



Gradient-based constrained optimization using a database of linear reduced-order models

Youngsoo Choi^a, Gabriele Boncoraglio^b, Spenser Anderson^b, David Amsallem^c, Charbel Farhat^{b,c,d,*}

^a Computational Engineering Division, Lawrence Livermore National Laboratory, Livermore, CA 94550, United States of America¹

^b Department of Aeronautics and Astronautics, Stanford University, Stanford, CA 94305, United States of America

^c Department of Mechanical Engineering, Stanford University, Stanford, CA 94305, United States of America

^d Institute for Computational and Mathematical Engineering, Stanford University, Stanford, CA 94305, United States of America

^e Facebook, Menlo Park, CA 94025, United States of America

ARTICLE INFO

Article history:

Received 16 November 2019

Received in revised form 12 August 2020

Accepted 13 August 2020

Available online 20 August 2020

Keywords:

Constrained optimization

Flutter

Gradient-based optimization

Interpolation on a matrix manifold

Model reduction

Parameter sampling

ABSTRACT

A methodology grounded in model reduction is presented for accelerating the gradient-based solution of a family of linear or nonlinear constrained optimization problems where the constraints include at least one linear Partial Differential Equation (PDE). A key component of this methodology is the construction, during an offline phase, of a database of pointwise, linear, Projection-based Reduced-Order Models (PROM)s associated with a design parameter space and the linear PDE(s). A parameter sampling procedure based on an appropriate saturation assumption is proposed to maximize the efficiency of such a database of PROMs. A real-time method is also presented for interpolating at any queried but unsampled parameter vector in the design parameter space the relevant sensitivities of a PROM. The practical feasibility, computational advantages, and performance of the proposed methodology are demonstrated for several realistic, nonlinear, aerodynamic shape optimization problems governed by linear aeroelastic constraints.

© 2020 Elsevier Inc. All rights reserved.

1. Introduction

Partial Differential Equation (PDE)-based constrained optimization problems arise in numerous engineering applications including automatic design, optimal control, and inverse identification. In practice, such problems are typically solved using a Nested ANalysis and Design (NAND) approach, which incurs the solution of many instances of the constraining parametric PDE, for different sampled parameter vectors in the design parameter space denoted here by \mathcal{D} . In many industrial settings, such a PDE may model, for example, a Computational Fluid Dynamics (CFD) or Computational Structural Dynamics (CSD) problem, its discretization is high-dimensional and therefore computationally intensive, and hence its enforcement can represent a significant fraction of the total cost of the optimization problem. Of course, this situation is exacerbated when the optimization problem of interest is constrained by multiple such PDEs. For this reason, PDE-constrained optimization problems are often solved using a less computationally intensive surrogate model such as, for example, a set of response surfaces for some Quantities of Interest (QoIs), or a Projection-based Reduced-Order Model (PROM).

* Corresponding author.

E-mail address: cfarhat@stanford.edu (C. Farhat).

¹ Lawrence Livermore National Laboratory is operated by Lawrence Livermore National Security, LLC, for the U.S. Department of Energy, National Nuclear Security Administration under Contract DE-AC52-07NA27344 and LLNL-JRNL-796678.

In general, a PROM seeks an approximate solution of the underlying High-Dimensional Model (HDM) within a subspace of the high-dimensional state space that is spanned by a Reduced-Order Basis (ROB) [1,2]. In the context of an optimization problem, an HDM is typically a μ -parametric HDM, where $\mu \in \mathcal{D} \subset \mathbb{R}^{N_{\mathcal{D}}}$ is a design parameter vector, and $N_{\mathcal{D}}$ denotes the dimension of the design parameter space \mathcal{D} . It is well-known that a PROM constructed for a specific design parameter vector $\mu \in \mathcal{D}$, which is referred to in this paper as a *pointwise* PROM, does not necessarily perform well at another design parameter vector μ' in this space (for example, see [3,4]). To address this issue, at least three approaches have been proposed so far.

In the first approach, a global ROB is constructed offline such that its associated PROM is accurate in the entire design parameter space, or at least a large region of this space [5–7]. A drawback of this approach, particularly for high-dimensional design parameter spaces, is that it may lead to a ROB that is prohibitively large, or sufficiently large to offset the computational advantages of the associated PROM. For nonlinear problems – in this context, the discretized nonlinear PDE constraints – the concept of a cluster of local, low-dimensional ROB's proposed in [8] addresses this issue. In this concept, locality refers to a region of the solution manifold, and the most appropriate local ROB is selected to approximate the solution when it evolves on its manifold.

Unlike the first approach which is comprehensive, the second approach was motivated primarily by optimization problems. In this approach, an initial ROB is *progressively* adapted so that it remains accurate not necessarily in the entire design parameter space, but in those regions of this space that are visited by the optimization algorithm [9–12]. Specifically, as new parameter vectors are queried by the optimizer in the design parameter space, additional solution snapshots are computed only when needed to maintain accuracy, and the dimension of the PROM is adjusted to incorporate the information gathered from these snapshots. This approach is suitable when the traditional offline-online computational framework is not particularly advantageous, and the minimization of the overall computational cost is instead preferred. On the other hand, it is not necessarily advantageous when the offline-online approach is desired, and the objective is to accelerate the execution of the online phase.

For linear problems – and in the context of this paper, specifically for the discretized linear PDE constraints – the parameter dependence of a PROM associated with a μ -parametric HDM can be addressed using the concept of a database of PROMs introduced in [13,14]. This approach can be summarized as follows. First, the design parameter space of interest is sampled offline using an appropriate greedy procedure, and a linear, pointwise PROM is constructed at each sampled design parameter vector using any preferred model reduction technique. Then, at each queried but unsampled design parameter vector, a PROM is constructed by interpolating the pre-computed PROMs on one or multiple matrix manifolds that characterize them. For linear problems, this approach is attractive for two reasons: the dimension of each pointwise PROM can be kept small while maintaining local accuracy, due to the restriction of the scope of such a PROM to a single parameter vector in the design parameter space (and perhaps a relatively small neighborhood); and the aforementioned interpolation on matrix manifolds can be performed in real time.

This paper focuses on the family of constrained optimization problems where the objective function may be linear or nonlinear but is time-independent, and the PDE constraint – and in the case of multiple PDE constraints, at least one of them – is linear and its discretization is high-dimensional. For such problems, which abound in engineering applications, the concept of a database of pointwise, linear PROMs is particularly attractive when multiple optimization instances must be carried out as in robust optimization, multi-objective optimization, and global optimization with multiple initial solutions. The work described here advances the state of the art of this concept, particularly for the solution of the aforementioned class of PDE-constrained optimization problems by a gradient-based NAND approach. Specifically, it contributes a parameter sampling procedure based on a saturation assumption for the construction of a cost-efficient PROM database. It also contributes a novel algorithm for interpolating the sensitivities of a pointwise, linear PROM with respect to parameter variations. Both of these contributions are independent of the specific method chosen for constructing a PROM. Therefore, they are applicable in the context of the Proper Orthogonal Decomposition (POD) method, the Reduced Basis (RB) method, the Proper Generalized Decomposition (PGD) method, and many other model reduction methods. Finally, the work described in this paper demonstrates the potential of the concept of a database of pointwise, linear PROMs equipped with these contributions for the identified class of applications. For this last purpose, it reports on the application of the proposed computational framework to the aerodynamic shape optimization of a nonlinear aeroelastic wing under a linear stability (flutter) constraint – a problem which is normally untractable when linearized CFD and CSD models are preferred for representing the PDE constraint.

2. PDE-constrained optimization

2.1. Problem formulation

Consider a time-independent but otherwise general PDE-constrained optimization problem, where the objective function may be linear or nonlinear and at least one of the parametric PDE constraints is linear. This problem may be also governed by additional linear or nonlinear algebraic constraints. The main objective of this paper is to present a computational framework for accelerating, in this context, the enforcement of any parametric, linear, PDE constraint whose discretization is high-dimensional. For simplicity, but without any loss of generality, assume that the optimization problem contains only one

such PDE constraint (the extension of the proposed computational framework to multiple parametric, linear, PDE constraints is straightforward).

The general form of a discretized, time-independent, μ -parametric, linear PDE can be written as

$$\mathbf{R}_L(\mathbf{w}_L(\mu), \mu) = \mathbf{A}(\mu)\mathbf{w}_L(\mu) - \mathbf{b}(\mu) = \mathbf{0} \quad (1)$$

where: the design parameter vector $\mu \in \mathcal{D} \subset \mathbb{R}^{N_D}$ is referred to throughout the remainder of this paper as the parameter “point” in \mathcal{D} , $\mathbf{A}(\mu) \in \mathbb{R}^{N_{w_L} \times N_{w_L}}$ and $\mathbf{b}(\mu) \in \mathbb{R}^{N_{w_L}}$ are the governing parametric matrix and right-hand side arising from the discretization of this PDE by N_{w_L} degrees of freedom (dof)s; and $\mathbf{w}_L(\mu) \in \mathbb{R}^{N_{w_L}}$ denotes its parametric solution. The parameter point μ may represent, for example, the material properties, shape design parameters, or boundary conditions of a physical system of interest.

For the sake of simplicity, it is also assumed here and in the remainder of this paper that when μ is varied, the size and topology of the computational mesh underlying the discrete equation (1) remain unchanged. The extension of the computational framework presented in this paper to the case where this size and topology vary when μ is varied can be performed using the computational approaches presented in [15].

The general PDE-constrained optimization problem outlined above can be formulated as follows

$$\begin{aligned} \min_{\mathbf{w}_L, \mathbf{w}_{NL}, \mu} \quad & f(\mathbf{w}_L(\mu), \mathbf{w}_{NL}(\mu), \mu) \\ \text{s.t.} \quad & \mathbf{c}(\mathbf{w}_L(\mu), \mathbf{w}_{NL}(\mu), \mu) \leq \mathbf{0} \\ & \mathbf{R}_L(\mathbf{w}_L(\mu), \mu) = \mathbf{A}(\mu)\mathbf{w}_L(\mu) - \mathbf{b}(\mu) = \mathbf{0} \\ & \mathbf{R}_{NL}(\mathbf{w}_{NL}(\mu), \mu) = \mathbf{0} \end{aligned} \quad (2)$$

where: f denotes the objective function and is assumed here for simplicity to be a scalar function; $\mathbf{w}_{NL} \in \mathbb{R}^{N_{w_{NL}}}$ is the $N_{w_{NL}}$ -dimensional solution vector of the set of discretized nonlinear PDEs $\mathbf{R}_{NL}(\cdot, \cdot) = \mathbf{0}$ that may be present or absent from (2); and $\mathbf{c}(\cdot, \cdot)$ is a general set of linear and/or nonlinear algebraic constraints including, for instance, linear constraints on \mathbf{w}_L and/or \mathbf{w}_{NL} or box constraints on μ .

