A locally adaptive non-intrusive block reduced basis method for scattering applications using the boundary element method

Philip Edel, Jérôme Simon

A locally adaptive non-intrusive block reduced basis method for scattering applications using the boundary element method

P. Edel[a,b], J. Simon[c,*]

[a] *DTIS, ONERA, Université Paris Saclay, F-91123 Palaiseau - France*
[b] *Sorbonne Université, CNRS, Laboratoire Jacques-Louis Lions, F-75005 Paris - France*
[c] *DEMR, ONERA, Université Paris Saclay, F-91123 Palaiseau - France*

Abstract

In computational electromagnetism, the radar signature of a target usually requires solving repeated electromagnetic scattering problems associated to plane waves illuminating the target with varying frequency and incident angle. When the problem is large scale, strategies based on repeated solver calls usually lead to prohibitive computational costs. This is especially the case when the solver relies on an integral equation discretized using the boundary element method (BEM), as this amounts to solving numerous complex, unsymmetrical and fully populated linear systems. In this work, reduced order models (ROMs) are built in order to rapidly and accurately approximate the solutions for illuminating waves with frequencies and incident angles within bands of interest. In the context of the BEM, the success of a ROM essentially depends on the ability to decouple the frequency from the Green kernels of the underlying integral equation. In this work, we present a methodology for achieving such decoupling that combines the Empirical Interpolation Method (EIM) with the notion of local adaptivity. We use our approximation of the frequency-dependent BEM operator in a locally adaptive non-intrusive reduced basis method. The proposed strategy is fully non-intrusive, in the sense that it only requires the ability to perform matrix-vector products with standard BEM operators. We illustrate our methodology on real-life electromagnetic scattering problems solved by the Combined Field Integral Equation (CFIE) and with matrix-vector products accelerated with the fast-multipole method (FMM).

Keywords: Boundary element method, Model order reduction, Reduced basis method, Combined field integral equation, Frequency sweep
2010 MSC: 65R20, 78M15, 78M34

*Corresponding author
Email address: jerome.simon@onera.fr (J. Simon)

1. Introduction

The Boundary Element Method (BEM) is a well-established numerical method for solving acoustic or electromagnetic scattering problems using integral formulations [1, 2]. Although dedicated methods such as the Fast-Multiple-Method (FMM) [3], Adaptive Cross Approximation (ACA) [4], $\mathcal{H}$-matrices [5] or more recently tensor methods [6] perform well on large problems, reducing the computational cost related to construction and resolution of the BEM linear system, frequency sweep analysis with the BEM is still a highly time and memory-consuming task. This is due to the fact that the BEM system is frequency-dependent and must be re-constructed and solved for each value of the frequency in the band of interest.

Building an efficient reduced order model (ROM) for the BEM is essentially finding a low-dimensional approximation space onto which to project the BEM system. The dimension of the projected system is small compared to the number of degrees of freedom in the BEM. Consequently, the computational burden can be relieved by solving the projected system, rather than the BEM system, at all the desired frequencies and incident angles. The key to the success of ROMs is that the solution of the projected system, despite being small and very efficient to solve, can be (provided that the ROM is adequately built) a very accurate approximation for the much more costly solution of the BEM system [7, 8, 9].

An approach that has shown great success is constructing low-dimensional approximation spaces as Krylov subspaces, iteratively enriched via the Arnoldi or Lanczos process [10, 11, 12]. A very popular alternative for building ROMs is the proper orthogonal decomposition (POD) as in reported in [13]. The POD consists in computing a large number of BEM solutions at different frequencies and incident angles during a so-called *exploration phase* and then performing a singular value decomposition (SVD), retaining only the most dominant modes. A variant is the greedy approach [14, 9, 15], which iteratively builds a so-called *reduced basis* (RB) made of a few BEM solutions that span a low-dimensional approximation subspace. The greedy variant aims at keeping the reduced basis as small as possible, thus minimizing the computational costs, while maintaining a strict control over the error of the ROM via residual-based error estimators.

In order to achieve significant speed-ups with ROMs, it is necessary to be able to decompose the overall computational effort into the so-called *offline* and *online* phases [7, 8]. The offline phase essentially consists in constructing the low-dimensional approximation space and building the ROM. In practice, this phase represents most of the overall computational effort. The online phase consists in evaluating the ROM at all the desired parameter values (i.e., solving the projected system at all the desired frequencies and incident angles). This phase is not computationally expensive, because the projected system has a small size, typically less than a hundred. Solving such a small system, even numerous times at all the desired values of the parameters, can be done very efficiently using direct solvers. Unfortunately, the ability to decompose the computational effort into an offline and an online phase is absolutely non-trivial with parametrized integral equations discretized using the BEM. In order to

obtain an efficient offline/online decomposition, it is necessary to approximate the frequency-dependent kernels of the integral equation using kernels which are independent from the frequency. To address this challenge, series expansion of the Green kernels have been proposed [16, 17]. This yields an approximation of the BEM frequency-dependent matrix expressed as the weighted sum of a few frequency-independent matrix terms, with frequency-dependent weighs known analytically, usually under the form of polynomials in the frequency. The drawback of this method is that the series expansion is only valid in the neighborhood of the expansion point. An alternative is to apply the Empirical Interpolation Method (EIM) [18] to the Green kernels, which recovers a similar approximation for the BEM frequency-dependent matrix, with frequency-dependent weighs obtained though imposing some interpolation constraints [9, 19].

All these methods have similar drawbacks: *(i)* due to the non-trivial dependency in the frequency, the number of frequency-independent matrix terms can be quite large, which can compromise the overall performance of the ROM and *(ii)* the frequency-independent matrix terms do not coincide with standard integral operators, thus new dedicated assembly or matrix-vector product routines must be implemented in order to handle these non-standard operators, which may represent a significant human investment. Some effort has been made to overcome these drawbacks. An attempt to control the number of frequency-independent matrix terms by means of an *hp*-type approach is reported in [9], while an original method for approximating each frequency-independent matrix term as a known integral operator is presented in [20].

In this work, we introduce a local approximation for the BEM frequency-dependent operator, in which the number Q of frequency-independent terms is decided by the user. In opposition to the *hp*-type approach proposed in [9], we achieve this goal relying on a *locally adaptive* approach introduced in Ref. [21]. Each frequency-independent matrix term is a standard integral operator in the fashion of [20], thus our method is non-intrusive, in the sense that there is no need to implement new assembly or matrix-vector product routines. More specifically, our method constructs a decomposition of the frequency window of interest into sub-windows. The BEM frequency-dependent operator is approximated over each sub-window by a so-called *non-intrusive local affine approximation*. Each non-intrusive local affine approximation is a fact a linear combination of Q standard integral operators, thus the two issues *(i)* and *(ii)* previously mentioned are simultaneously addressed.

The paper is structured as follows. First, in Section 2, we introduce our model electromagnetic scattering problem, solved by the Combined Field Integral Equation (CFIE) and approximated using the BEM. In Section 3, we review the Galerkin reduced basis method, relying on the EIM to decouple the frequency from the Green kernels. This detailed review serves to highlight the two announced issues; namely, *(i)* a large number of frequency-independent terms to deal with and *(ii)* the fact that these terms are non-standard integral operators. Section 4 introduces the novel *non-intrusive local affine approximations* and explains how they overcome our two issues. Next, in Section 5, we propose a revisited version of the reduced basis method, specifically tailored for

frequency sweep analysis with the BEM. Finally, Section 6 is devoted to some numerical illustrations on a real-life electromagnetic scattering problem solved with the CFIE.

2. The CFIE for electromagnetic scattering problems

In this work, we consider an electromagnetic plane wave propagating in the direction $\widehat{\boldsymbol{d}} \in \mathbb{R}^3$ with polarization $\widehat{\boldsymbol{p}} \in \mathbb{R}^3$ and frequency f (the hats indicate that the vectors are unitary). The direction depends on the incident angle. More precisely, the direction is expressed in spherical coordinates as $\widehat{\boldsymbol{d}} = (\sin\theta\cos\phi,\, \sin\theta\sin\phi,\, \cos\theta)$, with θ the polar angle and ϕ the azimuthal angle. In order to remember this dependency, we shall denote the direction $\widehat{\boldsymbol{d}}(\nu)$, where $\nu = (\theta,\, \phi)$ shall be abusively designated as the incident angle. Introducing the wavenumber $\mu = \frac{2\pi f}{c}$, where $c \approx 3 \times 10^8$m/s is the speed of light in free space, the plane electric and magnetic waves are respectively given for all $\boldsymbol{x} \in \mathbb{R}^3$ by

$$\boldsymbol{e}^{\mathrm{inc}}(\boldsymbol{x};\mu,\nu) = \widehat{\boldsymbol{p}} e^{i\mu\boldsymbol{x}\cdot\widehat{\boldsymbol{d}}(\nu)}, \quad \boldsymbol{h}^{\mathrm{inc}}(\boldsymbol{x};\mu,\nu) = (\widehat{\boldsymbol{d}}(\nu) \times \widehat{\boldsymbol{p}}) e^{i\mu\boldsymbol{x}\cdot\widehat{\boldsymbol{d}}(\nu)}. \tag{1}$$

Let $\Omega \subset \mathbb{R}^3$ be an bounded domain representing the target (i.e., the scattering object) and denote Γ its boundary. We assume that Γ is smooth enough for the exterior normal $\widehat{\boldsymbol{n}}$ to be well-defined. Due to the presence of the scattering object, there exists a scattered electromagnetic field $(\boldsymbol{e}^{\mathrm{sca}}, \boldsymbol{h}^{\mathrm{sca}})$, such that the total electromagnetic field $(\boldsymbol{e}, \boldsymbol{h})$ made up of the incident and scattered fields following $\boldsymbol{e} = \boldsymbol{e}^{\mathrm{inc}} + \boldsymbol{e}^{\mathrm{sca}}$ and $\boldsymbol{h} = \boldsymbol{h}^{\mathrm{inc}} + \boldsymbol{h}^{\mathrm{sca}}$ satisfies the time-harmonic Maxwell equations with Silver-Muller radiation condition at infinity given by

$$\begin{cases} \operatorname{curl} \boldsymbol{e} - i\mu\boldsymbol{h} = 0 \quad \text{in } \mathbb{R}^3 \setminus \overline{\Omega}, \\ \operatorname{curl} \boldsymbol{h} + i\mu\boldsymbol{e} = 0 \quad \text{in } \mathbb{R}^3 \setminus \overline{\Omega}, \\ \lim\limits_{|\boldsymbol{x}|\to\infty} (\boldsymbol{h}^{\mathrm{sca}}(\boldsymbol{x}) \times \boldsymbol{x} - |\boldsymbol{x}|\boldsymbol{e}^{\mathrm{sca}}(\boldsymbol{x})) = 0. \end{cases} \tag{2}$$

In this work, we shall assume that the scattering object is perfectly conducting thus the tangential component of the total electric field vanishes, yielding the boundary condition $\boldsymbol{e}^{\mathrm{sca}} \times \widehat{\boldsymbol{n}} = -\boldsymbol{e}^{\mathrm{inc}} \times \widehat{\boldsymbol{n}}$ on Γ.

In order to numerically solve the electromagnetic scattering problem, we use the BEM to discretize the weak form associated to two well-known integral equations; namely the Electric Field Integral Equation (EFIE) and the Magnetic Field Integral Equation (MFIE) [2]. First, the surface Γ is meshed with triangle elements. We denote $\mathcal{T}_h$ the set of all triangles in the mesh. Given a triangle $E \in \mathcal{T}_h$, we define the zero-th order local Raviart-Thomas space of complex-valued functions defined on E as $\boldsymbol{RT}_0(E) = \{\boldsymbol{v} : \boldsymbol{x} \in E \mapsto \boldsymbol{\alpha} + \beta\boldsymbol{x} \mid \boldsymbol{\alpha} \in \mathbb{C}^2,\ \beta \in \mathbb{C}\}$. Following the standard BEM, we choose as boundary element approximation space the global Raviart-Thomas space defined by

$$\boldsymbol{V}_h = \{\boldsymbol{v} \in \boldsymbol{H}^0_{\mathrm{div}}(\Gamma) \mid \boldsymbol{v}|_E \in \boldsymbol{RT}_0(E),\ \forall E \in \mathcal{T}_h\}, \tag{3}$$

where $\boldsymbol{H}^0_{\mathrm{div}}(\Gamma) = \{\boldsymbol{v} \in \boldsymbol{L}^2_t(\Gamma) \mid \mathrm{div}_\Gamma\, \boldsymbol{v} \in L^2(\Gamma)\}$ with $L^2(\Gamma)$ the Sobolev space of square-integrable complex-valued functions defined on Γ and $\boldsymbol{L}^2_t(\Gamma)$ the classical Sobolev space comprised of complex-valued functions $\boldsymbol{v} \in \boldsymbol{L}^2(\Gamma) = [L^2(\Gamma)]^3$ that are tangential to Γ, *i.e.*, such that $\boldsymbol{v} \cdot \widehat{\boldsymbol{n}}|_\Gamma = 0$. Note that the boundary element approximation space $\boldsymbol{V}_h$ is finite dimensional. We shall denote $\mathcal{N}_h$ its dimension (equal to the number of triangles in the mesh). It is also worth noting that the discrete approximation space $\boldsymbol{V}_h$ is conforming in the sense $\boldsymbol{V}_h \subset \boldsymbol{H}^{-1/2}_{\mathrm{div}}(\Gamma)$, see *e.g.* [15] for more details on the functional setting.

