galerkin method for the Helmholtz equation with high wave number, *SIAM J. Numer. Anal.* 51 (4) (2013) 2166–2188, <https://doi.org/10.1137/120883451>.
- [25] J.M. Melenk, S. Sauter, Wavenumber explicit convergence analysis for Galerkin discretizations of the Helmholtz equation, *SIAM J. Numer. Anal.* 49 (3) (2011) 1210–1243, <https://doi.org/10.1137/090776202>.
- [26] J.M. Melenk, A. Parsania, S. Sauter, General DG-methods for highly indefinite Helmholtz problems, *J. Sci. Comput.* 57 (3) (2013) 536–581, <https://doi.org/10.1007/s10915-013-9726-8>.
- [27] D. Lafontaine, E.A. Spence, J. Wunsch, Wavenumber-explicit convergence of the hp-FEM for the full-space heterogeneous Helmholtz equation with smooth coefficients, *Comput. Math. Appl.* 113 (2022) 59–69, <https://doi.org/10.1016/J.CAMWA.2022.03.007>.
- [28] X. Feng, H. Wu, Discontinuous Galerkin methods for the Helmholtz equation with large wave number, *SIAM J. Numer. Anal.* 47 (4) (2009) 2872–2896, <https://doi.org/10.1137/080737538>.
- [29] X. Feng, H. Wu, hp-Discontinuous Galerkin methods for the Helmholtz equation with large wave number, *Math. Comput.* 80 (276) (2011) 1997–2024, <https://doi.org/10.1090/S0025-5718-2011-02475-0>.
- [30] T. Strauszer-Caussade, L.M. Faria, A. Fernandez-Lado, C. Pérez-Arancibia, Windowed Green function method for wave scattering by periodic arrays of 2D obstacles, *Stud. Appl. Math.* 150 (1) (2023) 277–315, <https://doi.org/10.1111/sapm.12540>.
- [31] M.J. Gander, H. Zhang, A class of iterative solvers for the Helmholtz equation: factorizations, sweeping preconditioners, source transfer, single layer potentials, polarized traces, and optimized Schwarz methods, *SIAM Rev.* 61 (1) (2019) 3–76, <https://doi.org/10.1137/16M109781X>.
- [32] M.J. Gander, I.G. Graham, E.A. Spence, Applying GMRES to the Helmholtz equation with shifted Laplacian preconditioning: what is the largest shift for which wavenumber-independent convergence is guaranteed?, *Numer. Math.* 131 (3) (2015) 567–614, <https://doi.org/10.1007/s00211-015-0700-2>.
- [33] C.J. Gittelsohn, R. Hiptmair, I. Perugia, Plane wave discontinuous Galerkin methods: analysis of the h-version, *ESAIM: Math. Model. Numer. Anal.* 43 (2) (2009) 297–331, <https://doi.org/10.1051/m2an/2009002>.
- [34] R. Hiptmair, A. Moiola, I. Perugia, A survey of Trefftz methods for the Helmholtz equation, in: *Building Bridges: Connections and Challenges in Modern Approaches to Numerical Partial Differential Equations*, Springer International Publishing, 2016, pp. 237–279.
- [35] L. Yuan, W. Gong, A space–time Trefftz DG scheme for the time-dependent Maxwell equations in anisotropic media, *Math. Comput. Simul.* 211 (2023) 445–469, <https://doi.org/10.1016/j.matcom.2023.04.015>.
- [36] C. Lehrenfeld, P. Stocker, Embedded Trefftz discontinuous Galerkin methods, *Int. J. Numer. Methods Eng.* 124 (17) (2023) 3637–3661, <https://doi.org/10.1002/nme.7258>.
- [37] P.L. Lederer, C. Lehrenfeld, P. Stocker, Trefftz discontinuous Galerkin discretization for the Stokes problem, *Numer. Math.* 156 (3) (2024) 979–1013, <https://doi.org/10.1007/s00211-024-01404-z>.
- [38] Z. Lu, A. Cemeliloglu, J.J.W. Van der Vegt, Y. Xu, Discontinuous Galerkin approximations for computing electromagnetic Bloch modes in photonic crystals, *J. Sci. Comput.* 70 (2) (2017) 922–964, <https://doi.org/10.1007/s10915-016-0270-1>.
- [39] H. Brandsmeier, K. Schmidt, C. Schwab, A multiscale hp-FEM for 2D photonic crystal bands, *J. Comput. Phys.* 230 (2) (2011) 349–374, <https://doi.org/10.1016/J.JCP.2010.09.018>.
- [40] M. Kozon, *Theoretical and computational analysis of wave confinement in periodic media*, Ph.D. thesis University of Twente, Enschede, the Netherlands, 1 2023.
- [41] L.J.C. van Willenswaard, S. Smeets, N. Renaud, M. Schlottbom, J.J.W. van der Vegt, W.L. Vos, Non-utopian optical properties computed of a tomographically reconstructed real photonic nanostructure, *Opt. Express* 32 (18) (2024) 32028, <https://doi.org/10.1364/OE.519464>.
- [42] Y. Chen, T.Y. Hou, Y. Wang, Exponentially convergent multiscale methods for 2D high frequency heterogeneous Helmholtz equations, *Multiscale Model. Simul.* 21 (3) (2023) 849–883, <https://doi.org/10.1137/22M1507802>.