Because this paper focuses on the NAND solution approach, problem (2) is viewed here as an optimization problem defined over the parameter point $\mu \in \mathcal{D}$ only; the state vectors \mathbf{w}_L and \mathbf{w}_{NL} are considered as functions of μ that are implicitly defined by the discretized PDE constraints. When the optimizer queries a value of μ , an appropriate PDE solver is applied to compute each of \mathbf{w}_L and \mathbf{w}_{NL} in a nested fashion. This is in contrast to the solution of problem (2) by the Simultaneous Analysis and Design (SAND) approach, where the state variables of (2) are considered in this case as optimization parameters.

Remark. While the computational framework proposed in this paper for accelerating the enforcement of the discretized, parametric, linear PDE constraint $\mathbf{A}(\mu)\mathbf{w}_L(\mu) - \mathbf{b}(\mu) = \mathbf{0}$ is described in the context of the solution of problem (2) by the NAND approach, it is equally applicable in the context of the solution of this problem by the SAND approach. However, its performance may not scale well with the substantially larger number of optimization parameters that is typically encountered in the latter case.

2.2. Gradient-based optimization

Because the discretized PDE constraints introduced above may lead to large-scale systems of equations, a gradient-based method is preferred for the solution of problem (2). To this end, this section reviews the computation of the relevant gradients, in order to keep this paper as self-contained as possible.

Let

$$\mathbf{w} = \begin{bmatrix} \mathbf{w}_L \\ \mathbf{w}_{NL} \end{bmatrix} \in \mathbb{R}^{N_{w_L} + N_{w_{NL}}}$$

and

$$\mathbf{R}(\mathbf{w}(\mu), \mu) = \begin{bmatrix} \mathbf{R}_L(\mathbf{w}_L(\mu), \mu) \\ \mathbf{R}_{NL}(\mathbf{w}_{NL}(\mu), \mu) \end{bmatrix} = \mathbf{0} \quad (3)$$

A gradient-based optimization algorithm incurs the computation of the first derivatives of the objective function $f(\mathbf{w}(\mu), \mu)$, and those of the constraints $c_i(\mathbf{w}(\mu), \mu)$, $i = 1, \dots, N_c$. Let q denote the generic representation of f , any constraint c_i , and any QoI whose derivatives with respect to the components of the parameter point $\mu \in \mathcal{D} \subset \mathbb{R}^{N_D}$ must be computed. Using the chain rule, the following relations can be derived

$$\frac{dq}{d\mu_i}(\mathbf{w}(\mu), \mu) = \frac{\partial q}{\partial \mu_i}(\mathbf{w}(\mu), \mu) + \frac{\partial q}{\partial \mathbf{w}}(\mathbf{w}(\mu), \mu) \frac{\partial \mathbf{w}}{\partial \mu_i}(\mu), \quad i = 1, \dots, N_D \quad (4)$$

For a given q , the computation of the partial derivatives $\frac{\partial q}{\partial \mathbf{w}}(\mathbf{w}(\boldsymbol{\mu}), \boldsymbol{\mu})$ and $\frac{\partial q}{\partial \mu_i}(\mathbf{w}(\boldsymbol{\mu}), \boldsymbol{\mu})$ is straightforward. On the other hand, the computation of $\frac{\partial \mathbf{w}}{\partial \mu_i}(\boldsymbol{\mu})$ warrants some attention.

The differentiation of the discretized, parametric, PDE constraint (3) with respect to any parameter component μ_i can be written as

$$\frac{d\mathbf{R}}{d\mu_i}(\mathbf{w}(\boldsymbol{\mu}), \boldsymbol{\mu}) = \frac{\partial \mathbf{R}}{\partial \mu_i}(\mathbf{w}(\boldsymbol{\mu}), \boldsymbol{\mu}) + \frac{\partial \mathbf{R}}{\partial \mathbf{w}} \frac{\partial \mathbf{w}}{\partial \mu_i}(\mathbf{w}(\boldsymbol{\mu}), \boldsymbol{\mu}) = \mathbf{0}, \quad i = 1, \dots, N_{\mathcal{D}}$$

which leads to

$$\frac{\partial \mathbf{w}}{\partial \mu_i}(\mathbf{w}(\boldsymbol{\mu}), \boldsymbol{\mu}) = - \left[\frac{\partial \mathbf{R}}{\partial \mathbf{w}} \right]^{-1} \frac{\partial \mathbf{R}}{\partial \mu_i}(\mathbf{w}(\boldsymbol{\mu}), \boldsymbol{\mu}), \quad i = 1, \dots, N_{\mathcal{D}} \quad (5)$$

The differentiation of the discretized, parametric, linear constraint (1) with respect to any parameter component μ_i leads to the following linear system of equations governing the sensitivities of \mathbf{w}

$$\mathbf{A}(\boldsymbol{\mu}) \frac{\partial \mathbf{w}}{\partial \mu_i}(\boldsymbol{\mu}) = \frac{\partial \mathbf{b}}{\partial \mu_i}(\boldsymbol{\mu}) - \frac{\partial \mathbf{A}}{\partial \mu_i}(\boldsymbol{\mu}) \mathbf{w}(\boldsymbol{\mu}), \quad i = 1, \dots, N_{\mathcal{D}}$$

These state vector sensitivities can be used to compute the gradients of the objective function and constraints. Indeed, substituting the solution of (5) into (4) leads to

$$\frac{dq}{d\mu_i} = \frac{\partial q}{\partial \mu_i} - \frac{\partial q}{\partial \mathbf{w}} \left(\left[\frac{\partial \mathbf{R}}{\partial \mathbf{w}} \right]^{-1} \frac{\partial \mathbf{R}}{\partial \mu_i} \right), \quad i = 1, \dots, N_{\mathcal{D}} \quad (6)$$

which can also be written as

$$\frac{dq}{d\mu_i} = \frac{\partial q}{\partial \mu_i} - \left[\frac{\partial \mathbf{R}}{\partial \mathbf{w}}^{-T} \frac{\partial q}{\partial \mathbf{w}}^T \right]^T \frac{\partial \mathbf{R}}{\partial \mu_i}, \quad i = 1, \dots, N_{\mathcal{D}} \quad (7)$$

where the superscript denotes here and in the remainder of this paper the transpose operation. Equations (6) and (7) describe the *direct* and *adjoint* approaches for computing the set of sensitivities $\left\{ \frac{dq}{d\mu_i}(\mathbf{w}(\boldsymbol{\mu}), \boldsymbol{\mu}) \right\}_{i=1}^{N_{\mathcal{D}}}$, respectively:

1. In the direct approach, the state sensitivities $\frac{\partial \mathbf{w}}{\partial \mu_i}(\boldsymbol{\mu})$ are first computed by solving the linear system of equations (5) for each parameter component μ_i , then the sensitivities $\frac{dq}{d\mu_i}$ are evaluated using (4).
2. In the adjoint approach, the adjoint vector $\lambda_q(\boldsymbol{\mu})$ is first computed by solving the linear system of equations $\frac{\partial \mathbf{R}}{\partial \mathbf{w}}^T \lambda_q(\boldsymbol{\mu}) = \frac{\partial q}{\partial \mathbf{w}}(\mathbf{w}(\boldsymbol{\mu}), \boldsymbol{\mu})$ for $\lambda_q(\boldsymbol{\mu})$, then all sensitivities $\frac{dq}{d\mu_i}$ are computed using (7).

The direct and adjoint approaches require $N_{\mathcal{D}}$ and $N_c + 1$ “solves”, respectively. Hence, if $N_{\mathcal{D}} \leq 1 + N_c$, the direct approach is preferable; otherwise, the adjoint approach is preferable.

Using the gradients computed above, problem (2) can be solved, for example, by a nonlinear optimization algorithm such as Sequential Quadratic Programming (SQP) [16] equipped with a quasi-Newton approximation of the Hessian matrix, the trust-region method [17], or the interior-point method [18]. In either case, the treatment of the discretized PDE constraints is computationally intensive. For example, computing the sensitivities of the linear PDE's state vector with respect to the optimization parameters requires the solution of the linear system of equations (1) and that of another, large-scale linear system of equations of size N_{w_L} . To significantly reduce this cost, model reduction can be applied to (1), and reduced-order versions of $\frac{\partial \mathbf{b}}{\partial \mu_i}(\boldsymbol{\mu})$ and $\frac{\partial \mathbf{A}}{\partial \mu_i}(\boldsymbol{\mu})$, $i = 1, \dots, N_{\mathcal{D}}$, can be computed instead of their high-dimensional counterparts.

3. Parametric model order reduction

Again, the focus of this paper is on accelerating the solution of the optimization problem (2), in the case where the discretized, parametric, PDE constraint (1) is high-dimensional. In this case, solving the optimization problem (2) can be prohibitively expensive. For this purpose, model reduction is applied to the discrete system of linear constraints in (3) in order to reduce its dimension, which reduces the cost of enforcing these constraints and therefore reduces the cost of solving problem (2).

Algorithm 1 Proper orthogonal decomposition.**Input:** Snapshot matrix $\mathbf{X}(\boldsymbol{\mu}) \in \mathbb{R}^{N_w \times N_s}$; desired ROB dimension N_{w_r} **Output:** ROB $\mathbf{V}(\boldsymbol{\mu})$ 1: Compute the thin SVD of $\mathbf{X}(\boldsymbol{\mu})$: $\mathbf{X}(\boldsymbol{\mu}) = \mathbf{U}\mathbf{\Sigma}\mathbf{D}$, where $\mathbf{U} = [\mathbf{u}_1 \ \mathbf{u}_2 \ \cdots \ \mathbf{u}_{N_s}]$ 2: $\mathbf{V}(\boldsymbol{\mu}) = [\mathbf{u}_1 \ \mathbf{u}_2 \ \cdots \ \mathbf{u}_{N_{w_r}}]$ **3.1. Projection-based model order reduction and representation**

Projection-based model order reduction reduces the dimension of a $\boldsymbol{\mu}$ -parametric HDM such as (1), which governs the state vector \mathbf{w}_L , by performing the subspace approximation

$$\mathbf{w}_L(\boldsymbol{\mu}) \approx \mathbf{V}(\boldsymbol{\mu})\mathbf{w}_{L_r}(\boldsymbol{\mu}) \quad (8)$$

where $\mathbf{V}(\boldsymbol{\mu}) \in \mathbb{R}^{N_{w_L} \times N_{w_r}}$ is a *right* ROB of dimension $N_{w_r} \ll N_{w_L}$, and $\mathbf{w}_{L_r} \in \mathbb{R}^{N_{w_{L_r}}}$ is a vector of generalized coordinates referred to as the reduced-order state vector. The notation $\mathbf{V}(\boldsymbol{\mu})$ emphasizes that \mathbf{V} is a pointwise ROB – that is, a ROB constructed at the parameter point $\boldsymbol{\mu} \in \mathcal{D}$ – rather than a global ROB constructed to deliver accurate subspace approximations at any parameter point in the design parameter space. Typically, $\mathbf{V}(\boldsymbol{\mu})$ is constructed using the POD method of snapshots [2], which can be summarized as the collection of solution snapshots into a matrix, followed by the compression of this matrix using the Singular Value Decomposition (SVD) algorithm (see Algorithm 1).