In a Galerkin context, the discrete EFIE operator $\boldsymbol{T}(\mu)$ is expressed as

$$\begin{aligned} \forall \boldsymbol{v}, \boldsymbol{w} \in \boldsymbol{V}_h, \quad \langle \boldsymbol{T}(\mu)\boldsymbol{v}, \boldsymbol{w}\rangle = i\mu \int_\Gamma \overline{\boldsymbol{w}(\boldsymbol{x})} \cdot \int_\Gamma \mathcal{G}(\boldsymbol{x},\boldsymbol{y};\mu)\boldsymbol{v}(\boldsymbol{y})\mathrm{d}\Gamma_{\boldsymbol{y}}\mathrm{d}\Gamma_{\boldsymbol{x}} \\ -\frac{i}{\mu}\int_\Gamma \mathrm{div}_{\Gamma,\boldsymbol{x}}\, \overline{\boldsymbol{w}(\boldsymbol{x})} \int_\Gamma \mathcal{G}(\boldsymbol{x},\boldsymbol{y};\mu)\mathrm{div}_{\Gamma,\boldsymbol{y}}\, \boldsymbol{v}(\boldsymbol{y})\mathrm{d}\Gamma_{\boldsymbol{y}}\mathrm{d}\Gamma_{\boldsymbol{x}}, \end{aligned} \tag{4}$$

where $\mathcal{G}(\boldsymbol{x},\boldsymbol{y};\mu) = \frac{e^{i\mu|\boldsymbol{x}-\boldsymbol{y}|}}{4\pi|\boldsymbol{x}-\boldsymbol{y}|}$, $\boldsymbol{x} \neq \boldsymbol{y}$, denotes the outgoing fundamental solution at the wavenumber μ and $|\cdot|$ denotes the Euclidian norm in $\mathbb{R}^3$. The associated right-hand side (RHS) is $\boldsymbol{b}^{\mathrm{E}}(\mu,\nu)$ given by

$$\forall \boldsymbol{w} \in \boldsymbol{V}_h, \quad \langle \boldsymbol{b}^{\mathrm{E}}(\mu,\nu), \boldsymbol{w}\rangle = \int_\Gamma \boldsymbol{e}^{\mathrm{inc}}(\mu,\nu) \cdot \overline{\boldsymbol{w}}\mathrm{d}\Gamma. \tag{5}$$

Notice that from the expression (1) of the incident electric field, the RHS depends on both the wavenumber and incident angle.

Still in the Galerkin context, the discrete MFIE operator $\frac{1}{2}\boldsymbol{I} + \boldsymbol{K}(\mu)$ is given by

$$\begin{aligned} \forall \boldsymbol{v}, \boldsymbol{w} \in \boldsymbol{V}_h, \quad \langle \left(\frac{1}{2}\boldsymbol{I} + \boldsymbol{K}(\mu)\right)\boldsymbol{v}, \boldsymbol{w}\rangle = \frac{1}{2}\int_\Gamma \overline{\boldsymbol{w}(\boldsymbol{x})} \cdot \boldsymbol{v}(\boldsymbol{x})\mathrm{d}\Gamma_{\boldsymbol{x}} \\ + \int_\Gamma \overline{\boldsymbol{w}(\boldsymbol{x})} \cdot \left(\widehat{\boldsymbol{n}}(\boldsymbol{y}) \times \int_\Gamma \partial_{\widehat{\boldsymbol{n}}(\boldsymbol{y})}\mathcal{G}(\boldsymbol{x},\boldsymbol{y};\mu) \times \boldsymbol{v}(\boldsymbol{y})\mathrm{d}\Gamma_{\boldsymbol{y}}\right)\mathrm{d}\Gamma_{\boldsymbol{x}}. \end{aligned} \tag{6}$$

The associated RHS is $\boldsymbol{b}^{\mathrm{M}}(\mu,\nu)$ given by

$$\forall \boldsymbol{w} \in \boldsymbol{V}_h, \quad \langle \boldsymbol{b}^{\mathrm{M}}(\mu,\nu), \boldsymbol{w}\rangle = \int_\Gamma \widehat{\boldsymbol{n}} \times \boldsymbol{h}^{\mathrm{inc}}(\mu,\nu) \cdot \overline{\boldsymbol{w}}\mathrm{d}\Gamma. \tag{7}$$

This being set the discrete CFIE operator is given by $\boldsymbol{A}(\mu) = (1-c)\boldsymbol{T}(\mu) + c(\frac{1}{2}\boldsymbol{I} + \boldsymbol{K}(\mu))$ with $c \in]0,1[$ and the associated right-hand side is expressed as $\boldsymbol{b}(\mu,\nu) = (1-c)\boldsymbol{b}^{\mathrm{E}}(\mu,\nu) + c\boldsymbol{b}^{M}(\mu,\nu)$. In this work, we set $c = 0.5$. Given a basis $\{\boldsymbol{\phi}_j\}_{1\leq j\leq \mathcal{N}_h}$ for $\boldsymbol{V}_h$ the discrete CFIE problem writes as the linear system of equations

$$\mathbf{A}(\mu)\mathbf{u}(\mu,\nu) = \mathbf{f}(\mu,\nu), \tag{8}$$

where $\mathbf{A}(\mu) \in \mathbb{C}^{\mathcal{N}_h \times \mathcal{N}_h}$ is the fully-populated, non-hermitian and wavenumber-dependent matrix with coefficients $\langle \boldsymbol{A}(\mu)\boldsymbol{\phi}_j, \boldsymbol{\phi}_i\rangle$ for $1 \leq i,j \leq \mathcal{N}_h$ and $\mathbf{f}(\mu,\nu) \in \mathbb{C}^{\mathcal{N}_h}$ is the RHS vector with coefficients $\langle \boldsymbol{b}(\mu,\nu), \boldsymbol{\phi}_j\rangle$ for $1 \leq j \leq \mathcal{N}_h$.

Thanks to the Stratton-Chu integral representation formulas [1], the solution vector $\mathbf{u}(\mu,\nu) \in \mathbb{C}^{\mathcal{N}_h}$ can be used to recover an approximation for the total electromagnetic fields anywhere in the exterior domain $\mathbb{R}^3 \setminus \overline{\Omega}$. In particular, the radar cross section (RCS) can be derived from the amplitude of the electric field infinitely far away from the target.

3. Review of the RB method

Under the paradigm of direct solvers, solving Eq.(8) requires $\mathcal{O}(\mathcal{N}_h^2)$ operations to assemble the matrix, $\mathcal{O}(\mathcal{N}_h^3)$ operations to compute the LU-decomposition and finally $\mathcal{O}(\mathcal{N}_h^2)$ to perform a forward-backward triangular process for obtaining the solution. Noticing that the same LU-decomposition can be used to solve multiple right-hand sides, the overall complexity for solving ℓ incident angles at a fixed value of the wavenumber is $\mathcal{O}(\mathcal{N}_h^3 + \ell\mathcal{N}_h^2)$. In the frequency sweep context, a new assembly phase and a new LU-decomposition is required for each and every value of μ over the range of interest $[\mu_{\min}, \mu_{\max}]$, which may represents a time and memory-consuming task. This discussion remains valid under the paradigm of iterative solvers, because repeated calls to iterative solvers quickly lead to prohibitive computational costs. In this work, we turn to model order reduction techniques [22, 7, 8] in view of reducing the overall costs of such simulation campaigns.

3.1. RB approximation

First, let us choose a norm in which to measure the solutions to Eq. (8). Namely, this norm will be denoted $\|\cdot\|_{\mathbf{H}}$ and be given by $\|\mathbf{u}\|_{\mathbf{H}} = \sqrt{\mathbf{u}^*\mathbf{H}\mathbf{u}}$ where the superscript $*$ designates the conjugate transpose operation and $\mathbf{H} \in \mathbb{C}^{\mathcal{N}_h \times \mathcal{N}_h}$ is a user-defined hermitian, positive-definite matrix. For instance, $\mathbf{H}$ could be the identity matrix, in which case $\|\cdot\|_{\mathbf{H}}$ is the standard Euclidian norm. Another relevant choice for $\mathbf{H}$ is the mass matrix with coefficients $\int_\Gamma \boldsymbol{\phi}_j \cdot \overline{\boldsymbol{\phi}}_i \mathrm{d}\Gamma$ for $1 \leq i, j \leq \mathcal{N}_h$, in which case the $\|\cdot\|_{\mathbf{H}}$ norm is the $\boldsymbol{L}^2(\Gamma)$ norm. The choice of norm is discussed more extensively in [15].

The reduced basis (RB) method [23, 24] consists in computing few, say $N \ll \mathcal{N}_h$, solutions to Eq. (8) at N chosen points $(\mu^{(1)}, \nu^{(1)}), \ldots, (\mu^{(N)}, \nu^{(N)})$. Using a standard Gram-Schmidt procedure, the solutions $\mathbf{u}(\mu^{(1)}, \nu^{(1)}), \ldots, \mathbf{u}(\mu^{(N)}, \nu^{(N)})$ are $\mathbf{H}$-orthonormalized and concatenated in a matrix $\mathbf{P} \in \mathbb{C}^{\mathcal{N}_h \times N}$ such that

$$\mathrm{Colspan}(\mathbf{P}) = \mathrm{Span}\{\mathbf{u}(\mu^{(1)}, \nu^{(1)}), \ldots, \mathbf{u}(\mu^{(N)}, \nu^{(N)})\} \quad \text{and} \quad \mathbf{P}^*\mathbf{H}\mathbf{P} = \mathbf{I}. \tag{9}$$

For all wavenumber μ and incident angle ν, a RB approximation $\mathbf{u}_N(\mu,\nu) \in \mathbb{C}^{\mathcal{N}_h}$ is defined as

$$\mathbf{u}_N(\mu,\nu) = \mathbf{P}\mathbf{x}_N(\mu,\nu), \tag{10}$$

where $\mathbf{x}_N(\mu,\nu) \in \mathbb{C}^N$ is the solution to the following projected problem

$$\mathbf{P}^*\mathbf{A}(\mu)\mathbf{P}\mathbf{x}_N(\mu,\nu) = \mathbf{P}^*\mathbf{f}(\mu,\nu). \tag{11}$$

The projected problem Eq. (11) is of size $N \times N$. Recalling that $N \ll \mathcal{N}_h$, this is a major reduction compared to the BEM linear system Eq. (8), which is $\mathcal{N}_h \times \mathcal{N}_h$. However, at this stage, the solution $\mathbf{x}_N(\mu,\nu) \in \mathbb{C}^N$ of the projected problem cannot be computed with a complexity independent from the number of degrees of freedom $\mathcal{N}_h$, because assembling the projected system matrix $\mathbf{P}^*\mathbf{A}(\mu)\mathbf{P}$ requires having computed the N matrix-vector products $\mathbf{A}(\mu)\mathbf{P}$. Similarly, assembling the projected RHS $\mathbf{P}^*\mathbf{f}(\mu,\nu)$ requires having assembled the RHS $\mathbf{f}(\mu,\nu)$ beforehand.

The computational advantage of the RB method becomes clear in the situation where the matrix and RHS are affine with respect to the varying parameters (μ, ν) [25, 26], that is

$$\mathbf{A}(\mu) = \sum_{q=1}^{Q^a} \theta_q^a(\mu)\mathbf{A}_q, \quad \mathbf{f}(\mu,\nu) = \sum_{q=1}^{Q^f} \theta_q^f(\mu,\nu)\mathbf{f}_q, \tag{12}$$

where the $\mathbf{A}_q$'s are Q^a wavenumber-independent $\mathcal{N}_h \times \mathcal{N}_h$ matrices; the $\mathbf{f}_q$'s are Q^f $\mathcal{N}_h$-dimensional vectors independent from μ and from ν; the θ_q^a and θ_q^f's are respectively μ-dependent and (μ,ν)-dependent complex coefficients. Indeed, in this situation the projected system matrix and RHS write

$$\mathbf{P}^*\mathbf{A}(\mu)\mathbf{P} = \sum_{q=1}^{Q^a} \theta_q^a(\mu)\boxed{\mathbf{P}^*\mathbf{A}_q\mathbf{P}}, \quad \mathbf{P}^*\mathbf{f}(\mu,\nu) = \sum_{q=1}^{Q^f} \theta_q^f(\mu,\nu)\boxed{\mathbf{P}^*\mathbf{f}_q}. \tag{13}$$

The boxed quantities are independent from μ and from ν; thus they can be computed once and for all during the so-called *offline* phase. Next, the projected linear system Eq. (11) can be assembled for any value of μ and ν with $\mathcal{O}(N^2Q^a + NQ^f)$ operations during the so-called *online* phase. With this offline/online strategy, the reduced linear system Eq. (11) can be efficiently assembled and solved during the online stage with an operation count completely independent from the number of degrees of freedom $\mathcal{N}_h$.