- [43] E. Giammatteo, A. Heinlein, M. Schlottbom, An extension of the approximate component mode synthesis method to the heterogeneous Helmholtz equation, *IMA J. Numer. Anal.* 3 (2023).
- [44] M. Chupeng, C. Alber, R. Scheichl, Wavenumber explicit convergence of a multiscale generalized finite element method for heterogeneous Helmholtz problems, *SIAM J. Numer. Anal.* 61 (3) (2023) 1546–1584, <https://doi.org/10.1137/21M1466748>.
- [45] D. Peterseim, Eliminating the pollution effect in Helmholtz problems by local subscale correction, *Math. Comput.* 86 (305) (2017) 1005–1036, <https://doi.org/10.1090/mcom/3156>.
- [46] D. Peterseim, B. Verfürth, Computational high frequency scattering from high-contrast heterogeneous media, *Math. Comput.* 89 (326) (2020) 2649–2674, <https://doi.org/10.1090/mcom/3529>.
- [47] U.L. Hetmaniuk, R.B. Lehoucq, A special finite element method based on component mode synthesis, *ESAIM: Math. Model. Numer. Anal.* 44 (3) (2010) 401–420, <https://doi.org/10.1051/m2an/2010007>.
- [48] U. Hetmaniuk, A. Klawonn, Error estimates for a two-dimensional special finite element method based on component mode synthesis, *Electron. Trans. Numer. Anal.* 41 (2014) 109–132.
- [49] Z.V. Vardeny, A. Nahata, A. Agrawal, Optics of photonic quasicrystals, *Nat. Photonics* 7 (3) (2013) 177–187.
- [50] A. Heinlein, U. Hetmaniuk, A. Klawonn, O. Rheinbach, The approximate component mode synthesis special finite element method in two dimensions: parallel implementation and numerical results, *J. Comput. Appl. Math.* 289 (2015) 116–133, <https://doi.org/10.1016/j.cam.2015.02.053>.
- [51] I.G. Graham, S.A. Sauter, Stability and finite element error analysis for the Helmholtz equation with variable coefficients, *Math. Comput.* 89 (321) (2019) 105–138, <https://doi.org/10.1090/mcom/3457>.
- [52] J. Schoeberl, NGSolve Contributors, NGSolve/ngsolve: v6.2.2404, <https://doi.org/10.5281/zenodo.13685624>, 9 2024.
- [53] P.L. Lederer, E. Giammatteo, ACMS: release V1.0, <https://doi.org/10.5281/zenodo.13898694>, 10 2024.
- [54] Y. Chen, T.Y. Hou, Y. Wang, Exponential convergence for multiscale linear elliptic PDEs via adaptive edge basis functions, *Multiscale Model. Simul.* 19 (2) (2021) 980–1010, <https://doi.org/10.1137/20M1352922>.
- [55] P. Freese, M. Hauck, D. Peterseim, Super-localized orthogonal decomposition for high-frequency Helmholtz problems, *SIAM J. Sci. Comput.* 46 (4) (2024).
- [56] R. Courant, D. Hilbert, *Methods of Mathematical Physics vol. 1, vol. I*, Interscience Publishers, Inc., New York, N.Y., 1953.
- [57] E. Giammatteo, On the approximate component mode synthesis method for the approximation of the Helmholtz equation, Ph.D. thesis, University of Twente, 2025.

- [58] S.C. Brenner, L.R. Scott, *The Mathematical Theory of Finite Element Methods*, 3rd edition, Texts in Applied Mathematics, vol. 15, Springer, New York, 2008.
- [59] D. Braess, *Finite Elements: Theory, Fast Solvers, and Applications in Solid Mechanics*, Cambridge University Press, 2007.
- [60] A. Quarteroni, A. Valli, *Numerical Approximation of Partial Differential Equations*, vol. 23, Springer Science & Business Media, 2008.
- [61] R.J. Lipton, D.J. Rose, R.E. Tarjan, Generalized nested dissection, *SIAM J. Numer. Anal.* 16 (2) (1979) 346–358, <https://doi.org/10.1137/0716027>.
- [62] T.A. Davis, Algorithm 832, *ACM Trans. Math. Softw.* 30 (2) (2004) 196–199, <https://doi.org/10.1145/992200.992206>.
- [63] P. Virtanen, R. Gommers, T.E. Oliphant, M. Haberland, T. Reddy, D. Cournapeau, E. Burovski, P. Peterson, W. Weckesser, J. Bright, S.J. van der Walt, M. Brett, J. Wilson, K.J. Millman, N. Mayorov, A.R.J. Nelson, E. Jones, R. Kern, E. Larson, C.J. Carey, Í. Polat, Y. Feng, E.W. Moore, J. VanderPlas, D. Laxalde, J. Perktold, R. Cimrman, I. Henriksen, E.A. Quintero, C.R. Harris, A.M. Archibald, A.H. Ribeiro, F. Pedregosa, P. van Mulbregt, SciPy 1.0 Contributors, SciPy 1.0: fundamental algorithms for scientific computing in Python, *Nat. Methods* 17 (3) (2020) 261–272, <https://doi.org/10.1038/s41592-019-0686-2>.