Substituting (8) into (1) and pre-multiplying the resulting system of equations by $\mathbf{W}(\boldsymbol{\mu})^T$, where $\mathbf{W}(\boldsymbol{\mu}) \in \mathbb{R}^{N_{w_L} \times N_{w_{L_r}}}$ is known as the *left* ROB associated with the right ROB $\mathbf{V}(\boldsymbol{\mu})$, transforms this system into a square counterpart and leads to the parametric, linear, Petrov-Galerkin PROM

$$\mathbf{W}(\boldsymbol{\mu})^T \mathbf{A}(\boldsymbol{\mu}) \mathbf{V}(\boldsymbol{\mu}) \mathbf{w}_{L_r} = \mathbf{W}(\boldsymbol{\mu})^T \mathbf{b}(\boldsymbol{\mu}) \quad (9)$$

This PROM can be described by the low-dimensional matrix $\mathbf{A}_r(\boldsymbol{\mu}) = \mathbf{W}(\boldsymbol{\mu})^T \mathbf{A}(\boldsymbol{\mu}) \mathbf{V}(\boldsymbol{\mu}) \in \mathbb{R}^{N_{w_{L_r}} \times N_{w_{L_r}}}$, the low-dimensional vector $\mathbf{b}_r(\boldsymbol{\mu}) = \mathbf{W}(\boldsymbol{\mu})^T \mathbf{b}(\boldsymbol{\mu}) \in \mathbb{R}^{N_{w_{L_r}}}$, and therefore the doublet $(\mathbf{A}_r(\boldsymbol{\mu}), \mathbf{b}_r(\boldsymbol{\mu}))$. If $\mathbf{W}(\boldsymbol{\mu}) = \mathbf{V}(\boldsymbol{\mu})$, this PROM becomes a Galerkin PROM.

The sensitivities of the reduced-order state vector \mathbf{w}_{L_r} with respect to the parameter components $\mu_i, i = 1, \dots, N_{\mathcal{D}}$, can be computed by following the same approach described in Section 2.2. In this case, this approach entails the computation of the reduced-order sensitivity matrices and sensitivity vectors $\frac{\partial \mathbf{A}_r}{\partial \mu_i}(\boldsymbol{\mu}) \in \mathbb{R}^{N_{w_{L_r}} \times N_{w_{L_r}}}$ and $\frac{\partial \mathbf{b}_r}{\partial \mu_i}(\boldsymbol{\mu}) \in \mathbb{R}^{N_{w_{L_r}}}$, respectively, $i = 1, \dots, N_{\mathcal{D}}$. It follows that the PROM (9) and its associated reduced-order sensitivity equations can be collectively represented by the tuple of low-dimensional operators

$$\mathcal{A}_r(\boldsymbol{\mu}) = \left\{ \mathbf{A}_r(\boldsymbol{\mu}), \mathbf{b}_r(\boldsymbol{\mu}), \frac{\partial \mathbf{A}_r}{\partial \mu_i}(\boldsymbol{\mu}), \frac{\partial \mathbf{b}_r}{\partial \mu_i}(\boldsymbol{\mu}), \quad i = 1, \dots, N_{\mathcal{D}} \right\} \quad (10)$$

In summary, the subspace approximation (8) transforms the original constrained optimization problem (2) into the PROM-constrained optimization problem

$$\begin{array}{ll} \min_{\mathbf{w} \in \mathcal{W}, \boldsymbol{\mu} \in \mathcal{D}} & f(\mathbf{V}(\boldsymbol{\mu})\mathbf{w}_{L_r}(\boldsymbol{\mu}), \mathbf{w}_{NL}(\boldsymbol{\mu}), \boldsymbol{\mu}) \\ \text{s.t.} & \mathbf{c}(\mathbf{V}(\boldsymbol{\mu})\mathbf{w}_{L_r}(\boldsymbol{\mu}), \mathbf{w}_{NL}(\boldsymbol{\mu}), \boldsymbol{\mu}) \leq \mathbf{0} \\ & \mathbf{A}_r(\boldsymbol{\mu})\mathbf{w}_{L_r}(\boldsymbol{\mu}) = \mathbf{b}_r(\boldsymbol{\mu}) \\ & \mathbf{R}_{NL}(\mathbf{w}_{NL}(\boldsymbol{\mu}), \boldsymbol{\mu}) = \mathbf{0} \end{array} \quad (11)$$

where the discretized, parametric, linear PDE constraint and its associated sensitivity equations are low-dimensional. Just like the original constrained optimization problem (2), problem (11) can be solved using any preferred, gradient-based optimization algorithm.

3.2. Concept of a database of pointwise linear PROMs and associated sensitivities

As stated in Section 1, a PROM constructed at a parameter point $\boldsymbol{\mu} \in \mathcal{D}$ does not necessarily perform well at a different parameter point $\boldsymbol{\mu}'$ in this space. Addressing this issue by constructing a new PROM at each newly queried parameter point $\boldsymbol{\mu}'$ is computationally inefficient, as it entails: the computation of new solution snapshots at the parameter point $\boldsymbol{\mu}'$; the compression of these snapshots using, for example, the SVD algorithm in order to compute a new right ROB $\mathbf{V}(\boldsymbol{\mu}')$; and performing additional matrix-matrix and matrix-vector computations to compute, in the context of this paper, the tuple $\mathcal{A}_r(\boldsymbol{\mu}') = \left\{ \mathbf{A}_r(\boldsymbol{\mu}'), \mathbf{b}_r(\boldsymbol{\mu}'), \frac{\partial \mathbf{A}_r}{\partial \mu_i}(\boldsymbol{\mu}'), \frac{\partial \mathbf{b}_r}{\partial \mu_i}(\boldsymbol{\mu}'), \quad i = 1, \dots, N_{\mathcal{D}} \right\}$ representing the linear PROM at the parameter point $\boldsymbol{\mu}'$.

An alternative methodology for addressing efficiently the parameter dependence of a linear PROM based on the concept of a database of reduced-order information was proposed and developed in [19], and more recently refined in [15]. It consists in: constructing offline a database of pointwise, linear PROMs; equipping this database with a family of algorithms for interpolation on matrix manifolds [20,13,14]; and applying these algorithms online to build in real time a PROM at a queried but unsampled parameter point $\tilde{\mu} \in \mathcal{D} \subset \mathbb{R}^{N_D}$. In the context of the parametric, linear PROM (9), this offline-online methodology can be described as follows:

1. Offline

- Sample $N_{\mathcal{DB}}$ parameter points μ^j in the given design parameter space \mathcal{D} .
- At each sampled parameter point $\mu^j, j = 1, \dots, N_{\mathcal{DB}}$, construct a pointwise, linear PROM $\mathcal{A}_r(\mu^j) = \left\{ \mathbf{A}_r(\mu^j), \mathbf{b}_r(\mu^j), \frac{\partial \mathbf{A}_r}{\partial \mu_i}(\mu^j), \frac{\partial \mathbf{b}_r}{\partial \mu_i}(\mu^j), i = 1, \dots, N_D \right\}$.
- Store the set of pre-computed linear PROMs in a database \mathcal{DB} (in practice, the derivatives $\frac{\partial \mathbf{A}_r}{\partial \mu_i}(\mu^j)$ and $\frac{\partial \mathbf{b}_r}{\partial \mu_i}(\mu^j)$ need not be stored, because they can be computed online as explained in Section 5).

2. Online

For each queried but unsampled parameter point $\tilde{\mu}$, construct a pointwise, linear PROM $\mathcal{A}_r(\tilde{\mu}) = \left\{ \mathbf{A}_r(\tilde{\mu}), \mathbf{b}_r(\tilde{\mu}), \frac{\partial \mathbf{A}_r}{\partial \mu_i}(\tilde{\mu}), \frac{\partial \mathbf{b}_r}{\partial \mu_i}(\tilde{\mu}), i = 1, \dots, N_D \right\}$ as follows:

- For each reduced-order quantity in $\mathcal{A}_r(\mu)$, identify *a priori* a matrix manifold where it lies, for example, by identifying the most important algebraic properties characterizing this reduced-order quantity.
- Interpolate in real time the pre-computed reduced-order quantities on their identified manifolds, at the queried but unsampled parameter point $\tilde{\mu}$.

In order to extend the scope of this concept of a database of pointwise, linear PROMs to the real-time solution of the constraint equation (1), in view of accelerating the solution of the optimization problem (2), the state of the art of this concept is next advanced as follows. First, a computationally efficient parameter sampling procedure based on a saturation assumption is described. Then, a novel algorithm for interpolating in real time the sensitivities of a pointwise, linear PROM with respect to parameter variations is presented.

4. Parameter sampling using a residual-based error indicator

The typical construction of an $N_{\mathcal{DB}}$ -point database $\mathcal{DB} = \{\mu^j, \mathcal{A}_r(\mu^j)\}_{j=1}^{N_{\mathcal{DB}}}$ for the purpose of supporting the concept of interpolation on matrix manifolds, which addresses the parameter dependence of a linear PROM, is guided by a parameter sampling procedure that selects the parameter points in the design parameter space \mathcal{D} where to pre-compute linear PROMs of interest. Specifically, such a procedure should lead to a database \mathcal{DB} that has the smallest possible size $N_{\mathcal{DB}}$, in order to minimize as much as possible the computational cost of the offline stage. It should also populate this database in a manner that delivers interpolation accuracy in the entirety of \mathcal{D} – that is, $\forall \mu \in \mathcal{D}$.

Parameter sampling procedures can be classified in two categories: non-adaptive procedures that sample the given design parameter space *a priori*; and adaptive, error-informed procedures that incorporate a feedback in the sampling process. Non-adaptive parameter sampling procedures are typically easier to implement. Adaptive counterparts however can be expected to perform better at producing databases \mathcal{DB} that satisfy the desiderata outlined above.

The simplest non-adaptive parameter sampling procedure is perhaps the so-called Full Factorial (FF) sampling procedure, which samples the given design parameter space on an N_D -dimensional grid. In principle, this grid can be made sufficiently fine to deliver accurate approximations such as those to be performed here using interpolation on a matrix manifold. However, the FF sampling procedure rapidly becomes unfeasible, as the number of parameter points it requires for achieving accuracy grows exponentially with $N_{\mathcal{DB}}$. The alternative Latin Hypercube Sampling (LHS) procedure, which is a non-adaptive but sparse procedure, scales only linearly with $N_{\mathcal{DB}}$. However, because it has no explicit awareness of where the approximated solution will be inaccurate, it typically leads to oversampling when attempting to meet the aforementioned accuracy desideratum.