3.2. The need for non-intrusive local affine approximations

Inspection of Eqs. (4) and (6) reveals that the discrete EFIE and MFIE operators are non-affine because the Green kernel couples the spatial variables and the wavenumber. A well-known strategy consists in recovering affine approximations by applying the EIM to the wavenumber-dependent kernel [9, 19]. Observe that

$$\mathcal{G}(\boldsymbol{x},\boldsymbol{y};\mu) = \frac{e^{i\mu\xi}-1}{4\pi\xi} + \frac{1}{4\pi\xi}, \quad \xi = |\boldsymbol{x}-\boldsymbol{y}|. \tag{14}$$

We successively apply the EIM to the two functions $g_1^{\mathrm{ns}}(\xi;\mu) = i\mu\frac{e^{i\mu\xi}-1}{4\pi\xi}$ and $g_2^{\mathrm{ns}}(\xi;\mu) = \frac{-i}{\mu}\frac{e^{i\mu\xi}-1}{4\pi\xi}$. For ease of notation, we now use the notation $\star = 1, 2$. The EIM yields $M_\star \geq 1$ so-called *EIM basis functions* $h_1^{g\star}, \ldots, h_{M_\star}^{g\star}$ defined on $[0, \xi_{\max}]$, interpolation points $\{\xi_m^{g\star}\}_{1\leq m\leq M_\star}$ and a lower triangular *interpolation matrix* $\mathbf{B}^{g\star} \in \mathbb{C}^{M_\star \times M_\star}$ with unity diagonal. We refer the reader to the reference

[18] for a detailed explanation on how these quantities are determined. The EIM interpolant is given by

$$g_\star^{\rm ns}(\xi;\mu) \approx \widetilde{g_\star^{\rm ns}}(\xi;\mu) = \sum_{m=1}^{M_\star} \varsigma_m^{g\star}(\mu) h_m^{g\star}(\xi), \tag{15}$$

with complex coefficients $\boldsymbol{\varsigma}^{g\star}(\mu) = (\varsigma_1^{g\star}(\mu), \ldots, \varsigma_{M_\star}^{g\star}(\mu))^T \in \mathbb{C}^{M_\star}$ solution to the $M_\star \times M_\star$ linear system

$$\mathbf{B}^{g\star}\boldsymbol{\varsigma}^{g\star}(\mu) = \boldsymbol{\lambda}^{g\star}(\mu), \tag{16}$$

where $\boldsymbol{\lambda}^{g\star}(\mu) = (g_\star^{\rm ns}(\xi_1^{g\star};\mu), \cdots, g_\star^{\rm ns}(\xi_{M_\star}^{g\star};\mu))^T \in \mathbb{C}^{M_\star}$. Replacing $g_\star^{\rm ns}$ by its EIM approximation $\widetilde{g_\star^{\rm ns}}$ for $\star = 1, 2$ in the expression of the EFIE operator Eq. (4) yields an affine approximation $\widetilde{\boldsymbol{T}}(\mu)$ for the EFIE operator given by

$$\begin{aligned}
\langle \widetilde{\boldsymbol{T}}(\mu)\boldsymbol{v}, \boldsymbol{w}\rangle = & \sum_{m=1}^{M_1} \varsigma_m^{g1}(\mu) \int_\Gamma \overline{\boldsymbol{w}(\boldsymbol{x})} \cdot \int_\Gamma h_m^{g1}(|\boldsymbol{x}-\boldsymbol{y}|)\boldsymbol{v}(\boldsymbol{y}) \mathrm{d}\Gamma_{\boldsymbol{y}} \mathrm{d}\Gamma_{\boldsymbol{x}} \\
& + i\mu \int_\Gamma \overline{\boldsymbol{w}(\boldsymbol{x})} \cdot \int_\Gamma \frac{1}{4\pi|\boldsymbol{x}-\boldsymbol{y}|}\boldsymbol{v}(\boldsymbol{y}) \mathrm{d}\Gamma_{\boldsymbol{y}} \mathrm{d}\Gamma_{\boldsymbol{x}} \\
& + \sum_{m=1}^{M_2} \varsigma_m^{g2}(\mu) \int_\Gamma \mathrm{div}_{\Gamma,\boldsymbol{x}}\, \overline{\boldsymbol{w}(\boldsymbol{x})} \int_\Gamma h_m^{g2}(|\boldsymbol{x}-\boldsymbol{y}|) \mathrm{div}_{\Gamma,\boldsymbol{y}}\, \boldsymbol{v}(\boldsymbol{y}) \mathrm{d}\Gamma_{\boldsymbol{y}} \mathrm{d}\Gamma_{\boldsymbol{x}} \\
& - \frac{i}{\mu} \int_\Gamma \mathrm{div}_{\Gamma,\boldsymbol{x}}\, \overline{\boldsymbol{w}(\boldsymbol{x})} \int_\Gamma \frac{1}{4\pi|\boldsymbol{x}-\boldsymbol{y}|} \mathrm{div}_{\Gamma,\boldsymbol{y}}\, \boldsymbol{v}(\boldsymbol{y}) \mathrm{d}\Gamma_{\boldsymbol{y}} \mathrm{d}\Gamma_{\boldsymbol{x}}.
\end{aligned} \tag{17}$$

Remarking that all the integrated terms are independent from the wavenumber; it is clear that $\widetilde{\boldsymbol{T}}(\mu)$ is affine with $(M_1 + M_2 + 2)$ terms. We can address the MFIE operator in a similar way, by observing that

$$\partial_{\widehat{\boldsymbol{n}}(\boldsymbol{y})}\mathcal{G}(\boldsymbol{x},\boldsymbol{y};\mu) = (\psi^{\rm ns}(|\boldsymbol{x}-\boldsymbol{y}|;\mu) + \psi^{\rm s}(|\boldsymbol{x}-\boldsymbol{y}|)) \frac{(\boldsymbol{y}-\boldsymbol{x})\cdot\boldsymbol{n}(\boldsymbol{y})}{|\boldsymbol{x}-\boldsymbol{y}|}, \tag{18}$$

with $\psi^{\rm ns}(\cdot;\mu)$ and $\psi^{\rm s}$ defined by

$$\psi^{\rm ns}(\xi;\mu) = i\mu\frac{e^{i\mu\xi}-1}{4\pi\xi} - \frac{e^{i\mu\xi}-1-i\mu\xi}{4\pi\xi^2}, \quad \psi^{\rm s}(\xi) = -\frac{1}{4\pi\xi^2}. \tag{19}$$

An affine approximation for the MFIE operator defined by Eq. (6) can be straightforwardly obtained from applying a third EIM to the function $g_3^{\rm ns} = \psi^{\rm ns}$. With this strategy, we obtain an affine approximation for the MFIE operator with $M_3 + 2$ terms (*i.e.*, $M_3 + 1$ terms for $\boldsymbol{K}(\mu)$ and an additionnal term for the mass term $\frac{1}{2}\boldsymbol{I}$). Thus, we obtain an affine approximation for the CFIE operator with $Q^a = (M_1 + M_2 + M_3 + 4)$ terms under the form

$$\widetilde{\mathbf{A}}(\mu) = \sum_{q=1}^{Q^a} \sigma_q(\mu)\mathbf{A}_q \tag{20}$$

with complex coefficients $\boldsymbol{\sigma}(\mu) = (\sigma_1(\mu), \ldots, \sigma_{Q^a}(\mu)) \in \mathbb{C}^{Q^a}$ solution to the linear system $\mathbf{B}\boldsymbol{\sigma}(\mu) = \boldsymbol{\lambda}(\mu)$, where

$$\mathbf{B} = \begin{pmatrix} \mathbf{B}^{g1} & & & & & & & \\ & 1 & & & & & & \\ & & \mathbf{B}^{g2} & & & & & \\ & & & 1 & & & & \\ & & & & \mathbf{B}^{g3} & & & \\ & & & & & 1 & & \\ & & & & & & 1 & \end{pmatrix}, \quad \boldsymbol{\lambda}(\mu) = \begin{pmatrix} \boldsymbol{\lambda}^{g1}(\mu) \\ i\mu \\ \boldsymbol{\lambda}^{g2}(\mu) \\ \frac{-i}{\mu} \\ \boldsymbol{\lambda}^{g3}(\mu) \\ 1 \\ \frac{1}{2} \end{pmatrix}. \tag{21}$$

Notice that $\mathbf{B} \in \mathbb{C}^{Q^a \times Q^a}$ is lower triangular with unity diagonal. Remark also that the last row of the system is associated to the mass term in the MFIE operator. At this stage, two main issues arise:

- *Intrusivity issue:* the wavenumber-independent operator terms $\mathbf{A}_q$, $1 \leq q \leq Q^a$ in the affine decomposition Eq. (20) are non-standard integral operators (this is clear in Eq. (17)). As reported in [20], it is usually necessary to implement new routines within the computational code for performing matrix-vector products with these non-standard integral operators. Implementation of the matrix-vector product routines is a difficult task, because special treatment is required to adequately integrate the singular kernels.

- *Bandwidth issue:* the number of affine terms Q^a grows with the frequency band of interest and can be quite large, which can compromise the overall efficiency of the RB method, as reported in [9].

In the rest of this work, we circumvent the intrusivity issue by proposing a non-intrusive approach and we address the bandwidth issue by constructing multiple local affine approximations, each valid over a specific sub-window following the locally adaptive paradigm introduced in [21].

4. Non-intrusive local affine approximations

In the locally adaptive method [21], a number of terms $Q \geq 2$ is prescribed by the user. A set of $J \geq Q$ wavenumbers $\hat{\mu}_1 \leq \cdots \leq \hat{\mu}_J$ inside a given window of interest $[\mu_{\min}, \mu_{\max}]$ is selected (the selection procedure will be described shortly). Upon these J wavenumbers, one defines the $K = J - Q + 1$ sets

$$\mathcal{T}_k = \{\hat{\mu}_k, \ldots, \hat{\mu}_{k+Q-1}\}, \quad k \in \{1, \ldots, K\}, \tag{22}$$

each of cardinality Q. The *indicator function* $\mathcal{I} : [\mu_{\min}, \mu_{\max}] \to \{1, \ldots, K\}$ maps each wavenumber μ to the index k such that $\mathcal{T}_k$ is the set of the Q points among the $\hat{\mu}_j$'s that are closest to μ.

This being set, we introduce for all $k \in \{1, \ldots, K\}$ the k^{th} *non-intrusive local affine approximation* as

$$\mathbf{A}_k(\mu) = \sum_{q=1}^{Q} \theta_q^k(\mu) \mathbf{A}(\hat{\mu}_{k+q-1}), \tag{23}$$

where the θ_q^k's are wavenumber-dependent coefficients. As we shall see, $\mathbf{A}_k(\mu)$ will only be a good affine approximation for $\mathbf{A}(\mu)$ *locally* for values of μ in the sub-window $\mathcal{D}_k = \mathcal{I}^{-1}(k) = \{\mu \,|\, \mathcal{I}(\mu) = k\}$. We now explain the construction process in detail.

4.1. The wavenumber-dependent coefficients

First, we explain how the wavenumber-dependent coefficients θ_q^k in Eq. (23) are defined given J available wavenumbers $\hat{\mu}_1 \leq \cdots \leq \hat{\mu}_J$. For this purpose, let $\mu \in [\mu_{\min}, \mu_{\max}]$ and denote $k = \mathcal{I}(\mu)$. Then $\boldsymbol{\theta}^k(\mu) = (\theta_1^k(\mu), \ldots, \theta_Q^k(\mu))^T \in \mathbb{C}^Q$ is defined by

$$\boldsymbol{\theta}^k(\mu) = \underset{\boldsymbol{\theta} \in \mathbb{C}^Q}{\operatorname{argmin}} \left\| \boldsymbol{\sigma}(\mu) - \sum_{q=1}^{Q} \theta_q \boldsymbol{\sigma}(\hat{\mu}_{k+q-1}) \right\|_2, \tag{24}$$

with $\|\cdot\|_2$ denoting the Euclidian norm in $\mathbb{C}^{Q^a}$. Equivalently, the wavenumber-dependent coefficients satisfy

$$P_{\mathcal{T}_k}\left[\boldsymbol{\sigma}(\mu)\right] = \sum_{q=1}^{Q} \theta_q^k(\mu) \boldsymbol{\sigma}(\hat{\mu}_{k+q-1}), \tag{25}$$

where $P_{\mathcal{T}_k}[\cdot]$ denotes orthogonal projection from $\mathbb{C}^{Q^a}$ onto the Q-dimensional subspace spanned by the Q vectors $\{\boldsymbol{\sigma}(\hat{\mu}),\, \hat{\mu} \in \mathcal{T}_k\}$.

4.2. Construction using a localization procedure

The locally adaptive method provides an automatic procedure for selecting the wavenumbers $\hat{\mu}_1 \leq \cdots \leq \hat{\mu}_J$. There are two phases [21] : phase 1 selects $Q+1$ wavenumbers using a classical greedy strategy and phase 2 selects more wavenumbers following a locally adaptive strategy until a prescribed tolerance is reached for the worst projection error. We review each phase in detail.