It follows that for high-dimensional design parameter spaces, an adaptive approach where the parameter points are sampled in an incremental manner and the location of the next sampled parameter point is identified by drawing information from the already sampled parameter points is often necessary. Such an approach is typically iterative and referred to as a greedy procedure. It initializes the construction of the database by populating it with one or a small number of parameter points located, for example, in the center of \mathcal{D} and/or on its boundaries. At each iteration, it considers a set of candidate parameter points for sampling and selects that parameter point where some error estimator or indicator $e(\mu, \mathcal{DB})$ is maximized [7,21–26]. After n greedy iterations, it produces a database \mathcal{DB}_n that can be characterized as follows

$$\mathcal{DB}_n = \{\mathcal{DB}_{n-1}, \mu^n, \mathcal{A}_r(\mu^n)\} \quad \text{where} \quad \mu^n = \arg \max_{\mu \in \Xi} e(\mu, \mathcal{DB}_{n-1})$$

where $\Xi = \{\mu_c^1, \dots, \mu_c^{N_\Xi}\}$ denotes the set of N_Ξ candidate parameter points in the design parameter space \mathcal{D} . This set can be fixed throughout the iterations or refreshed at each iteration, but in any case, should be chosen large enough to faithfully represent \mathcal{D} . The iterations performed by such a greedy procedure are terminated when the error estimator or indicator falls below a specified threshold, or the pre-set maximum size of the database or maximum run-time allocated for the execution of the procedure has been reached.

The most pertinent error for an adaptive sampling procedure is the true error $e(\mu_c^j) = \|\mathbf{w}_L(\mu_c^j) - \mathbf{w}_L^j(\mu_c^j)\| = \|\mathbf{w}_L(\mu_c^j) - \tilde{\mathbf{V}}(\mu_c^j)\mathbf{w}_{L_r}^j(\mu_c^j)\|$, where $\mathbf{w}_L(\mu_c^j)$ and $\mathbf{w}_{L_r}^j(\mu_c^j)$ denote the solutions computed at the candidate parameter point $\mu_c^j \in \mathcal{D}$ of the HDM and PROM problems constructed and interpolated at μ_c^j , respectively, and $\tilde{\mathbf{V}}(\mu_c^j)$ is an approximation of $\mathbf{V}(\mu_c^j)$ that can be computed in real time (see Section 4.2). However, whereas interpolating the PROMs, solving the reduced-order problems, and reconstructing the corresponding high-dimensional solutions at the set of candidate parameter points is feasible, because each of these operations can be performed relatively fast if not in real time, the direct computation of the high-dimensional solutions for the entire set Ξ is typically prohibitively expensive. Consequently, the true error is replaced in practice by: an estimation of this error that is based on a proven error bound, when such a bound is available; or by an error indicator otherwise. Error bounds are known for the solutions of some relatively simple elliptic [27,26], parabolic [23], and hyperbolic [28] PDEs, as well as for the solutions of some simple linear time-invariant systems [29]. However, these bounds typically require the computation of an inf-sup constant, which is seldom a simple task, and are not usually tight [30]. For this reason, computationally feasible error indicators based on some norm of a relevant residual evaluated using the reconstructed high-dimensional solution are used instead [21,22]. In the context of this paper, such an error indicator can be expressed as $\|\mathbf{R}_L(\tilde{\mathbf{V}}(\mu_c^j)\mathbf{w}_{L_r}^j(\mu_c^j))\|$, where \mathbf{R}_L is defined in (2). Nevertheless, even the computation of such error indicators can rapidly become overwhelming when N_Ξ becomes very large, for example, because the dimension $N_{\mathcal{D}}$ of the design parameter space \mathcal{D} is increased. One approach for addressing this issue is to reduce as much as possible the number of parameter points that must be sampled to achieve the desired accuracy – that is, to minimize as much as possible the size $N_{\mathcal{DB}}$ of the database \mathcal{DB} . To this end, an accelerated parameter sampling procedure based on an appropriate saturation assumption is proposed next.

4.1. Accelerated parameter sampling procedure based on a saturation assumption

Here, the set of N_Ξ candidate parameter points in the design parameter space \mathcal{D} , $\Xi = \{\mu_c^1, \dots, \mu_c^{N_\Xi}\}$, is fixed throughout all iterations of the greedy procedure. In this case, this parameter sampling procedure can be further accelerated by using a saturation technique where previous evaluations of the error indicator are used to prevent, whenever possible, unnecessary additional evaluations of this indicator. This is possible if the accuracy of the database can be assumed to increase only to a certain extent as more parameter points are sampled, as this assumption makes it possible to omit the evaluation of the error indicator at some parameter points.

The incorporation of a saturation assumption in a greedy procedure was first proposed in [24]. Here, the efficiency of this technique is further improved by applying it at each m -th iteration of the greedy procedure not to Ξ , but to a random subset Π_m of Ξ . Specifically, a fixed number of N_Π parameter points are randomly selected in Ξ at each greedy iteration to assess the accuracy of the database being constructed, and a saturation constant is defined as follows.

Definition 1. Saturation Constant

Let $e(\mu; \mathcal{DB}_m)$ denote an error indicator depending on the design parameter vector μ and the database \mathcal{DB}_m . Let also $\{\mathcal{DB}_m\}_{m=1}^n$ denote a sequence of nested databases – that is, a set of databases satisfying $\mathcal{DB}_m \subset \mathcal{DB}_n$ for all $1 \leq m < n$. The saturation constant $\tau_s > 0$ is defined as

$$\tau_s = \sup_{1 \leq m < n, \mu \in \mathcal{D}} \frac{e(\mu; \mathcal{DB}_n)}{e(\mu; \mathcal{DB}_m)} \quad (12)$$

From the above definition, it follows that: $e(\mu; \mathcal{DB}_n) \leq \tau_s e(\mu; \mathcal{DB}_m)$, for all $1 \leq m < n$; and $\tau_s < 1$ implies that $e(\mu; \mathcal{DB}_n) < e(\mu; \mathcal{DB}_m)$ for all $1 \leq m < n$. In other words, $e(\mu; \cdot)$ strictly decreases as more parameter points are sampled during the construction of the database. Similarly, $\tau_s = 1$ implies a monotone decrease in $e(\mu; \cdot)$ as the number of sampled parameter points is increased during the construction of the database. If $\tau_s > 1$, $e(\mu; \cdot)$ may increase at some parameter point $\mu \in \mathcal{DB}$ for certain iterations of the greedy procedure.

Now, consider the additional two definitions:

- Let $e_{\text{profile}}(\mu_c)$ denote the most recently computed error indicator at the candidate parameter point $\mu_c \in \Pi_m \subset \Xi$, where Π_m is a subset of Ξ randomly selected during the m -th iteration of the greedy procedure. Note that because of the randomness aspect of the construction of the subset of candidate parameter points Π_m , μ_c might not have been selected at iteration $m - 1$ for error indication. Hence in general, $e_{\text{profile}}(\mu_c) \neq e(\mu_c; \mathcal{DB}_{m-1})$. Consequently, if μ_c was not previously selected for computing the error indicator, $e_{\text{profile}}(\mu_c)$ is set to ∞ .

Algorithm 2 Greedy procedure incorporating a randomized saturation assumption.

Input: A set of candidate parameter points $\Xi \subset \mathcal{D}$, its size N_Ξ , a tolerance $\epsilon_{tol} > 0$, a saturation constant τ_s , and two sizes N_Π and $N'_\Pi > N_\Pi$ for the two subsets of Ξ to be randomly selected at each greedy iteration

Output: PROM database \mathcal{DB}_m

```

1: Choose an initial parameter point  $\mu_c^1 \in \Xi$  and compute  $\mathcal{A}_r(\mu_c^1)$ 
2: Set  $\mathcal{DB}_1 = \{(\mu_c^1, \mathcal{A}_r(\mu_c^1))\}$ 
3: Set  $e_{\text{profile}}(\mu) = \infty$  for all  $\mu \in \Xi$ 
4: while TRUE do
5:   Generate a random subset  $\Pi_m$  of  $\Xi$  of dimension  $N_\Pi$ 
6:   Set  $e_{\text{max}}^{(l)} = 0$ 
7:   for  $\mu_c^{k_l} \in \Pi_m, l = 1, \dots, N_\Pi$  do
8:     if  $\tau_s e_{\text{profile}}(\mu_c^{k_l}) > e_{\text{max}}^{(l)}$  then
9:       Compute  $e(\mu_c^{k_l}; \mathcal{DB}_m)$  and set  $e_{\text{profile}}(\mu_c^{k_l}) = e(\mu_c^{k_l}; \mathcal{DB}_m)$ 
10:    if  $e(\mu_c^{k_l}; \mathcal{DB}_m) > e_{\text{max}}^{(l)}$  then
11:      Set  $e_{\text{max}}^{(l)} = e(\mu_c^{k_l}; \mathcal{DB}_m)$  and  $\mu_{m+1} = \mu_c^{k_l}$ 
12:    end if
13:  end if
14: end for
15: if  $e_{\text{max}}^{(l)} < \epsilon_{tol}$  then
16:   Check convergence using the larger subset of candidate parameter points of size  $N'_\Pi > N_\Pi$ 
17:   if convergence is reached then
18:     Terminate the parameter sampling procedure
19:   end if
20: end if
21: Compute  $\mathcal{A}_r(\mu_c^{m+1})$ 
22: Set  $\mathcal{DB}_{m+1} = \mathcal{DB}_m \cup \{(\mu_c^{m+1}, \mathcal{A}_r(\mu_c^{m+1}))\}$ 
23:  $m \leftarrow m + 1$ 
24: end while

```

- The following definition is associated with a subiterative process within the m -th iteration of the greedy procedure. Let $e_{\text{max}}^{(l)}(\mathcal{DB}_m)$ denote the current maximum indicated error at subiteration l of iteration m of the greedy procedure, where one error indication is computed at a candidate parameter point in $\Pi_m \subset \Xi$. This iterative maximum indicated error is computed as follows

$$e_{\text{max}}^{(l)}(\mathcal{DB}_m) = \begin{cases} 0 & \text{for } l = 1 (\text{initialization}) \\ \max_{j=k_1, \dots, k_{l-1}} e(\mu_c^j; \mathcal{DB}_m) & \text{for } 1 < l \leq N_\Pi \text{ and } \mu_c^j \in \Pi_m \end{cases}$$

where l also denotes a local indexing of a candidate parameter point within Π_m , and k_l denotes the corresponding global indexing of this candidate parameter point within $\Xi \supset \Pi_m$.