Phase 1. The first phase selecting $Q+1$ wavenumber following by a classical greedy procedure driven by the projection error. Namely, at iteration $J \geq 1$ a set $\mathcal{C}_J$ of J wavenumbers is available. Thus, the vector $\boldsymbol{\sigma}(\mu)$ and its orthogonal projection $P_{\mathcal{C}_J}\left[\boldsymbol{\sigma}(\mu)\right]$ onto the J-dimensional subspace $\operatorname{ColSpan}\{\boldsymbol{\sigma}(\hat{\mu}),\, \hat{\mu} \in \mathcal{C}_J\}$ can be computed for all $\mu \in \Xi$ (where $\Xi \subset [\mu_{\min}, \mu_{\min}]$ is a discrete set). Therefore, the maximizer of the projection error $\|\boldsymbol{\sigma}(\mu) - P_{\mathcal{C}_J}\left[\boldsymbol{\sigma}(\mu)\right]\|_2$ can be easily found by enumeration. The set $\mathcal{C}_J$ is then enriched by adding the maximizer wavenumber $\hat{\mu}^\star \in \Xi$. This greedy selection procedure continues until $Q+1$ wavenumbers are selected. This procedure is summarized by Alg. 1.

Algorithm 1 Classical greedy (phase 1 of localization procedure)

Choose a prescribed number of term Q and a discrete set $\Xi \subset [\mu_{\min}, \mu_{\max}]$
Pick a random $\hat{\mu}^\star \in \Xi$
Set $\mathcal{C}_1 = \{\hat{\mu}^\star\}$
for $J = 1, \dots Q$ **do**
    Find $\hat{\mu}^\star = \underset{\mu \in \Xi}{\operatorname{argmax}} \| \boldsymbol{\sigma}(\mu) - P_{\mathcal{C}_J} [\boldsymbol{\sigma}(\mu)] \|_2$
    Enrich $\mathcal{C}_{J+1} = \mathcal{C}_J \cup \{\hat{\mu}^\star\}$
end for

Phase 2. At the start of the second phase, $Q + 1$ wavenumbers are available from the first phase. Thus, Eq. (22) defines two sets $\mathcal{T}_k$, $k = 1, 2$. In this context, the indicator function maps each μ to the integer $k = \mathcal{I}(\mu)$ such that the projection of $\boldsymbol{\sigma}(\mu)$ must be performed onto the Q-dimensional subspace $\mathrm{ColSpan}\{\boldsymbol{\sigma}(\hat{\mu}), \hat{\mu} \in \mathcal{T}_k\}$. Thus, for any value of μ, the *local* projection error is given by $\| \boldsymbol{\sigma}(\mu) - P_{\mathcal{T}_k} [\boldsymbol{\sigma}(\mu)] \|_2$, with $k = \mathcal{I}(\mu)$. The locally adaptive strategy, summarized by Alg. 2, consists in selecting the wavenumbers that maximize the local projection error until a prescribed tolerance is reached on the maximal local projection error.

Algorithm 2 Locally adaptive strategy (phase 2 of localization procedure)

Choose a prescribed tolerance $tol > 0$
Obtain a set $\mathcal{C}_J = \{\hat{\mu}_j\}_{1 \leq j \leq J}$ with $J = Q + 1$ by Alg. 1
Set $K = 2$ and find $\mu^\star = \underset{\mu \in \Xi}{\operatorname{argmax}} \| \boldsymbol{\sigma}(\mu) - P_{\mathcal{T}_k} [\boldsymbol{\sigma}(\mu)] \|_2$, where $k = \mathcal{I}(\mu)$
Compute $\epsilon = \| \boldsymbol{\sigma}(\mu^\star) - P_{\mathcal{T}_{k^\star}} [\boldsymbol{\sigma}(\mu^\star)] \|_2$, where $k^\star = \mathcal{I}(\mu^\star)$
while $\epsilon > tol$ **do**
    Enrich $\mathcal{C}_{J+1} = \mathcal{C}_J \cup \{\hat{\mu}^\star\}$
    Update $J \leftarrow J + 1$ and $K \leftarrow K + 1$
    Find $\mu^\star = \underset{\mu \in \Xi}{\operatorname{argmax}} \| \boldsymbol{\sigma}(\mu) - P_{\mathcal{T}_k} [\boldsymbol{\sigma}(\mu)] \|_2$, where $k = \mathcal{I}(\mu)$
    Compute $\epsilon = \| \boldsymbol{\sigma}(\mu^\star) - P_{\mathcal{T}_{k^\star}} [\boldsymbol{\sigma}(\mu^\star)] \|_2$, where $k^\star = \mathcal{I}(\mu^\star)$
end while

4.3. Discussion

We have explained how local affine approximations in the form of Eq. (23) could be constructed following an automatic procedure, which is essentially the locally adaptive procedure described in [21]. At this point it is worth noticing that the proposed construction only requires the ability to evaluate $\mu \mapsto \boldsymbol{\sigma}(\mu)$, which, recalling Eq. (21), exclusively relies on the knowledge of the functions $g_\star^{\mathrm{ns}}$, $\star = 1, 2, 3$ and associated EIM interpolation matrices and interpolation points. Thus, the proposed construction is completely independent from the BEM discretization with $\mathcal{N}_h$ degrees of freedom.

The rationale behind the proposed construction is the following: at the end of Alg. 2, for any $\mu \in [\mu_{\min}, \mu_{\max}]$, the vector $\boldsymbol{\sigma}(\mu) \in \mathbb{C}^{Q^a}$ can be approximated

by its orthogonal projection $P_{\mathcal{T}_k}[\boldsymbol{\sigma}(\mu)]$ with $k = \mathcal{I}(\mu)$ with an error smaller than the prescribed tolerance *tol*. Replacing $\boldsymbol{\sigma}(\mu)$ by $P_{\mathcal{T}_k}[\boldsymbol{\sigma}(\mu)]$ in the affine approximation for the CFIE operator $\widetilde{\mathbf{A}}(\mu)$ given by Eq. (20) and recalling the expression Eq. (25) for $P_{\mathcal{T}_k}[\boldsymbol{\sigma}(\mu)]$, we get

$$\widetilde{\mathbf{A}}(\mu) = \sum_{q=1}^{Q^a} \sigma_q(\mu)\mathbf{A}_q \approx \sum_{q=1}^{Q^a}\sum_{p=1}^{Q} \theta_p^k(\mu)\sigma_q(\hat{\mu}_{k+p-1})\mathbf{A}_q. \tag{26}$$

Swapping the summations we obtain

$$\widetilde{\mathbf{A}}(\mu) \approx \sum_{p=1}^{Q} \theta_p^k(\mu) \sum_{q=1}^{Q^a} \sigma_q(\hat{\mu}_{k+p-1})\mathbf{A}_q = \sum_{p=1}^{Q} \theta_p^k(\mu)\widetilde{\mathbf{A}}(\hat{\mu}_{k+p-1}). \tag{27}$$

Omitting the tilde in the RHS of Eq. (27) yields the non-intrusive local approximation proposed in Eq. (23). The tilde can indeed be omitted, since $\widetilde{\mathbf{A}}$ is by design a good approximation for $\mathbf{A}$. Ultimately, we obtain that $\mathbf{A}(\mu) \approx \sum_{p=1}^{Q} \theta_p^k(\mu)\mathbf{A}(\hat{\mu}_{k+p-1})$ which corresponds to our initial non-intrusive local affine approximation statement Eq. (23).

5. Locally adaptive non-intrusive block RB method

We now revisit the classical RB method reviewed in Section 3, adapting it to the use of non-intrusive local affine approximations.

5.1. Block approach

Following the classical RB method presented in Section 3, the construction of a RB of size N requires N repeated solves of the BEM linear system (8) at parameter points $(\mu^{(1)}, \nu^{(1)}), \ldots, (\mu^{(N)}, \nu^{(N)})$. Here, we propose a different approach, the so-called *block* approach, that consists in choosing N_f wavenumber values denoted $\mu^{(1)}, \ldots, \mu^{(N_f)}$ and, for each wavenumber value $\mu^{(n)}$, an associated set of $I^{(n)} \geq 1$ incident angles $\{\nu_i^{(n)}\}_{1 \leq i \leq I^{(n)}}$. For convenience, let us assume that the N_f chosen wavenumbers are in increasing order, *i.e.*, $\mu^{(1)} \leq \cdots \leq \mu^{(N_f)}$ and for all $1 \leq n \leq N_f$ we introduce the notation

$$\mathbf{u}_i^{(n)} = \mathbf{u}(\mu^{(n)}, \nu_i^{(n)}), \quad 1 \leq i \leq I^{(n)}. \tag{28}$$

With the block approach, the overall number of computed BEM solutions is $N = I^{(1)} + \cdots + I^{(N_f)}$. In opposition to the classical approach, computing these N BEM solutions relies on only $N_f \leq N$ linear solves with multiple right-hand sides. Under the paradigm of direct solvers, recalling that the same LU-decomposition can be used to solve multiple right-hand sides, the block approach requires $\mathcal{O}(N_f\mathcal{N}_h^3 + N\mathcal{N}_h^2)$ operations to compute all the BEM solutions, which is more advantageous than the $\mathcal{O}(N\mathcal{N}_h^3 + N\mathcal{N}_h^2)$ operations required by the classical approach. Furthermore, we note that the block approach becomes increasingly advantageous as the number of incident angles per wavenumber

increases. This feature of the block approach can be maintained under the paradigm of iterative solvers provided that adequate block Krylov recycling strategies are used to efficiently solve multiple right-hand sides with the same matrix [27].

5.2. Locally adaptive approach

The classical RB approach consists in projecting the BEM system in a global RB of size N spanned by all N computed BEM solutions, as outlined in Section 3. In order to benefit from an efficient offline/online strategy and following standard practice [26], the matrix that is projected is not the BEM system matrix $\mathbf{A}(\mu)$ but the local affine approximation $\mathbf{A}_k(\mu)$ with $k = \mathcal{I}(\mu)$. Thus, the $N \times N$ matrix involved in the projected system matrix is $\mathbf{P}^* \mathbf{A}_k(\mu) \mathbf{P}$. Given the affine expression (23) of $\mathbf{A}_k(\mu)$, the online complexity for assembling the projected system matrix is $\mathcal{O}(QN^2)$. However, the success of the online phase requires that the J matrices $\mathbf{P}^* \mathbf{A}(\hat{\mu}_j) \mathbf{P}$, $\hat{\mu}_j \in \mathcal{C}_J$ (each matrix of size $N \times N$) have been pre-computed offline. These offline computations require performing N matrix-vector products with each BEM matrix $\mathbf{A}(\hat{\mu}_j)$, thus JN BEM matrix-vector products overall. In practice, these numerous BEM matrix-vector products can represent a significant computational effort, with the potential to compromise the overall efficiency of the RB method. In the following, we propose a locally adaptive approach that reduces the number of matrix-vector products required during the offline phase.

Following the locally adaptive approach, we let the user choose a number $N_f^{loc} \leq N_f$. Let us define $R = N_f - N_f^{loc} + 1$ and the discrete sets

$$\mathcal{T}^{(r)} = \{\mu^{(r)}, \ldots, \mu^{(r+N_f^{loc}-1)}\}, \quad r \in \{1, \ldots, R\}, \tag{29}$$

each of cardinality N_f^{loc}. We define the *indicator function* $\mathcal{J} : [\mu_{\min}, \mu_{\max}] \to \{1, \ldots, R\}$ mapping each wavenumber μ to the index r such that $\mathcal{T}^{(r)}$ is the set of the N_f^{loc} points among the $\mu^{(n)}$'s that are closest to μ.

This being set, we introduce for all $r \in \{1, \ldots, R\}$ the r^{th} local RB as $\mathbf{P}_r \in \mathbb{C}^{\mathcal{N}_h \times N_r}$ given by

$$\text{Colspan}(\mathbf{P}_r) = \text{Span}\{\mathbf{u}_i^{(\varrho)},\ 1 \leq i \leq I^{(\varrho)},\ r \leq \varrho \leq r + N_f^{loc} - 1\}, \tag{30}$$

such that $\mathbf{P}_r^* \mathbf{H} \mathbf{P}_r = \mathbf{I}$. In other words, the r^{th} local RB is the span of the BEM solutions at all the chosen incident angles associated to the wavenumbers ranging from $\mu^{(r)}$ to $\mu^{(r+N_f^{loc}-1)}$. The size N_r of the r^{th} local RB is given by

$$N_r = I^{(r)} + \cdots + I^{(r+N_f^{loc}-1)}. \tag{31}$$

As we shall see in the following, the r^{th} local RB will serve to build approximation of the BEM solution for values of μ in the sub-window $\mathcal{D}^{(r)} = \mathcal{J}^{-1}(r) = \{\mu \,|\, \mathcal{J}(\mu) = r\}$.