Based on the two above definitions and that of the saturation constant (12), it follows that if $\tau_s e_{\text{profile}}(\mu_c^i) < e_{\text{max}}^{(l)}(\mathcal{DB}_m)$, then $e(\mu_c^i; \mathcal{DB}_m)$ is guaranteed to be less than $e_{\text{max}}^{(l)}(\mathcal{DB}_m)$ prior to the evaluation of the error indicator at μ_c^i . This is because $e(\mu_c^i; \mathcal{DB}_m) \leq \tau_s e_{\text{profile}}(\mu_c^i)$, due to the definition of the saturation constant. Hence, it is not necessary to evaluate $e(\mu_c^i; \mathcal{DB}_m)$ if $\tau_s e_{\text{profile}}(\mu_c^i) < e_{\text{max}}^{(l)}(\mathcal{DB}_m)$. This implies that the simple inequality check between $\tau_s e_{\text{profile}}(\mu_c^i)$ and $e_{\text{max}}^{(l)}(\mathcal{DB}_m)$, whose values are readily available, can avoid some unnecessary evaluations of the error indicator. Then, convergence of the greedy procedure can be monitored as follows. A check is performed to verify whether $e_{\text{max}}^{(l)}(\mathcal{DB}_m) < \epsilon_{tol}$ or not, where ϵ_{tol} is a specified threshold tolerance for assessing the accuracy of the constructed database. For this purpose, the error indicator is evaluated at all parameter points of a larger subset of Ξ of size $N'_\Pi > N_\Pi$ that is also randomly selected. If $e_{\text{max}}^{(l)}(\mathcal{DB}_m) < \epsilon_{tol}$ holds for this larger subset of candidate parameters, the greedy procedure is terminated.

In general, the value of τ_s is set to 1 or 2: 1 for an aggressive approach; and 2 for a conservative approach.

Algorithm 2 summarizes the greedy procedure for parameter sampling equipped with the saturation technique.

4.2. Computation of the residual-based error indicator

To complete the description of the greedy parameter sampling procedure proposed in this paper for guiding the construction of a database of linear, pointwise PROMs, it remains to discuss one subtlety regarding the practical computation of a residual-based error indicator such as

$$\left\| \mathbf{R}_L \left(\mathbf{V}(\mu_c^j) \mathbf{W}_{L_r}^L(\mu_c^j) \right) \right\| \quad (13)$$

where all quantities appearing in (13) have been defined in the introduction part of Section 4. Note that neither the right ROB $\mathbf{V}(\mu)$ nor the left counterpart $\mathbf{W}(\mu)$ is part of the definition of the tuple $\mathcal{A}_r(\mu) = \left\{ \mathbf{A}_r(\mu), \mathbf{b}_r(\mu), \frac{\partial \mathbf{A}_r}{\partial \mu_i}(\mu), \frac{\partial \mathbf{b}_r}{\partial \mu_i}(\mu) \right\}$,

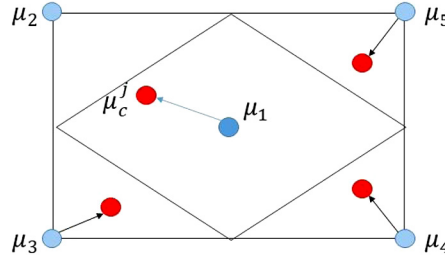


Fig. 1. Greedy parameter sampling procedure: evaluation of the residual-based error indicator using a constant extrapolation approach for approximating $\mathbf{V}_{\mu_c^j}$.

$i = 1, \dots, N_{\mathcal{D}} \}$ representing the linear PROM constructed or reconstructed at the parameter point μ . Note also that the computational complexity of the construction of a ROB such as $\mathbf{V}(\mu)$, for $\mu \in \mathcal{D}$, scales with the dimension of the HDM: therefore, the construction of such a ROB at each candidate parameter point $\mu_c^j \in \Xi$ is not feasible. Instead, a reasonable approximation $\tilde{\mathbf{V}}(\mu_c^j)$ of $\mathbf{V}(\mu_c^j)$ is sought-after here. For this purpose, computing $\tilde{\mathbf{V}}(\mu_c^j)$ by interpolating on the Grassman manifold [20] the pre-computed ROBs at the previously sampled parameter points, which could be stored in this case in the database \mathcal{DB} , is feasible, has been shown in [20] to deliver a good level of accuracy even in the context of a sparsely populated database \mathcal{DB} , but unfortunately requires the solution of an auxiliary problem whose computational complexity scales with the dimension of the HDM (see [20] for details). Hence, given that in the context of a greedy procedure the level of accuracy of an approximation is not paramount for the selection among the elements of Ξ of the best parameter point to be sampled, it is proposed here to store the aforementioned pre-computed ROBs in \mathcal{DB} and reconstruct $\tilde{\mathbf{V}}(\mu_c^j)$ for each $\mu_c^j \in \Xi$ using a constant extrapolation procedure based on the nearest sampled parameter point μ^j . This approach for computing $\tilde{\mathbf{V}}(\mu_c^j)$, for each $\mu_c^j \in \Xi$, leads to a procedure for constructing the database of interest \mathcal{DB} that can be summarized as follows

$$\mathcal{DB} = \left\{ \mu^j, \mathcal{A}_r(\mu^j), \mathbf{V}(\mu^j) \right\}_{j=1}^{N_{\mathcal{DB}}}$$

and $\mathbf{R}_L \left(\mathbf{V}(\mu_c^j) \mathbf{w}_{L_r}^l(\mu_c^j) \right) \approx \mathbf{R}_L \left(\mathbf{V}(\hat{\mu}) \mathbf{w}_{L_r}^l(\mu_c^j) \right)$, where $\hat{\mu} = \arg \min_{\mu' \in \mathcal{DB}} \|\mu' - \mu\|_2^2$

This approach is not only feasible, but also computationally efficient.

5. Interpolation on a matrix manifold

Suppose that an $N_{\mathcal{DB}}$ -point database $\mathcal{DB} = \left\{ \mu^j, \mathcal{A}_r(\mu^j), \mathbf{V}(\mu^j) \right\}_{j=1}^{N_{\mathcal{DB}}}$ is constructed as described above, for the purpose of accelerating the solution of the reduced-order optimization problem (11) by a gradient-based algorithm. At some point of the optimization process, such algorithm may query an unsampled parameter point $\tilde{\mu} \in \mathcal{D}$. Here, the idea is to construct in real time at $\tilde{\mu}$ the tuple $\mathcal{A}_r(\tilde{\mu})$ representing the linear PROM at this point, by interpolating the content of \mathcal{DB} .

To this end, it is first noted that the straightforward (or direct) interpolation of the operators defining $\mathcal{A}_r(\tilde{\mu})$ – that is, the approach where each entry of each operator is interpolated independently from the others – is bound to deliver inaccurate results for the following reasons:

1. Because the PROM operators stored in \mathcal{DB} may have been constructed using different Petrov-Galerkin bases – that is, for $i \neq j$, $(\{\mathbf{W}(\mu^i), \mathbf{V}(\mu^i)\}) \neq (\{\mathbf{W}(\mu^j), \mathbf{V}(\mu^j)\})$, the reduced-order operators defining the tuples $\mathcal{A}_r(\mu^j)$ may have been constructed for different generalized coordinate systems [14]. In this case, a correct interpolation requires first transforming the content of the database \mathcal{DB} into one that uses a single generalized coordinate system (Fig. 1).
2. The direct interpolation approach does not necessarily preserve some important properties of the PROM operators being interpolated, such as nonsingularity, orthogonality, or positive-semi-definiteness [14,20].

Both issues raised above can be addressed by first “rotating” the PROM operators defining each tuple $\mathcal{A}_r(\tilde{\mu})$ in order to achieve consistency in the generalized coordinate system, then interpolating these operators on appropriate matrix manifolds [14].

5.1. Interpolation of projection-based reduced-order models

5.1.1. Enforcement of consistency

Starting from the observation that any ROB $\mathbf{V}(\mu)$ remains a basis of the same subspace after multiplication by an orthogonal matrix \mathbf{Q} , the following classes of equivalence can be defined for each subspace $\mathcal{S}(\mu) = \text{range}(\mathbf{V}(\mu)\mathbf{Q})$

$$c(\mathbf{V}(\boldsymbol{\mu})) = \{\mathbf{V}(\boldsymbol{\mu})\mathbf{Q} \mid \mathbf{Q} \in \mathcal{O}(N_{w_{Lr}})\}$$

where $\mathcal{O}(N_{w_{Lr}})$ denotes the set of orthogonal matrices in $\mathbb{R}^{N_{w_{Lr}} \times N_{w_{Lr}}}$. Each equivalence class defines a set of bases leading to PROM operators that act in the same subspace, but use for this purpose different systems of generalized coordinates.

The corresponding classes of equivalence for the PROM operators themselves are

$$c\left(\mathbf{A}_r(\boldsymbol{\mu}), \mathbf{b}_r(\boldsymbol{\mu}), \frac{\partial \mathbf{A}_r}{\partial \mu_i}(\boldsymbol{\mu}), \frac{\partial \mathbf{b}_r}{\partial \mu_i}(\boldsymbol{\mu})\right) = \left\{ \mathbf{Q}^T \mathbf{A}_r(\boldsymbol{\mu}) \mathbf{Q}, \mathbf{Q}^T \mathbf{b}_r(\boldsymbol{\mu}), \mathbf{Q}^T \frac{\partial \mathbf{A}_r}{\partial \mu_i}(\boldsymbol{\mu}) \mathbf{Q}, \mathbf{Q}^T \frac{\partial \mathbf{b}_r}{\partial \mu_i}(\boldsymbol{\mu}) \mid \mathbf{Q} \in \mathcal{O}(N_{w_{Lr}}) \right\}$$

Hence, given a reference ROB $\mathbf{V}(\boldsymbol{\mu}^\circ)$ and any other ROB of interest $\mathbf{V}(\boldsymbol{\mu}^j)$, the latter ROB can be transformed into an equivalent basis that uses the same system of generalized coordinates as that used by $\mathbf{V}(\boldsymbol{\mu}^\circ)$, through its right multiplication by the rotation matrix \mathbf{Q} that solves the following minimization problem

$$\min_{\mathbf{Q} \in \mathcal{O}(N_{w_{Lr}})} \left\| \mathbf{V}(\boldsymbol{\mu}^j) \mathbf{Q} - \mathbf{V}(\boldsymbol{\mu}^\circ) \right\|_F^2$$

This problem, which is known as the orthogonal Procrustes problem, has the following analytical solution

$$\mathbf{Q}^j = \mathbf{U}^j \mathbf{Z}^{jT} \quad (14)$$

where \mathbf{U}^j and \mathbf{Z}^j are given by the SVD of the matrix product $\mathbf{V}(\boldsymbol{\mu}^j)^T \mathbf{V}(\boldsymbol{\mu}^\circ)$ – that is,

$$\mathbf{V}(\boldsymbol{\mu}^j)^T \mathbf{V}(\boldsymbol{\mu}^\circ) = \mathbf{U}^j \boldsymbol{\Sigma}^j \mathbf{Z}^{jT}$$

The corresponding transformation of the tuple $\mathcal{A}_r(\boldsymbol{\mu}^j)$ is given by

$$\mathcal{A}_r^*(\boldsymbol{\mu}^j) = \left\{ \underbrace{\mathbf{Q}^{jT} \mathbf{A}_r(\boldsymbol{\mu}^j) \mathbf{Q}^j}_{\mathbf{A}_r^*(\boldsymbol{\mu}^j)}, \underbrace{\mathbf{Q}^{jT} \mathbf{b}_r(\boldsymbol{\mu}^j)}_{\mathbf{b}_r^*(\boldsymbol{\mu}^j)}, \underbrace{\mathbf{Q}^{jT} \frac{\partial \mathbf{A}_r}{\partial \mu_i}(\boldsymbol{\mu}^j) \mathbf{Q}^j}_{\frac{\partial \mathbf{A}_r^*}{\partial \mu_i}(\boldsymbol{\mu}^j)}, \underbrace{\mathbf{Q}^{jT} \frac{\partial \mathbf{b}_r}{\partial \mu_i}(\boldsymbol{\mu}^j)}_{\frac{\partial \mathbf{b}_r^*}{\partial \mu_i}(\boldsymbol{\mu}^j)}, \quad i = 1, \dots, N_D \right\} \quad (15)$$

In summary, given an $N_{\mathcal{DB}}$ -point database $\mathcal{DB} = \{\boldsymbol{\mu}^j, \mathcal{A}_r(\boldsymbol{\mu}^j), \mathbf{V}(\boldsymbol{\mu}^j)\}_{j=1}^{N_{\mathcal{DB}}}$ where all tuples $\mathcal{A}_r(\boldsymbol{\mu}^j)$ are not necessarily consistent – in the sense that they are not constructed for the same generalized coordinate system – the following procedure may be applied to transform the content of this database into an equivalent one that is consistent: first, identify a reference parameter $\boldsymbol{\mu}^\circ$; then, rotate $\mathbf{V}(\boldsymbol{\mu}^j)$ and all PROMs stored in the database using (14) and (15).

5.1.2. Logarithm and exponential maps and interpolation in the tangent space

Let

$$\mathcal{DB}^* = \left\{ \boldsymbol{\mu}^j, \mathcal{A}_r^*(\boldsymbol{\mu}^j), \mathbf{V}^*(\boldsymbol{\mu}^j) \right\}_{j=1}^{N_{\mathcal{DB}}} \quad (16)$$

where $\mathbf{V}^*(\boldsymbol{\mu}^j) = \mathbf{V}(\boldsymbol{\mu}^j) \mathbf{Q}^j$ (see (14)), be an $N_{\mathcal{DB}}$ -point database of consistent, linear PROMs. For each reduced-order operator contributing to the definition of a tuple \mathcal{A}_r^* representing a linear PROM – which in this case, is a reduced-order matrix – let $\mathcal{M} \subset \mathbb{R}^{N_{w_{Lr}} \times M_{w_{Lr}}}$ denote the most important, known, matrix manifold on which this operator lies. For this purpose, any reduced-order vector contributing to the definition of the tuple \mathcal{A}_r^* is viewed here and throughout the remainder of this paper, where appropriate, as a single column matrix (hence, $M_{w_{Lr}} = N_{w_{Lr}}$ or $M_{w_{Lr}} = 1$). Such a manifold \mathcal{M} can be identified, for example, by considering the most important algebraic property characterizing the reduced-order matrix to be interpolated (for example, symmetric positive definiteness in the case of the reduced-order mass matrix). Then, when interpolating this reduced-order matrix at a queried but unsampled parameter point $\tilde{\boldsymbol{\mu}} \in \mathcal{D}$ – which contributes to the interpolation of the linear PROM represented by the tuple $\mathcal{A}_r(\tilde{\boldsymbol{\mu}})$ – this property can be preserved by performing the interpolation on the identified matrix manifold \mathcal{M} . Such an interpolation is described below for a generic reduced-order matrix of the tuple $\mathcal{A}_r^*(\boldsymbol{\mu})$, and the manifold of interest \mathcal{M} on which it lies.

Let $\mathbf{X} \in \mathcal{M}$ denote a reference point on the manifold \mathcal{M} , which in this case is differentiable. Let also $\{\mathbf{Y}^j \in \mathcal{M}\}_{j=1}^{N_{\mathcal{DB}}}$ denote a set of additional points on the manifold located in a neighborhood of \mathbf{X} , $\mathcal{N}(\mathbf{X})$. Here and throughout the remainder of this section, the superscript j refers to $\boldsymbol{\mu}^j \in \mathcal{D} \subset \mathbb{R}^{N_D}$. Then, $\boldsymbol{\Gamma}^j = \text{Log}_{\mathbf{X}}(\mathbf{Y}^j)$ is a mapping of $\mathbf{Y}^j \in \mathcal{M}$ onto the tangent space to \mathcal{M} at \mathbf{X} , $\mathcal{T}_{\mathbf{X}}\mathcal{M}$. The neighborhood $\mathcal{N}(\mathbf{X})$ is a subset of \mathcal{M} defined such that the equation $\text{Exp}_{\mathbf{X}}(\boldsymbol{\Gamma}^j) = \mathbf{Y}^j$ has a unique solution satisfying $\text{Log}_{\mathbf{X}}(\mathbf{Y}^j) = \boldsymbol{\Gamma}^j$, where $\text{Exp}_{\mathbf{X}}(\boldsymbol{\Gamma})$ is known as the exponential map: it maps elements of the tangent space $\mathcal{T}_{\mathbf{X}}\mathcal{M}$ back onto the manifold \mathcal{M} . In other words, the scope of $\mathcal{N}(\mathbf{X})$ is such that the inverse operation

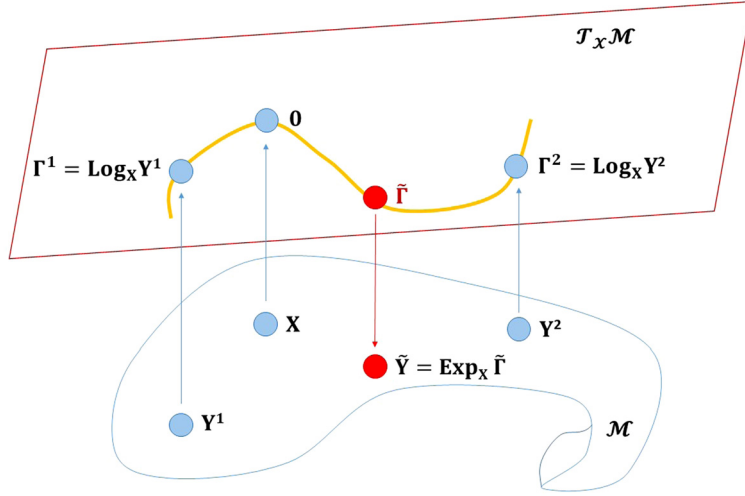


Fig. 2. Interpolation on a matrix manifold: tangent space at a reference point and logarithm and exponential maps.

Table 1

Logarithm and exponential maps for some well-known matrix manifolds.

Manifold	$\mathcal{R}^{M \times N}$	Nonsingular matrices	SPD matrices
$\text{Log}_X(\mathbf{Y})$	$\mathbf{Y} - \mathbf{X}$	$\log(\mathbf{Y}\mathbf{X}^{-1})$	$\log(\mathbf{X}^{-1/2}\mathbf{Y}\mathbf{X}^{-1/2})$
$\text{Exp}_X(\Gamma)$	$\mathbf{X} + \Gamma$	$\exp(\Gamma)\mathbf{X}$	$\mathbf{X}^{1/2}\exp(\Gamma)\mathbf{X}^{1/2}$

Algorithm 3 Interpolation on a matrix manifold \mathcal{M} .

Input: $N_{\mathcal{DB}}$ matrices $\mathbf{Y}^1, \dots, \mathbf{Y}^{N_{\mathcal{DB}}}$ belonging to a matrix manifold \mathcal{M} ; queried but unsampled parameter point $\tilde{\boldsymbol{\mu}} \in \mathcal{D} \subset \mathbb{R}^{N_{\mathcal{D}}}$

Output: Interpolated matrix $\tilde{\mathbf{Y}} \in \mathcal{M}$

- 1: Select a reference point on \mathcal{M} : for example, $\mathbf{X} = \mathbf{Y}^1$
 - 2: **for** $j = 1, \dots, N_{\mathcal{DB}}$ **do**
 - 3: Compute $\Gamma^j = \text{Log}_X(\mathbf{Y}^j)$
 - 4: **end for**
 - 5: Compute each entry of $\tilde{\Gamma}$ by interpolating the corresponding entries of Γ^j , $j = 1, \dots, N_{\mathcal{DB}}$
 - 6: Compute $\tilde{\mathbf{Y}} = \mathbf{Y}(\tilde{\boldsymbol{\mu}}) = \text{Exp}_X(\tilde{\Gamma})$
-

of the logarithm map has a well-defined unique output. The tangent space $\mathcal{T}_X \mathcal{M}$ being a vector space, any multi-variate interpolation approach ($N_{\mathcal{D}} \geq 1$) that is valid in a vector space can be used to interpolate the quantities Γ^j . Let $\tilde{\Gamma}$ denote the outcome of such an interpolation in $\mathcal{T}_X \mathcal{M}$. Next, $\tilde{\Gamma}$ is mapped back onto \mathcal{M} via the exponential map to obtain the final outcome of the interpolation process, $\tilde{\mathbf{Y}}(\tilde{\boldsymbol{\mu}}) = \text{Exp}_X(\tilde{\Gamma})$. This process is illustrated in Fig. 2.

Table 1 gives the expressions of the logarithm and exponential maps for various matrix manifolds of practical importance. The reader can observe that these expressions can be evaluated in real time. Algorithm 3 summarizes the procedure of interpolation on a matrix manifold described above. Again, the reader can observe that this algorithm can be processed in real time. It can be furthermore encapsulated in the following expression

$$\tilde{\mathbf{Y}} = \mathbf{Y}(\tilde{\boldsymbol{\mu}}) = \text{Exp}_X \left[\mathcal{I} \left(\tilde{\boldsymbol{\mu}}; \left\{ \text{Log}_X(\mathbf{Y}^j) \right\}_{j=1}^{N_{\mathcal{DB}}} \right) \right] \quad (17)$$

where $\tilde{\boldsymbol{\mu}} \in \mathcal{D}$ denotes as before a queried but unsampled point of the parameter space \mathcal{D} , \mathcal{I} is a multi-variate interpolation operator that intervenes in the tangent space $\mathcal{T}_X \mathcal{M}$, and $\tilde{\mathbf{Y}}$ is the interpolated reduced-order matrix of interest. (For further details on the theory and practice of interpolation on a matrix manifold in the context of linear PROMs, the reader is referred to [14,20].)