Indeed, for all wavenumber μ and all incident angle ν, we define the RB approximation $\mathbf{u}^{\text{rb}}(\mu, \nu) \in \mathbb{C}^{\mathcal{N}_h}$ as

$$\mathbf{u}^{\text{rb}}(\mu, \nu) = \mathbf{P}_r \mathbf{x}(\mu, \nu), \quad r = \mathcal{J}(\mu), \tag{32}$$

where $\mathbf{x}(\mu,\nu) \in \mathbb{C}^{N_r}$ is the solution to the following projected problem

$$\mathbf{P}_r^* \mathbf{A}_k(\mu) \mathbf{P}_r \mathbf{x}(\mu,\nu) = \mathbf{P}_r^* \widetilde{\mathbf{f}}(\mu,\nu), \quad k = \mathcal{I}(\mu),\ r = \mathcal{J}(\mu), \tag{33}$$

where $\mathbf{A}_k(\mu)$ is our non-intrusive local affine approximation given by Eq. (23) and $\widetilde{\mathbf{f}}(\mu,\nu)$ is an affine approximation of $\mathbf{f}(\mu,\nu)$ with Q^f terms, obtained by applying the EIM to the plane wave Eq. (1). The projected problem Eq. (33) is of size $N_r \times N_r$ and can be assembled in $\mathcal{O}(QN_r^2)$ operations following an adequate offline/online strategy. The next section provides a detailed description of the quantities which must be computed during the offline phase.

5.3. The offline phase in detail

In this section, we describe the offline phase in detail. We shall proceed in two steps: first, we detail the reduced matrices that need to be pre-computed and then we detail the necessary BEM matrix-vector operations that need to be performed when pre-computing these reduced matrices.

The reduced matrices that need to be pre-computed. Let us consider a fixed index $r \in \{1,\dots,R\}$. Let $k_{\min}^{(r)}$ (resp. $k_{\max}^{(r)}$) denote the smallest (resp. largest) index $k \in \{1,\dots,K\}$ such that

$$\mathcal{D}^{(r)} \cap \mathcal{D}_k \neq \emptyset \tag{34}$$

and set $\ell^{(r)} = k_{\max}^{(r)} - k_{\min}^{(r)}$. This being set, we find that it suffises to pre-compute the $Q + \ell^{(r)}$ reduced matrices $\mathbf{P}_r^* \mathbf{A}(\hat{\mu}_j) \mathbf{P}_r$ for $j \in \{k_{\min}^{(r)}, \dots, k_{\min}^{(r)} + Q + \ell^{(r)} - 1\}$ in order to be able to efficiently assemble for any $\mu \in \mathcal{D}^{(r)}$ the projected system matrix $\mathbf{P}_r^* \mathbf{A}_k(\mu) \mathbf{P}_r$ with $k = \mathcal{I}(\mu)$.

We conclude from these considerations that the overall number of reduced matrices that need to be precomputed offline is $RQ + \sum_{r=1}^{R} \ell^{(r)}$. An illustration is provided on figure 1.

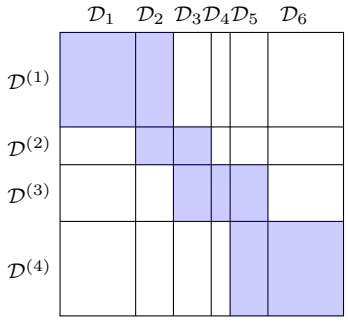

Figure 1: Schematic view of a case $K = 6$, $R = 4$.

For all $r = 1,\dots,R$ and all $k = 1,\dots,K$, non-empty sets $\mathcal{D}^{(r)} \cap \mathcal{D}_k$ are colored. In this situation : $k_{\min}^{(1)} = 1, k_{\min}^{(2)} = 2, k_{\min}^{(3)} = 3, k_{\min}^{(4)} = 5$ and $\ell^{(1)} = 1, \ell^{(2)} = 1$, $\ell^{(3)} = 2$, $\ell^{(4)} = 1$. Thus one needs to pre-compute:

- for $r = 1$, the $Q + 1$ reduced matrices $\mathbf{P}_1^* \mathbf{A}(\hat{\mu}_1)\mathbf{P}_1, \ldots, \mathbf{P}_1^* \mathbf{A}(\hat{\mu}_{Q+1})\mathbf{P}_1$,
- for $r = 2$, the $Q + 1$ reduced matrices $\mathbf{P}_2^* \mathbf{A}(\hat{\mu}_2)\mathbf{P}_2, \ldots, \mathbf{P}_2^* \mathbf{A}(\hat{\mu}_{Q+2})\mathbf{P}_2$,
- for $r = 3$ the $Q + 2$ reduced matrices $\mathbf{P}_3^* \mathbf{A}(\hat{\mu}_3)\mathbf{P}_3, \ldots, \mathbf{P}_3^* \mathbf{A}(\hat{\mu}_{Q+3})\mathbf{P}_3$
- and for $r = 4$ the $Q+1$ reduced matrices $\mathbf{P}_3^* \mathbf{A}(\hat{\mu}_5)\mathbf{P}_3, \ldots, \mathbf{P}_3^* \mathbf{A}(\hat{\mu}_{Q+5})\mathbf{P}_3$.

It is worth noting that $J = K + Q - 1 = Q + 5$.
In this situation : $r_1^{\min} = r_2^{\min} = \cdots = r_{Q+1}^{\min} = 1, r_{Q+2}^{\min} = 2, r_{Q+3}^{\min} = 3, r_{Q+4}^{\min} = 3, r_{Q+5}^{\min} = 4$.

The BEM matrix-vector products that need to be performed. We now determine the number of BEM matrix-vector products that need to be performed when pre-computing the reduced matrices. To this end, let us now consider a fixed index $j \in \{1, \ldots, J\}$. Clearly, $\hat{\mu}_j$ is a member of all the sets $\mathcal{T}_k$ for all $k \in \{\max(1, j - Q + 1), \ldots, \min(K, j)\}$. Let $r_j^{\min}$ (resp. $r_j^{\max}$) denote the smallest (resp. largest) index $r \in \{1, \ldots, R\}$ such that

$$\mathcal{D}^{(r)} \cap \left(\bigcup_{\max(1,j-Q+1) \leq k \leq \min(K,j)} \mathcal{D}_k \right) \neq \emptyset. \tag{35}$$

and set $s_j = r_j^{\max} - r_j^{\min}$. Thus, we find that the BEM matrix $\mathbf{A}(\hat{\mu}_j)$ is involved in the $s_j + 1$ reduced matrices $\mathbf{P}_r^* \mathbf{A}(\hat{\mu}_j)\mathbf{P}_r$ for all $r \in \{r_j^{\min}, \ldots, r_j^{\max}\}$. A logical way to proceed would be to successively compute the matrix-vector products $\mathbf{A}(\hat{\mu}_j)\mathbf{P}_r$ for $r \in \{r_j^{\min}, \ldots, r_j^{\max}\}$. Given that each local RB $\mathbf{P}_r$ has N_r columns, the number of matrix-vector products required with the BEM matrix $\mathbf{A}(\hat{\mu}_j)$ would be $\sum_{r=r_j^{\min}}^{r_j^{\max}} N_r$ under this approach.

We now propose an alternative way to proceed that reduces the number of matrix-vector products required with the BEM matrix $\mathbf{A}(\hat{\mu}_j)$ to $I^{(r_j^{\min})} + \cdots + I^{(r_j^{\max} + N_f^{loc} - 1)}$.

Let us introduce for $r^{\min} \in \{1, \ldots, R\}$ and $r^{\max} \geq r^{\min}$ the notation $\mathbf{U}_{r^{\min}:r^{\max}}$ for the matrix that concatenates the BEM solutions at all chosen incident angles associated to the wavenumbers ranging from $\mu^{(r^{\min})}$ to $\mu^{(r^{\max} + N_f^{loc} - 1)}$

$$\mathbf{U}_{r^{\min}:r^{\max}} = \left[\mathbf{u}_1^{(r^{\min})} | \cdots | \mathbf{u}_{I^{(r^{\min})}}^{(r^{\min})} | \cdots | \mathbf{u}_1^{(r^{\max} + N_f^{loc} - 1)} | \cdots | \mathbf{u}_{I^{(r^{\max} + N_f^{loc} - 1)}}^{(r^{\max} + N_f^{loc} - 1)} \right]. \tag{36}$$

Notice that $\mathbf{U}_{r^{\min}:r^{\max}}$ has $I^{(r^{\min})} + \cdots + I^{(r^{\max} + N_f^{loc} - 1)}$ columns. In particular, for $r^{\max} = r^{\min} = r$ we shall use the shortcut notation $\mathbf{U}_{r:r} = \mathbf{U}_r$. It is worth noticing that $\mathbf{U}_r$ has N_r columns, with N_r defined by (31). Furthermore, the columns of $\mathbf{U}_r$ and the columns of $\mathbf{P}_r$ span the same subspace (compare Eqs. (30) and (36)). Indeed, it is clear that

$$\mathbf{P}_r \mathbf{R}_r = \mathbf{U}_r, \tag{37}$$

with $\mathbf{R}_r = \mathbf{P}_r^* \mathbf{H} \mathbf{U}_r$ an invertible $N_r \times N_r$ matrix (this matrix is further upper triangular following the standard Gram-Schmidt procedure or the QR-decomposition). Rather than successively computing the matrix-vector products $\mathbf{A}(\hat{\mu}_j)\mathbf{P}_r$ for $r \in \{r_j^{\min}, \ldots, r_j^{\max}\}$, we can equivalently compute the matrix-vector products $\mathbf{A}(\hat{\mu}_j)\mathbf{U}_{r_j^{\min}:r_j^{\max}}$ and recover the quantities $\mathbf{A}(\hat{\mu}_j)\mathbf{P}_r$ for $r \in \{r_j^{\min}, \ldots, r_j^{\max}\}$ by the formula

$$\mathbf{A}(\hat{\mu}_j)\mathbf{P}_r = \mathbf{A}(\hat{\mu}_j)\mathbf{U}_r \mathbf{R}_r^{-1}. \tag{38}$$

Thus, the number of matrix-vector products required with the BEM matrix $\mathbf{A}(\hat{\mu}_j)$ is reduced to $I^{(r_j^{\min})} + \cdots + I^{(r_j^{\max}+N_f^{loc}-1)}$ as announced above. This number of matrix-vector products is $\leq N$ (with equality only in the case $r_j^{\min} = 1$, $r_j^{\max} = R$), therefore the use of local affine approximations combined with a locally adaptive RB strategy leads to a reduction in overall number of matrix-vector products compared to the classical RB approach.

6. Numerical illustration

6.1. The fighter aircraft test case

In the section, we present numerical results obtained on a target with the geometry of a simplified fighter aircraft. From nose to tail the length of the aircraft is about 15m while the wings span about 8m. This target is meshed with $\mathcal{N}_h = 145,000$ degrees of freedom. Following the well-known rule of letting 8 degrees of freedom per wavelength, this mesh is fit for frequencies up to 650MHz. Figure 2 shows the surface currents solved with the BEM on this mesh at the frequency 650MHz for two illuminating waves with the same direction but two different polarization. In this section, we propose to solve the scattering problem for frequencies ranging from 350MHz to 650MHz, a fixed polar angle $\theta = 90°$ (which corresponds to the plane of the wings), the azimuthal angle ϕ ranging from 160° to 200° and a fixed polarizations.

6.1.1. Preliminary step: the non-intrusive local affine approximations

The preliminary step for building an efficient ROM for this parametrized problem is constructing the non-intrusive local affine approximations. For this purpose, we first apply the EIM [18] to the three functions $g_1^{\text{ns}}, g_2^{\text{ns}}$ and g_3^{ns} defined in §3.2. Given the size of the target, we let $\xi_{\max} = 15\text{m}$. We draw attention to the fact that $\xi_{\max}$ is the only information that is needed on the target. In particular, the mesh of the target with $\mathcal{N}_h$ degrees of freedom is not used at this point. Prescribing a tolerance 10^{-6} on the EIM approximation error, we find $M_1 = 31, M_2 = 28$ and $M_3 = 30$. Thus, we have obtained an affine approximation for the CFIE operator in the form of Eq. (20) with $Q^a = 93$ terms. As anticipated in §3.2, Q^a is quite large.