5.2. Interpolation of reduced-order model sensitivities

The solution of the reduced-order optimization problem (11) by a gradient-based algorithm requires the computation of the sensitivities of $\mathbf{A}_r^*(\boldsymbol{\mu})$ and $\mathbf{b}_r^*(\boldsymbol{\mu})$ with respect to the design parameters μ_i , $i = 1, \dots, N_{\mathcal{D}}$. Two approaches can be considered for computing these sensitivities in real time, at a queried but unsampled parameter point $\tilde{\boldsymbol{\mu}} \in \mathcal{D} \subset \mathbb{R}^{N_{\mathcal{D}}}$:

- *Finite differencing approach.* For each parameter component $\tilde{\mu}_i$, $i = 1, \dots, N_{\mathcal{D}}$, interpolate each set of pre-computed and transformed reduced-order matrices $\{\mathbf{A}_r^*(\boldsymbol{\mu}^j)\}_{j=1}^{N_{\mathcal{DB}}}$ and $\{\mathbf{b}_r^*(\boldsymbol{\mu}^j)\}_{j=1}^{N_{\mathcal{DB}}}$ on an appropriate matrix manifold \mathcal{M} , at the

Table 2

Shape sensitivities of the exponential maps associated with some well-known matrix manifolds.

Manifold	$\mathcal{R}^{M \times N}$	Nonsingular matrices	SPD matrices
$\frac{\partial \text{Exp}_{\mathbf{X}}(\Gamma(\boldsymbol{\mu}))}{\partial \mu_i}$	$\frac{\partial \Gamma(\boldsymbol{\mu})}{\partial \mu_i}$	$\frac{\partial \exp(\Gamma(\boldsymbol{\mu}))}{\partial \mu_i} \mathbf{X}$	$\mathbf{X}^{1/2} \frac{\partial \exp(\Gamma(\boldsymbol{\mu}))}{\partial \mu_i} \mathbf{X}^{1/2}$

parameter points $\tilde{\boldsymbol{\mu}} + \boldsymbol{\varepsilon}_i$ and/or $\tilde{\boldsymbol{\mu}} - \boldsymbol{\varepsilon}_i$, where $\boldsymbol{\varepsilon}_i \in \mathbb{R}^{N_{\mathcal{D}}}$ is the vector whose entries are zero except the i -th entry which satisfies $\varepsilon_i \neq 0$ and $|\varepsilon_i| \ll |\tilde{\mu}_i|$. If the aforementioned interpolations are performed at the point $\tilde{\boldsymbol{\mu}} + \boldsymbol{\varepsilon}_i$ or $\tilde{\boldsymbol{\mu}} - \boldsymbol{\varepsilon}_i$ only, interpolate also each aforementioned set of reduced-order matrices on \mathcal{M} , at the parameter point $\tilde{\boldsymbol{\mu}}$. Then, approximate $\frac{\partial \mathbf{A}_r}{\partial \mu_i}(\tilde{\boldsymbol{\mu}})$ and $\frac{\partial \mathbf{b}_r}{\partial \mu_i}(\tilde{\boldsymbol{\mu}})$ using finite differencing.

- *Analytical approach.* For each parameter component $\tilde{\mu}_i$, $i = 1, \dots, N_{\mathcal{D}}$, interpolate each set of pre-computed and transformed reduced-order sensitivity matrices $\left\{ \frac{\partial \mathbf{A}_r^*}{\partial \mu_i}(\boldsymbol{\mu}^j) \right\}_{j=1}^{N_{\mathcal{DB}}}$ and $\left\{ \frac{\partial \mathbf{b}_r^*}{\partial \mu_i}(\boldsymbol{\mu}^j) \right\}_{j=1}^{N_{\mathcal{DB}}}$ on an appropriate matrix manifold \mathcal{M} , at the parameter point $\tilde{\boldsymbol{\mu}}$.

The first approach outlined above is known to lack robustness with respect to the choice of the small quantity ε_i . At each queried but unsampled parameter point $\tilde{\boldsymbol{\mu}} \in \mathcal{D} \subset \mathbb{R}^{N_{\mathcal{D}}}$, it incurs $2(1 + N_{\mathcal{D}})$ interpolations if first-order forward or backward differencing is adopted, and $2(2N_{\mathcal{D}})$ interpolations if second-order central differencing is used instead. The second approach mentioned above is independent of ε_i : it requires $2N_{\mathcal{D}}$ interpolations. Hence, it is the preferred approach. However, it requires the development of an appropriate algorithm for interpolating each set $\left\{ \frac{\partial \mathbf{A}_r^*}{\partial \mu_i}(\boldsymbol{\mu}^j) \right\}_{j=1}^{N_{\mathcal{DB}}}$ and $\left\{ \frac{\partial \mathbf{b}_r^*}{\partial \mu_i}(\boldsymbol{\mu}^j) \right\}_{j=1}^{N_{\mathcal{DB}}}$ on

an appropriate matrix manifold, at $\tilde{\boldsymbol{\mu}}$, in order to approximate $\frac{\partial \mathbf{A}_r}{\partial \mu_i}(\tilde{\boldsymbol{\mu}})$ and $\frac{\partial \mathbf{b}_r}{\partial \mu_i}(\tilde{\boldsymbol{\mu}})$, respectively (for the same reason explained above for the interpolation of the other components of the tuple $\mathcal{A}_r^*(\boldsymbol{\mu})$). Such an algorithm is proposed below, using the same notation as in Section 5.1.2.

First, note that each quantity $\text{Log}_{\mathbf{X}}(\mathbf{Y}^j)$, where $\mathbf{Y}(\boldsymbol{\mu}^j) \in \mathcal{M}$ may represent here the $N_{w_{Lr}} \times N_{w_{Lr}}$ reduced-order matrix $\mathbf{A}_r^*(\boldsymbol{\mu}^j)$ or the $N_{w_{Lr}} \times 1$ reduced-order matrix $\mathbf{b}_r^*(\boldsymbol{\mu}^j)$ and $j = 1, \dots, N_{\mathcal{DB}}$, depends on $\boldsymbol{\mu}^j \in \mathcal{D}$ but does not depend on the arbitrary parameter point $\boldsymbol{\mu}$. Hence, the partial derivative (17) with respect to an arbitrary parameter component μ_i can be computed at any arbitrary parameter point $\boldsymbol{\mu} \in \mathcal{D} \subset \mathbb{R}^{N_{\mathcal{D}}}$ as

$$\frac{\partial \mathbf{Y}}{\partial \mu_i}(\boldsymbol{\mu}) = \frac{\partial \text{Exp}_{\mathbf{X}} \mathcal{I}}{\partial \mathcal{I}} \frac{\partial \mathcal{I}}{\partial \mu_i} \left(\boldsymbol{\mu}; \{\text{Log}_{\mathbf{X}}(\mathbf{Y}^j)\}_{j=1}^{N_{\mathcal{DB}}} \right), \quad i = 1, \dots, N_{\mathcal{D}} \quad (18)$$

The objective of the algorithm presented here is to approximate the above partial derivative at a queried but unsampled parameter point $\tilde{\boldsymbol{\mu}}$ using interpolation on a matrix manifold. To this end, Table 2 provides, for several practical and relevant matrix manifolds, the partial derivatives of the exponential map $\frac{\partial \text{Exp}_{\mathbf{X}} \Gamma(\boldsymbol{\mu})}{\partial \mu_i}$, $i = 1, \dots, N_{\mathcal{D}}$. Note that $\frac{\partial \Gamma(\boldsymbol{\mu})}{\partial \mu_i} \in \mathcal{T}_{\mathbf{X}} \mathcal{M}$ can be computed analytically, but the evaluation of the partial derivative $\frac{\partial \exp(\Gamma(\boldsymbol{\mu}))}{\partial \mu_i}$ requires special attention. In order to compute this derivative, consider the matrix

$$\mathbf{B}_i(\Gamma(\boldsymbol{\mu})) = \begin{bmatrix} \Gamma(\boldsymbol{\mu}) & \frac{\partial \Gamma(\boldsymbol{\mu})}{\partial \mu_i} \\ \mathbf{0} & \Gamma(\boldsymbol{\mu}) \end{bmatrix} \in \mathbb{R}^{2N_{w_{Lr}} \times 2N_{w_{Lr}}}, \quad i = 1, \dots, N_{\mathcal{D}} \quad (19)$$

In [31], it is shown that the exponential of such a matrix is the matrix

$$\exp(\mathbf{B}_i(\Gamma(\boldsymbol{\mu}))) = \begin{bmatrix} \exp(\Gamma(\boldsymbol{\mu})) & \frac{\partial \exp(\Gamma(\boldsymbol{\mu}))}{\partial \mu_i} \\ \mathbf{0} & \exp(\Gamma(\boldsymbol{\mu})) \end{bmatrix} \in \mathbb{R}^{2N_{w_{Lr}} \times 2N_{w_{Lr}}}, \quad i = 1, \dots, N_{\mathcal{D}}, \quad i = 1, \dots, N_{\mathcal{D}} \quad (20)$$

From (19) and (20), it follows that $\frac{\partial \exp(\Gamma(\boldsymbol{\mu}))}{\partial \mu_i}$ can be simply obtained by computing the exponential of the matrix $\mathbf{B}_i(\Gamma(\boldsymbol{\mu}))$, then extracting from $\exp(\mathbf{B}_i(\Gamma(\boldsymbol{\mu})))$ its (1,2)-block. This computation can be performed in real time as $\mathbf{B}_i \in \mathbb{R}^{2N_{w_{Lr}} \times 2N_{w_{Lr}}}$ is a reduced-order matrix.

Therefore, the proposed real-time algorithm for interpolating the quantity $\frac{d\mathbf{Y}}{d\mu_i}(\boldsymbol{\mu})$ (18) on a matrix manifold, at a queried but unsampled parameter point $\tilde{\boldsymbol{\mu}}$, where $\mathbf{Y}(\boldsymbol{\mu}) \in \mathcal{M}$ may represent the $N_{w_{Lr}} \times N_{w_{Lr}}$ reduced-order matrix $\mathbf{A}_r^*(\boldsymbol{\mu})$

or the $N_{w_{lr}} \times 1$ reduced-order matrix $\mathbf{b}_r^*(\boldsymbol{\mu})$ (and therefore $\frac{d\mathbf{Y}}{d\mu_i}(\boldsymbol{\mu})$ may represent the reduced-order sensitivity matrix $\frac{d\mathbf{A}_r^*}{d\mu_i}(\boldsymbol{\mu})$ or $\frac{d\mathbf{b}_r^*}{d\mu_i}(\boldsymbol{\mu})$) can be described as follows:

- Let \mathcal{M} be the matrix manifold to which $\mathbf{Y}(\boldsymbol{\mu})$ belongs.
- Select a reference point on \mathcal{M} : for example, $\mathbf{X} = \mathbf{Y}^1$.
- Compute $\boldsymbol{\Gamma}^j = \text{Log}_{\mathbf{Y}^1}(\mathbf{Y}^j) \in \mathcal{T}_{\mathbf{Y}^1}\mathcal{M}$, $j = 1, \dots, N_{\mathcal{DB}}$.
- Compute also analytically $\frac{\partial \boldsymbol{\Gamma}(\boldsymbol{\mu}^j)}{\partial \mu_i}$, $j = 1, \dots, N_{\mathcal{DB}}$.
- Interpolate $\{\boldsymbol{\Gamma}^j\}_{j=1}^{N_{\mathcal{DB}}}$ in the tangent space $\mathcal{T}_{\mathbf{Y}^1}\mathcal{M}$, at the point $\tilde{\boldsymbol{\mu}} \in \mathcal{D}$, to obtain an approximation of $\boldsymbol{\Gamma}(\tilde{\boldsymbol{\mu}})$.
- For each parameter component $\tilde{\mu}_i$:
 - Form the reduced-order matrix $\mathbf{B}_i(\boldsymbol{\Gamma}(\tilde{\boldsymbol{\mu}}))$ (19).
 - Compute the reduced-order matrix $\exp(\mathbf{B}_i(\boldsymbol{\Gamma}(\tilde{\boldsymbol{\mu}})))$ (20).
 - Extract the (1, 2)-block of $\exp(\mathbf{B}_i(\boldsymbol{\Gamma}(\tilde{\boldsymbol{\mu}})))$ and identify $\frac{d\mathbf{Y}}{d\mu_i}(\tilde{\boldsymbol{\mu}})$ with this block.