Next, we apply the locally adaptive procedure explained in §4 to transform the *global and intrusive* affine approximation with Q^a terms (*global* in the sense that it is valid over 350MHz to 650MHz) into K *local and non-intrusive* affine

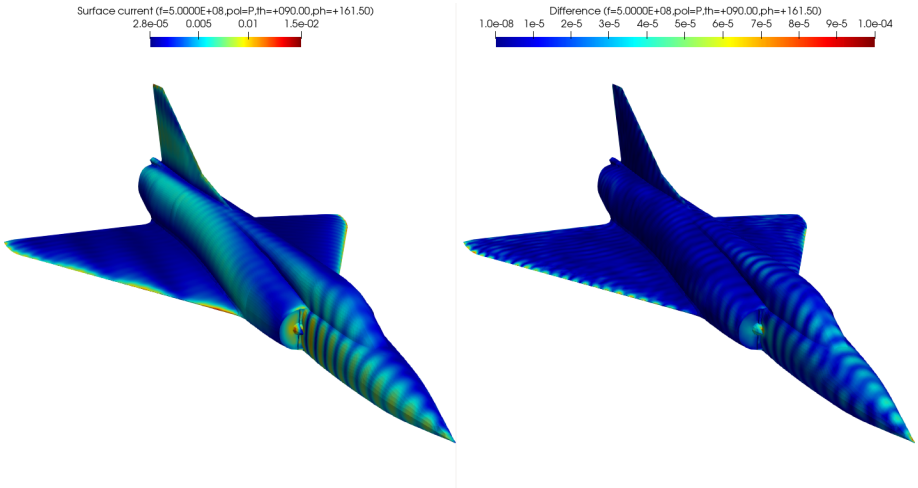

Figure 2: Surface currents on the target of the fighter aircraft test case for the illuminating wave with frequency $f = 650$MHz, polar angle $\theta = 90°$ (which corresponds to the xy plane of the wings) azimuthal angle $\phi = 160°$ for the two polarizations $\widehat{\boldsymbol{p}} = \widehat{\boldsymbol{e}}_\theta$ (left) and $\widehat{\boldsymbol{p}} = \widehat{\boldsymbol{e}}_\phi$ (right).

approximations. To this end, we prescribe a number of terms Q and a tolerance *tol* and let the algorithms 1 and 2 automatically build the K sub-windows and associated local non-intrusive affine approximations. For now, we choose $Q = 8$, $tol = 10^{-1}$, but we shall discuss the choice of Q later in §6.1.3. With these choices, the locally adaptive procedure terminates with $K = 24$, which means that a set $J = 31$ wavenumbers have been selected. Figure 3 shows the projection error with respect to the frequency at a few selected iterations of the greedy or locally adaptive procedures.

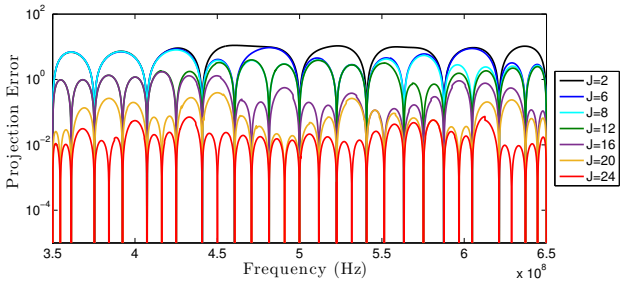

Figure 3: The projection error $\mu \mapsto \|\boldsymbol{\sigma}(\mu) - P_{\mathcal{C}_J}[\boldsymbol{\sigma}(\mu)]\|$ for selected values of J, plotted with respect to the frequency $f = \frac{c\mu}{2\pi}$ in Hz. The convergence to the prescribed tolerance $tol = 0.1$ occurs at $J = 24$.

6.1.2. Construction of the ROM

Following the block approach described in §5.1, we need to choose $N_f \geq 1$ wavenumbers $\mu^{(1)} \leq \cdots \leq \mu^{(N_f)}$ and, for each wavenumber $\mu^{(n)}$, we need to choose a set of $I^{(n)} \geq 1$ associated incident angles $\{\nu_i^{(n)}\}_{1 \leq i \leq I^{(n)}}$. In the present test case the polar angle is fixed to $\theta = 90°$, thus we only need to choose for each n the set of azimuthal angles $\{\phi_i^{(n)}\}_{1 \leq i \leq I^{(n)}}$. We propose to choose the number of frequencies N_f following

$$N_f = \left\lceil \frac{C_f(f_{\max} - f_{\min})}{\delta f} \right\rceil, \quad \delta f = \frac{c}{2L} \tag{39}$$

where $\lceil \cdot \rceil$ denotes the ceiling operation, C_f is a constant, $f_{\max} = 650$MHz is the maximal and $f_{\min} = 350$MHz is the minimum frequency, $c \approx 3 \times 10^8$m/s is the speed of light and $L = 18$m is the characteristic size of the target. While Shannon's sampling criterion [28] corresponds to the choice of constant $C_f = 1$. Based on our experience, we advocate using $C_f = \frac{2}{3}$.

Regarding the choice of incident angles, we explore two possible strategies:

- *The adaptive grid strategy.* This approach is motivated by the observations made in [9]. The number of associated incident angles increases with the frequency. In detail, $I^{(n)}$ is an increasing function of n. We propose to define number $I^{(n)}$ as

$$I^{(n)} = \left\lceil \frac{C_\phi(\phi_{\max} - \phi_{\min})}{\delta\phi^{(n)}} \right\rceil, \quad \delta\phi^{(n)} = \frac{c}{2Lf^{(n)}}, \tag{40}$$

  where C_ϕ is a constant, $f^{(n)} = \frac{\mu^{(n)}c}{2\pi}$ is the n^{th} frequency. The choice $C_\phi = 1$ corresponds to Shannon's sampling criterion at the frequency $f^{(n)}$, but based on some numerical tests, we advocate using $C_\phi = 3$.

- *The uniform grid strategy.* For simplicity, the set of associated incident angles is the same for all N_f wavenumbers, in other words $I^{(1)} = \cdots = I^{(N_f)} \equiv I$ and $\phi_i^{(1)} = \cdots = \phi_i^{(N_f)}$ for all $1 \leq i \leq I$. The number I is chosen using the formula (40), where $\delta\phi^{(n)}$ is replaced by the n-independent quantity $\delta\phi = \frac{c}{2Lf_{\max}}$.

For the present aircraft test, Eq. (40) gives $N_f = 24$. Figure 4 shows the adaptive and uniform grids.

The last ingredient for building the locally adaptive non-intrusive ROM is the choice of a local number of frequencies N_f^{loc}. For now, we set $N_f^{loc} = 10$. The choice of N_f^{loc} will be discussed later in §6.1.3. In this situation, the elapsed time for building the ROM is 2h38min with the adaptive grid and 3h04min with the uniform grid. These elapsed times take into account : (i) the preliminary step (three EIMs and the locally adaptive procedure), (ii) the N_f calls to the BEM linear solver for computing the BEM solutions following the block approach described in §5.1 and (iii) the computation of all the necessary reduced matrices with the strategy explained in §5.3. For solving the BEM,

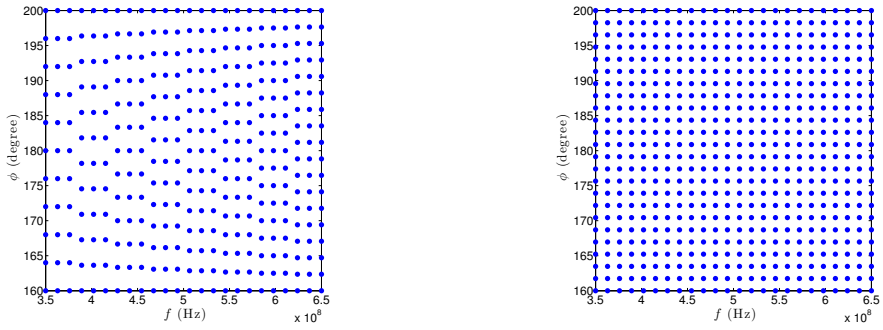

Figure 4: Adaptive grid (left) and uniform grid (right) for the fighter aircraft test case.

we use a multiple right-hand-side Generalized Conjugate Residual method [29]. We use the Multi-Level Fast-Multipole Method (MLFMM) for performing all matrix-vector operations [30]. Let us emphasize that the MLFMM is not only used within the iterative solver for performing matrix-vector operations at each iteration, it is also used for performing all matrix-vector operations listed in §5.3 which are required for the offline phase. All the computations run in parallel on 8 CPUs.

The ROM being built, it is now possible to compute very efficient approximation for the BEM solution for any frequency within $350 - 650$MHz and any azimuthal angle within $160 - 200°$. We compute the monostatic RCS on a cartesian grid with 121 frequencies and 161 azimuthal angles in 53min elapsed time with adaptive grid strategy and 58min elapsed time with the caratesian grid strategy. Notice that these online elapsed times include : (i) computing the reduced solutions, (ii) expressing the reduced solutions as $\mathcal{N}_h$-dimensional vectors to obtain the approximate surface currents and (iii) radiating the approximate surface currents to compute the RCS. Step (i) is fully $\mathcal{N}_h$-independent and thus very efficient, while steps (ii) and (iii) have a complexity dependent on $\mathcal{N}_h$. This dependency in $\mathcal{N}_h$ still allows online performance, because radiating the surface currents is a simple post-processing task which is not time-consuming compared to solving the BEM, for instance. Indeed, for comparison, it takes 7h13min to compute the RCS on the same 121×161 cartesian grid when resolving the surface currents by solving the BEM rather than the ROM.

For assessing the accuracy of the ROM, we compare the RCS obtained with the ROM and the RCS obtained with the BEM on grid 51×81 cartesian grid, see figure 5. Given a frequency f and an azimuthal angle ϕ, let $RCS(f, \phi)$ be the reference RCS obtained from the surface currents resolved with the BEM and let $\widetilde{RCS}(f, \phi)$ denote the approximate RCS obtained from the surface currents resolved with the ROM. On figure 5 (right), we have plotted the quantity $(f, \phi) \mapsto 20 \log_{10} \Delta(f, \phi)$, where $\Delta(f, \phi)$ is the relative error given by

$$\Delta(f, \phi) = \frac{|RCS(f, \phi) - \widetilde{RCS}(f, \phi)|}{|RCS(f, \phi)|}. \tag{41}$$

The logarithm provides a measure of the error of the ROM in dB, which is the usual unit for radarists. We find that this error never exceeds -20dB, which confirms the excellent accuracy of the ROM.

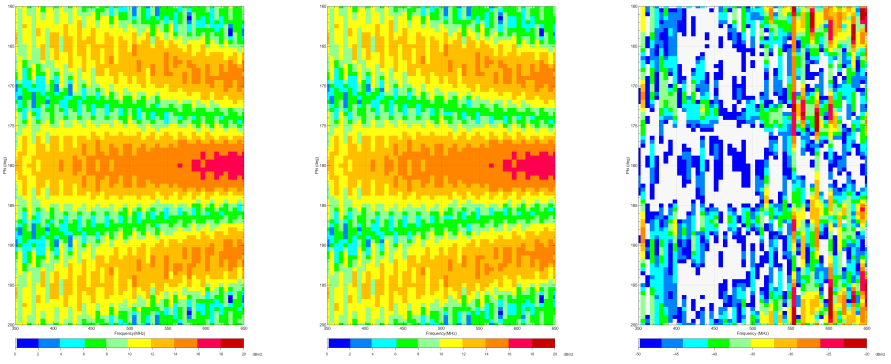

Figure 5: Monostatic RCS in the (f, ϕ) plane for the aircraft test case on a cartesian grid with 51 frequencies and 81 azimuthal angles. From left to right: reference, approximation (with $Q = 8$, $N_f^{loc} = 10$, adaptive grid strategy) and the error in dB.

6.1.3. Different choices for Q and N_f^{loc}

So far, we have successfully built two ROMs, with the adaptive and cartesian grid strategies using the fixed parameters $Q = 8$, $N_f^{loc} = 10$. In this section, we now explore the choice of the two parameters Q and N_f^{loc} in more detail.

We start by applying the locally adaptive procedure prescribing different number of terms $Q = 6, 8, 10, 12$. The procedure terminates with $K = 34, 24, 19, 15$ sub-windows respectively, which means that Algs. 1 and 2 have selected $J = 39, 31, 28, 26$ frequencies respectively. In the present case, it is to be observed that dividing the prescribed number of terms by two (from 12 to 6) does roughly multiply the number of sub-windows by a factor 2, while the number of selected frequencies is multiplied by a factor $\frac{3}{2}$. This shows that even though the locally adaptive procedure builds many sub-windows, yet the number of selected frequencies does not grow too fast.

We now explore different choices for N_f^{loc}, namely $N_f^{loc} = 6, 8, 10, 12$. For each combination of Q and N_f^{loc}, two domain-decompositions of the frequency window co-exist: the domain-decomposition in K sub-windows $\mathcal{D}_k$, $1 \leq k \leq K$, which is exclusively due to the choice of Q and the domain-decomposition in R sub-windows $\mathcal{D}^{(r)}$, $1 \leq r \leq R$, which is uniquely determined by the choice of N_f^{loc}. As explained in §5.3, the complexity of the offline phase essentially involves the (k, r)-couples for which the intersection $\mathcal{D}_k \cap \mathcal{D}^{(r)}$ is non-empty. In order to illustrate this, we have plotted on figure 6 what we shall call the *domain-decomposition skeleton*. This consists in the $K \times R$ grid where the (k, r)-couples for which the set $\mathcal{D}_k \cap \mathcal{D}^{(r)}$ is non-empty are marked by a black tile.

As can be seen on figure 6, the domain-decomposition skeleton always takes the form of an approximate diagonal.

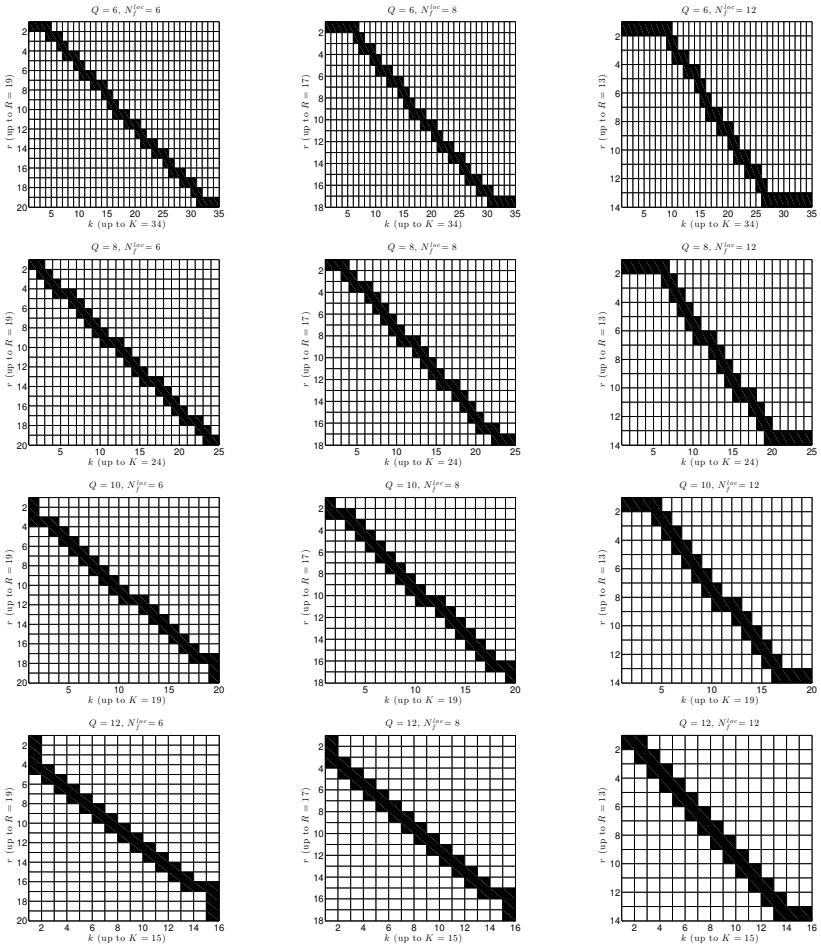

Figure 6: Domain-decomposition skeletons for the adaptive grid, $Q = 6, 8, 10, 12$ and $N_f^{loc} = 6, 8, 12$. Each row corresponds to a constant Q and each column to a constant N_f^{loc}. For each domain decomposition skeleton, the tile (k, r) is black when the sets $\mathcal{D}_k$ and $\mathcal{D}^{(r)}$ have a non empty intersection (see text).

We draw attention to two typical situations. The first corresponds to a situation where the domain-decomposition in K sub-windows is relatively more refined than the domain-decomposition in R sub-windows. Thus, for a fixed r, there may be plenty of different values of k such that $\mathcal{D}_k \cap \mathcal{D}^{(r)}$ is non-empty, resulting in a horizontal line in the domain-decomposition skeleton as can be observed for instance for $Q = 6$, $N_f^{loc} = 12$ (top right of figure 6). The second typical situation corresponds to a domain-decomposition in R sub-windows relatively more refined than the domain-decomposition in K sub-windows. In

this case, for a fixed k, there may be plenty of different values of r such that $\mathcal{D}_k \cap \mathcal{D}^{(r)}$ is non-empty, resulting in a vertical line in the domain-decomposition skeleton. The case $Q = 12$, $N_f^{loc} = 6$ (bottom left of figure 6) nicely illustrates this situation. We conclude that exploring the different choices $Q = 6, 8, 10, 12$ and $N_f^{loc} = 6, 8, 10, 12$ is a good sampling of the various possible situations.

For each combination of Q and N_f^{loc}, a different ROM is built. Table 1 consigns the number of matrix-vector products performed during each of the different offline phases. It is to be observed that for a fixed Q the number of matrix-vector products increases with N_f^{loc}. This means that the number of matrix-vector products performed with each matrix $\mathbf{A}(\hat{\mu}_j)$, $1 \leq j \leq J$ increases with N_f^{loc} (the number J of matrices is constant for a fixed Q). Table 1 also shows that for a fixed N_f^{loc} the overall number of matrix-vector products first decreases with Q until some threshold value of Q after which the number of matrix-vector products increases again. The existence of this threshold value of Q is related to the fact that, under a decreasing Q, the number J of different matrices increases while the number of matrix-vector products required with each matrix $\mathbf{A}(\hat{\mu}_j)$, $1 \leq j \leq J$ decreases (because each $s_j = r_j^{\max} - r_j^{\min}$ decreases, see §5.3). In the present test case, the threshold value of Q is found to be around 6 or 8 depending on N_f^{loc}.

	Adaptive grid				Uniform grid			
$Q \backslash N_f^{loc}$	6	8	10	12	6	8	10	12
6	4959	5917	6876	7834	6156	7344	8532	9720
8	4789	5457	6126	6794	5940	6768	7596	8424
10	5048	5571	6093	6616	6264	6912	7560	8208
12	5406	5816	6226	6636	6696	7200	7704	8208

Table 1: Number of matrix-vector products performed during the offline phase for different values of Q and N_f^{loc} using the adaptive and uniform grid strategies.

	Adaptive grid				Uniform grid			
$Q \backslash N_f^{loc}$	6	8	10	12	6	8	10	12
6	2:08	2:19	2:33	2:44	2:20	2:39	2:56	3:15
8	2:03	2:16	2:25	2:37	2:19	2:34	2:53	3:07
10	2:03	2:17	2:27	2:37	2:27	2:38	2:52	3:09
12	2:07	2:19	2:30	2:40	2:23	2:39	2:56	3:22

Table 2: Elapsed times ($h\!:\!mm$, with h the number of hours and mm the number of minutes) for the offline phase with different values of Q and N_f^{loc} using the adaptive and uniform grid strategies.

Table 2 shows the elapsed times for the different offline phases. A quick comparison with Table 1 confirms that the offline phase is more time-consuming if there are more matrix-vector products to perform, which is the expected behavior. However, a close comparison reveals that the elapsed times are not always

exactly correlated to the overall number of matrix-vector products presented in Table 1. For instance, the situation with the uniform grid $Q = 10$, $N_f^{loc} = 6$ which requires 6264 matrix-vector products is surprisingly more time-consuming than the situation with $Q = 12$, $N_f^{loc} = 6$ which yet requires more (6696) matrix-vector products. We believe that this can be explained by our computational strategy for performing the matrix-vector operations in practice. Rather than computing the n matrix-vector products $\mathbf{A}\mathbf{v}_1,\ldots,\mathbf{A}\mathbf{v}_n$ successively with the same matrix $\mathbf{A}$, we prefer computing a single matrix-matrix product $\mathbf{AV}$, where $\mathbf{V}$ is the matrix whose n columns correspond to the n vectors $\mathbf{v}_1,\ldots,\mathbf{v}_n$. This strategy favors the cases where the number of matrix-vector products required with each matrix $\mathbf{A}(\hat{\mu}_j)$, $1 \leq j \leq J$ is large, i.e., these are the cases that exhibit long horizontal lines in their domain decomposition skeleton.

Table 3 shows the maximum relative error for the each of the constructed ROMs. These results show that the error decreases with increasing N_f^{loc}. This is related to the fact that larger values of N_f^{loc} are associated with richer approximation spaces with better approximation properties [21]. Moreover, we find that the error is only slightly smaller with the uniform than with the adaptive grid strategy. Indeed, it is to be observed that the two strategies provide very similar results in terms of errors. Given that the adaptive grid strategy is able to recover these results much faster (see Table 2), this strategy is more competitive than the uniform grid strategy in practice. Finally, Table 3 confirms that the ROM approximation error does not depend on Q as announced.

	Adaptive grid				Uniform grid			
$Q \backslash N_f^{loc}$	6	8	10	12	6	8	10	12
6	0.78	0.68	0.66	0.61	0.72	0.63	0.62	0.57
8	0.78	0.68	0.66	0.61	0.71	0.62	0.61	0.56
10	0.79	0.68	0.66	0.61	0.72	0.63	0.62	0.57
12	0.79	0.69	0.67	0.62	0.72	0.64	0.63	0.59

Table 3: Maximum relative error $\max \Delta(f, \phi)$ in % over a 121×161 grid for different values of Q and N_f^{loc} using the adaptive and uniform grid strategies.

6.2. The UAV test case

We present numerical results on a target with the geometry of a UAV. Compared to the fighter aircraft test case, this target is much smaller (about 3.2m long and 2.9m wingspan). This target is meshed with about $\mathcal{N}_h = 290,000$ degrees of freedom following the rule of letting 8 degrees of freedom per wavelength. The frequency interval of interest for this target is $2.5 - 4$GHz and the range of interest for the azimuthal angle ϕ is -15 to $+15°$. The surface currents on the UAV target for an arbitrary point (f, ϕ) within these bounds is shown on figure 8. Notice the presence of feature on the right wing that is absent on the left wing that breaks the symmetry of the target.

We build the local non-intrusive affine approximations of the CFIE operator using $Q = 8$ and prescribing the tolerance $tol = 10^{-1}$ in the locally adaptive

procedure. The procedure stops with a set of $J = 57$ selected wavenumbers (i.e., $K = 50$ sub-windows). For constructing the ROM, we choose $N_f = 30$, which is the number of frequencies stemming from the Shannon criterion Eq. (39) with $f_{\min} = 2.5$GHz, $f_{\max} = 4$GHz and the advocated constant $C_f = \frac{3}{2}$. The number of incident angles to be considered for each frequency value is chosen following the adaptive grid strategy Eq. (40) with the advocated constant $C_\phi = 3$. The grid that is used for constructing the ROM is shown on figure 7.

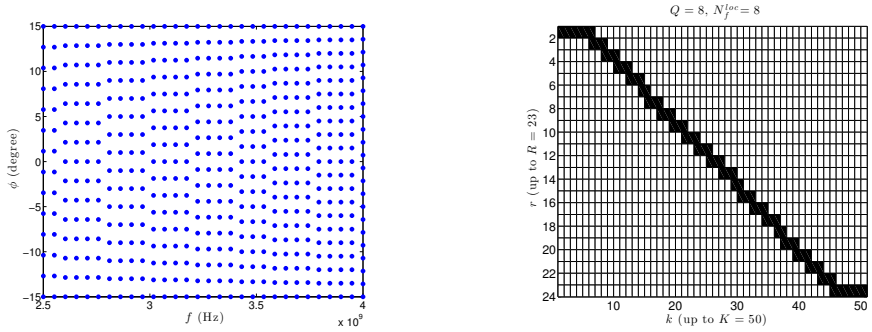

Figure 7: For the UAV test case : adaptive grid (left) and skeleton obtained with the choice $Q = 8$, $N_f^{loc} = 8$.

Finally, we choose the local number of frequencies $N_f^{loc} = 8$. Based on our previous study, this is a rather low value for N_f^{loc}, that is expected to favor performance over quality of approximation. This choices yields a domain-decomposition into $R = 30$ sub-windows. The domain-decomposition skeleton is shown on figure 7. It is to be observed that the domain-decomposition skeleton is slightly more refined horizontally than vertically which, as already mentioned, is an advantage when matrix-vector operations with the same matrix are performed using blocks of vectors.

Table 4 consigns the elapsed times for two different polarizations of the incident plane wave (T-T and P-P). The offline elapsed time is the elapsed time for constructing the ROM, while the online elapsed time is the elapsed time for computing the reduced solutions on a cartesian grid with 101 frequencies and 121 incident angles and post-processing these reduced solutions for obtaining the RCS. First, let us notice that the BEM takes slightly longer to solve for the T-T than for the P-P polarization. This is reflected in both the offline elapsed times and the BEM RCS recovery times consigned in Table 4. However, it is to be observed that once the ROMs are built, the online RCS recovery times are the same for both the T-T and P-P polarization. This illustrates the fact that the online complexities are indeed strictly identical for both polarizations.

Let us now consider the *instantaneous speed-up*, defined as the ratio between the elapsed time for recovering the RCS without the ROM (i.e., with the BEM) and with the ROM. The instantaneous speed-up is the usual measure of speed-up [9, 19, 31], but let us emphasize that the instantaneous speed-up does not take the offline elapsed time into account. In fact, the instantaneous speed-up

Polarization	Offline	Online (101×121)	BEM (101×121)
T-T	11h08	1h07	26h43
P-P	10h09	1h08	23h04

Table 4: Elapsed times for the UAV test case : $Q = 8$, $N_f^{loc} = 8$.

must be interpreted as a measure of the speed-up once the ROM is constructed and readily available. In the present case, we find an instantaneous speed-up of about 24 for the T-T polarization and about 20 for the P-P polarization. In the present situation where both polarization have identical online costs, the larger instantaneous speed-up for the T-T polarization comes as a consequence of the T-T polarization being more costly to solve with BEM.

In terms of quality of approximation, we find the maximum of the relative error $\Delta(f, \phi)$ (see by eq. (41)) over the 101×121 (f, ϕ) cartesian grid to be 2.95% for the T-T polarization (3.95% for the P-P polarization). Figure 8 shows the difference between the BEM surface currents and the ROM surface currents on the target at an arbitrary point (f, ϕ). In this example, the error in the surface currents is relatively well distributed on the target with no particular hot spots and the magnitude of the error is observed to be below 1%.

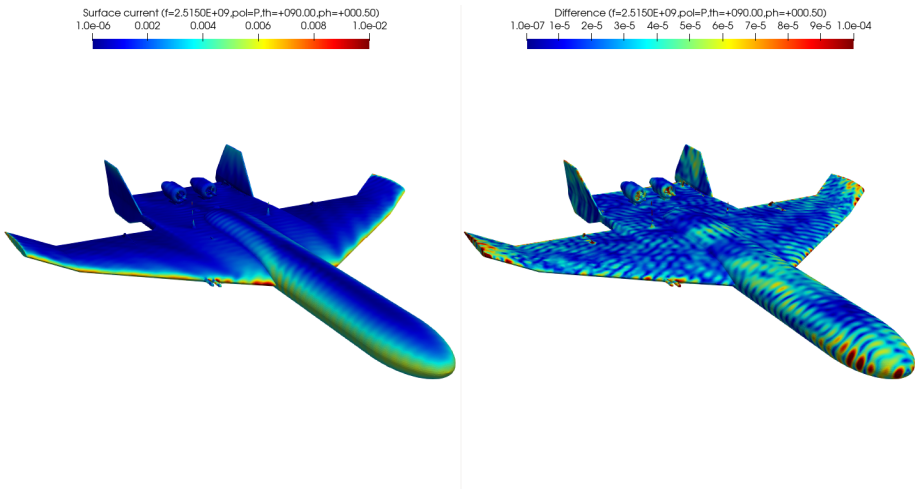

Figure 8: Surface currents computed with the BEM on the UAV target (left) and difference between BEM and ROM surface currents (right). The frequency is $f = 2.515$GHz and the azimuthal angle $\phi = +0.5°$ (this (f, ϕ) couple does not coincide with any point in the grid shown on figure 7).

Such levels of error may seem high (they are compared to our previous test case), but they are more than acceptable, for example in the context of holographic imaging. An holographic image is obtained by post-processing the RCS. Figure 9 shows two holographic images, the one on the left obtained using the BEM only and the one on the right using the ROM. Clearly, these two images

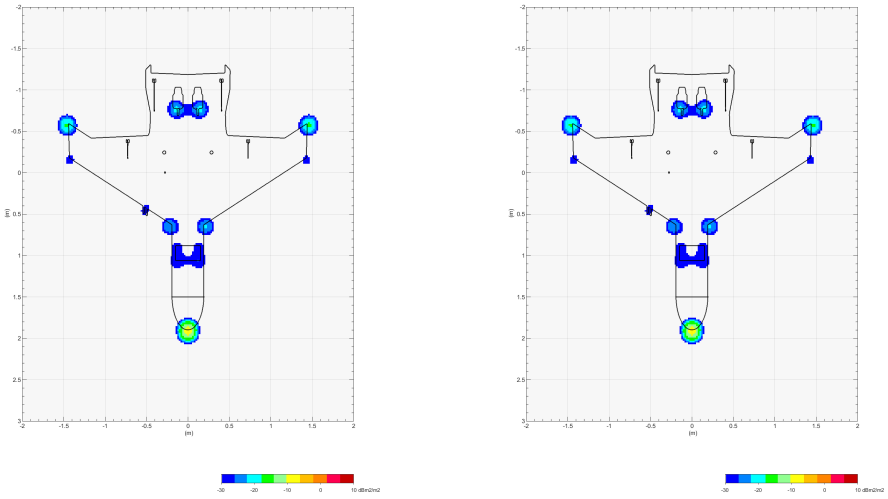

Figure 9: Holographic images in the UAV test case : computed using the BEM (left) and using the ROM (right).

cannot be distinguished, yet the ROM allows a faster recovery by a factor more than two overall (taking both offline and online phases into account).

7. Conclusions and perspectives

In this paper, we have presented a non-intrusive method for constructing efficient reduced order models for parametrized electromagnetic scattering problems solved using surface integral equations.

Our main contribution is to propose *non-intrusive local affine approximations* for approximating the parametrized integral operator. We have chosen to present our approach on the CFIE integral operator in electromagnetic scattering, but of course our approach is general and could be applied on any other integral operator arising in electromagnetism or acoustics. The original method proposed in this paper is especially thought for, but not limited to, combined formulations, for which multiple kernels need to be dealt with. In this context, the standard reduced basis approach reviewed in section 3, fails to provide an efficient affine approximation, because the number of affine terms is usually very large [9]. Combining the ideas found in Ref. [20] to achieve non-intrusivity and the locally adaptive concept found in Ref. [21] naturally yields the non-intrusive local affine approximations. An efficient construction procedure has been proposed in section 4. In section 5, the classical reduced basis method has been adapted to use the non-intrusive local affine approximations. Our proposition was to use local reduced basis approximation spaces, rather than a

single reduced basis approximation space following [21]. We have shown that the combining the non-intrusive local affine approximations with local reduced basis approximation spaces enabled a potential reduction in the number of matrix-vector products that need to be performed during the offline phase, as well as a reduction in the memory requirements, with potentially less reduced matrices making up the reduced order model.

We have illustrated our approach in section 6, on two industrially relevant test cases. All the BEM solutions required to construct the different ROMs have been computed using a FMM solver and all the required BEM matrix-vector products have been performed with the FMM. Hence, the fully-populated BEM matrix is never assembled nor stored, thus enabling efficiency on large-scale problems. On the fighter aircraft test case, we have explored the influence of the two user-defined parameters Q (the number of affine terms in each non-intrusive local affine approximation) and N_f^{loc} (the number of different frequencies to be considered in each local reduced basis approximation space). We have shown the quality of the ROM strongly depended on N_f^{loc} but not at all on Q. Our numerical experiment illustrate that a good choice of Q can indeed minimize the number of matrix-vector products that need to be performed offline and help reduce the computational costs. For constructing the reduced basis approximation spaces, we have proposed and compared two different strategies : the uniform grid strategy (constant number of incident angles per frequency) and the adaptive grid strategy (number of incident angles increasing with the frequency). We have further proposed to rely on Shannon's criterion to determine the number of frequencies and incident angles to be used. Numerical results have confirmed the soundness of this choice and the relative superiority of the adaptive versus the uniform grid strategy (already observed in [9]) which shows a more interesting balance between performance and the quality of approximation. Finally, we have shown our method in action on the UAV test case, which is a test case of industrial size and interest. We have shown significant speed-ups with our method and have confirmed that the approximation errors were sufficiently low for the reconstructed holographic image to be indistinguishable from the reference holographic image.

The method that we have described is non-intrusive in the sense that the BEM solver used to compute the BEM solutions and perform the BEM matrix-vector products can be *any* BEM solver. Our future work will include analyzing the performance of the locally adaptive non-intrusive block reduced basis method for different BEM solvers, such as ACA or H-matrix solvers.

References

References

[1] D. Colton, R. Kress, Integral equations in scattering theory, Pure and Applied Mathematics, John Wiley and Sons, New York 228.

[2] J.-C. Nédélec, Acoustic and electromagnetic equations: integral representations for harmonic problems, Vol. 144, Springer, 2001.

[3] V. Rokhlin, Rapid solution of integral equations of scattering theory in two dimensions, Journal of Computational Physics 86 (2) (1990) 414–439.

[4] S. Kurz, O. Rain, S. Rjasanow, The adaptive cross-approximation technique for the 3d boundary-element method, IEEE transactions on Magnetics 38 (2) (2002) 421–424.

[5] W. Hackbusch, A sparse matrix arithmetic based on h-matrices. part i: Introduction to h-matrices, Computing 62 (2) (1999) 89–108.

[6] I. Oseledets, E. Tyrtyshnikov, Tt-cross approximation for multidimensional arrays, Linear Algebra and its Applications 432 (1) (2010) 70–88.

[7] A. Quarteroni, A. Manzoni, F. Negri, Reduced basis methods for partial differential equations: an introduction, Vol. 92, Springer, 2015.

[8] J. S. Hesthaven, G. Rozza, B. Stamm, et al., Certified reduced basis methods for parametrized partial differential equations, Vol. 590, Springer, 2016.

[9] M. Fares, J. S. Hesthaven, Y. Maday, B. Stamm, The reduced basis method for the electric field integral equation, Journal of Computational Physics 230 (14) (2011) 5532–5555.

[10] Z. Bai, Krylov subspace techniques for reduced-order modeling of large-scale dynamical systems, Applied numerical mathematics 43 (1-2) (2002) 9–44.

[11] D. Panagiotopoulos, E. Deckers, W. Desmet, Krylov subspaces recycling based model order reduction for acoustic bem systems and an error estimator, Computer Methods in Applied Mechanics and Engineering 359 (2020) 112755.

[12] X. Xie, Y. Liu, An adaptive model order reduction method for boundary element-based multi-frequency acoustic wave problems, Computer Methods in Applied Mechanics and Engineering 373 (2021) 113532.

[13] H. Jiang, X. Zhang, X. Huang, Reduced-basis boundary element method for efficient broadband acoustic simulation, Journal of Sound and Vibration 456 (2019) 374–385.

[14] A. Buffa, Y. Maday, A. T. Patera, C. Prud'homme, G. Turinici, A priori convergence of the greedy algorithm for the parametrized reduced basis method, ESAIM: Mathematical modelling and numerical analysis 46 (3) (2012) 595–603.

[15] J. S. Hesthaven, B. Stamm, S. Zhang, Certified reduced basis method for the electric field integral equation, SIAM Journal on Scientific Computing 34 (3) (2012) A1777–A1799.

[16] S. Li, An efficient technique for multi-frequency acoustic analysis by boundary element method, Journal of sound and vibration 283 (3-5) (2005) 971–980.

[17] Z. Wang, Z. Zhao, Z. Liu, Q. Huang, A method for multi-frequency calculation of boundary integral equation in acoustics based on series expansion, Applied acoustics 70 (3) (2009) 459–468.

[18] M. Barrault, Y. Maday, N. C. Nguyen, A. T. Patera, An 'empirical interpolation' method: application to efficient reduced-basis discretization of partial differential equations, Comptes Rendus Mathematique 339 (9) (2004) 667–672.

[19] Y. Shi, X. Chen, Y. Tan, H. Jiang, S. Liu, Reduced-basis boundary element method for fast electromagnetic field computation, JOSA A 34 (12) (2017) 2231–2242.

[20] F. Casenave, A. Ern, T. Lelièvre, A nonintrusive reduced basis method applied to aeroacoustic simulations, Advances in Computational Mathematics 41 (5) (2015) 961–986.

[21] Y. Maday, B. Stamm, Locally adaptive greedy approximations for anisotropic parameter reduced basis spaces, SIAM Journal on Scientific Computing 35 (6) (2013) A2417–A2441.

[22] F. Chinesta, A. Huerta, G. Rozza, K. Willcox, Model order reduction, Encyclopedia of computational mechanics.

[23] Y. Maday, A. T. Patera, D. V. Rovas, A blackbox reduced-basis output bound method for noncoercive linear problems.

[24] G. Rozza, D. B. P. Huynh, A. T. Patera, Reduced basis approximation and a posteriori error estimation for affinely parametrized elliptic coercive partial differential equations, Archives of Computational Methods in Engineering 15 (3) (2007) 1.

[25] K. Veroy, C. Prud'Homme, D. Rovas, A. Patera, A posteriori error bounds for reduced-basis approximation of parametrized noncoercive and nonlinear elliptic partial differential equations, in: 16th AIAA Computational Fluid Dynamics Conference, 2003, p. 3847.

[26] M. A. Grepl, Y. Maday, N. C. Nguyen, A. T. Patera, Efficient reduced-basis treatment of nonaffine and nonlinear partial differential equations, ESAIM: Mathematical Modelling and Numerical Analysis 41 (3) (2007) 575–605.

[27] K. M. Soodhalter, E. de Sturler, M. E. Kilmer, A survey of subspace recycling iterative methods, GAMM-Mitteilungen 43 (4) (2020) e202000016.

[28] C. E. Shannon, Communication in the presence of noise, Proceedings of the IRE 37 (1) (1949) 10–21.

[29] P. Soudais, Iterative solution of a 3-d scattering problem from arbitrary shaped multidielectric and multiconducting bodies, IEEE transactions on antennas and propagation 42 (7) (1994) 954–959.

[30] J. Song, C.-C. Lu, W. C. Chew, Multilevel fast multipole algorithm for electromagnetic scattering by large complex objects, IEEE transactions on antennas and propagation 45 (10) (1997) 1488–1493.

[31] X. Xie, W. Wang, K. He, G. Li, Fast model order reduction boundary element method for large-scale acoustic systems involving surface impedance, Computer Methods in Applied Mechanics and Engineering 400 (2022) 115618.