Note that the above algorithm interpolates the values of the reduced-order sensitivity matrices $\frac{\partial \mathbf{Y}}{\partial \mu_i}(\tilde{\boldsymbol{\mu}})$, $i = 1, \dots, N_{\mathcal{D}}$, on the same matrix manifold where Algorithm 3 interpolates the reduced-order matrix $\mathbf{Y}(\tilde{\boldsymbol{\mu}})$. Furthermore, this algorithm shares with Algorithm 3 most of its steps. Hence, both algorithms can be combined to interpolate on the same matrix manifold $\mathbf{Y}(\tilde{\boldsymbol{\mu}})$ and its shape sensitivities $\frac{\partial \mathbf{Y}}{\partial \mu_i}(\tilde{\boldsymbol{\mu}})$, $i = 1, \dots, N_{\mathcal{D}}$.

5.3. Tessellation-free approach based on adaptive least-squares radial basis functions

It remains to specify the interpolation method to be applied in the tangent space $\mathcal{T}_{\mathbf{Y}^1}\mathcal{M}$. The most important requirement for such a method is to be practical and perform well in the presence of high-dimensional parameter spaces and scattered data. Tessellation methods or related methods based on computational meshes constitute good candidates for this purpose when $N_{\mathcal{D}} \leq 3$ (for example, see [19]). However, they are impractical if not impossible when $N_{\mathcal{D}} > 3$. High-dimensional parameter spaces call for mesh-free interpolation methods. An approach that is both popular and effective in this context is one that is based on radially symmetric basis functions that usually take the form of univariate functions of an Euclidean norm. Such an approach turns a multi-variate interpolation problem into one that is virtually one-dimensional, thereby simplifying the interpolation problem and guaranteeing some interesting mathematical properties.

In the context of the parameter space \mathcal{D} of dimension $N_{\mathcal{D}}$, a Radial Basis Function (RBF) is a real-valued function ϕ whose value at a parameter point $\boldsymbol{\mu}$ depends only on the distance from this point to some other point \mathbf{c} called the center, so that

$$\phi(\boldsymbol{\mu}) = \phi(\|\boldsymbol{\mu} - \mathbf{c}\|_2)$$

Popular RBFs include the Gaussian RBF characterized by $\phi(\boldsymbol{\mu}) = e^{-(\epsilon\|\boldsymbol{\mu} - \mathbf{c}\|)^2}$, and the inverse quadratic RBF characterized by $\phi(\boldsymbol{\mu}) = (1 + (\epsilon\|\boldsymbol{\mu} - \mathbf{c}\|)^2)^{-1}$, where ϵ is a calibration scalar. In this work, linear combinations of these RBFs are proposed to approximate a Qol such as a component of the tuple of low-dimensional operators $\mathcal{A}_r(\tilde{\boldsymbol{\mu}})$ (10). The center points \mathbf{c} are chosen as usual to be the interpolation data, and therefore the interpolant can be written as

$$P(\boldsymbol{\mu}) = \sum_{j=1}^{N_{\lambda}} \lambda_j \phi(\|\boldsymbol{\mu} - \boldsymbol{\mu}^j\|_2), \quad \boldsymbol{\mu} \in \mathcal{D} \subset \mathbb{R}^{N_{\mathcal{D}}}$$

where the number of basis functions N_{λ} is in general less than or equal to the number sampled parameter points – that is, $N_{\lambda} \leq N_{\mathcal{DB}}$. The coefficients λ_j , $j = 1, \dots, N_{\lambda}$ are typically determined by solving the linear least-squares problem

$$\min_{\boldsymbol{\lambda} \in \mathbb{R}^{N_{\lambda}}} \|\mathbf{f} - B\boldsymbol{\lambda}\|_2^2$$

where $B \in \mathbb{R}^{N \times N_{\lambda}}$ is the interpolation matrix defined by $B_{ij} = \phi(\|\boldsymbol{\mu}^i - \boldsymbol{\mu}^j\|)$, $\boldsymbol{\lambda} = (\lambda_1, \dots, \lambda_{N_{\lambda}})$, $\mathbf{f} = (f_1, \dots, f_N)$ is the vector storing the values of each component of the tuple $\mathcal{A}_r(\boldsymbol{\mu}^j)$ to be interpolated, $j = 1, \dots, N_{\mathcal{DB}}$, and therefore $N = N_{\mathcal{DB}}N_{w_{lr}}M_{w_{lr}}$ (see Section 5.1.2).

As already stated, the center points of the RBFs are typically chosen to be the entire set of sampled parameter points: this defines the straightforward implementation of an RBF-based interpolation method. For high-dimensional parameter spaces however, or when higher accuracy is desired, this implementation may incur an unacceptable computational cost, particularly in the context of real-time computations. Here, this cost is alleviated by focusing on the *adaptive least-squares*

implementation [32] of an RBF-based interpolation, where a smaller, more economical interpolant is iteratively constructed. In this adaptive implementation, an initial interpolant is constructed using a small subset of points in \mathcal{D} as center points. Then, the same greedy procedure used to construct the database \mathcal{DB} (see Section 4) but equipped with the interpolation error as an error indicator, is used to select additional center points for constructing the RBF interpolant. Though this process incurs the construction of several small interpolants, interpolation accuracy can be achieved using a substantially smaller number of center points than $N_{\mathcal{DB}}$, which makes this implementation more efficient than its straightforward counterpart.

6. Applications

In this section, the proposed fast approach for gradient-based constrained optimization based on a database of linear PROMs is illustrated with two examples pertaining to the aeroelastic design optimization of the same wing, under the same linearized flutter constraint. The only difference between the two examples is the dimension of the parameter space $N_{\mathcal{D}}$ – that is, the number of design parameters. All HDM-based computations are performed on 32 CPU cores (or simply, CPUs) of a Linux cluster. All HDM-based computations incurred by the offline construction of a PROM database are also performed on 32 CPUs. However, all PROM-based online computations are performed on a single CPU of the same cluster. In all cases, a CPU timing is computed and reported as the product of the corresponding wall-clock timing and the number of CPUs used to perform the corresponding computation.

Throughout the remainder of this section, the superscript $*$ is dropped from (16) and any related quantity for the sake of simplicity. Instead, it is assumed that all information pre-computed offline is rotated as described in Section 5.1.1 before it is stored in a database \mathcal{DB} , in order to enforce mathematical consistency.

6.1. Aeroelastic design optimization of a wing subject to a linearized flutter constraint

Flutter is one of the most important considerations in aircraft design. It is an aeroelastic instability that can be predicted by linearizing the governing fluid-structure equations of dynamic equilibrium around the steady-state equilibrium associated with the specified flight conditions, and solving an associated eigenvalue problem. When the fluid subsystem is modeled using CFD – for example, to improve fidelity in the transonic regime where shocks play an important role – the prediction process becomes very CPU intensive. For this reason, the aeroelastic design optimization of an aircraft wing under a linearized flutter constraint is a good candidate problem for the computational technology described in this paper.

6.1.1. Problem formulation and specification

To this end, the following constrained, aeroelastic, design optimization problem – formulated here for any generic aircraft wing – is considered:

$$\begin{aligned} & \underset{\boldsymbol{\mu} \in \mathcal{D}}{\text{maximize}} && \frac{L(\boldsymbol{\mu})}{D(\boldsymbol{\mu})} \\ & \text{subject to} && W(\boldsymbol{\mu}) \leq W_{\text{upper}} \\ & && \sigma_{\text{VM}}(\boldsymbol{\mu}) \leq \sigma_{\text{upper}} \\ & && \boldsymbol{\mu}_{\text{lower}} \leq \boldsymbol{\mu} \leq \boldsymbol{\mu}_{\text{upper}} \\ & && \boldsymbol{\zeta}_{\text{lower}} \leq \boldsymbol{\zeta}(\boldsymbol{\mu}) \end{aligned} \tag{21}$$

where: $L(\boldsymbol{\mu})$ and $D(\boldsymbol{\mu})$ denote the lift and drag, respectively; $W(\boldsymbol{\mu})$ denotes the weight of the wing, which is required to remain below an upper bound W_{upper} ; $\sigma_{\text{VM}}(\boldsymbol{\mu})$ is the von Mises stress in the wing at the steady-state equilibrium point around which linearization is performed, and is constrained not to exceed the maximum allowable yield stress σ_{upper} ; $\boldsymbol{\mu}_{\text{lower}}$ and $\boldsymbol{\mu}_{\text{upper}}$ are lower and upper bounds, respectively, that define bound constraints on $\boldsymbol{\mu}$ in order to avoid unrealistic designs; $\boldsymbol{\zeta}(\boldsymbol{\mu})$ is the vector of damping ratios, where each damping ratio is required to remain above a lower bound $\zeta_{\text{lower}} > 0$ in order to avoid flutter; and $\boldsymbol{\zeta}_{\text{lower}} = \zeta_{\text{lower}} \mathbf{1}$, where $\mathbf{1}$ denotes the unit vector of same dimension as $\boldsymbol{\zeta}$. Here, the flutter constraint is evaluated using a PROM database.

Two types of design parameters are considered and therefore $\boldsymbol{\mu}$ is written as $\boldsymbol{\mu} = (\boldsymbol{\mu}_s, \boldsymbol{\mu}_m)$: aerodynamic shape parameters represented by the subvector $\boldsymbol{\mu}_s \in \mathbb{R}^{N_{\mathcal{D}_s}}$, and structural parameters represented by $\boldsymbol{\mu}_m \in \mathbb{R}^{N_{\mathcal{D}_m}}$. The aerodynamic shape parameters $\boldsymbol{\mu}_s$ affect both the shape of the structure, and the fluid flow through the transmission conditions enforced at the fluid-structure interface. The parameter subvector $\boldsymbol{\mu}_m$ contains the material properties of the structure as well as the internal shape parameters associated with its “dry” elements – that is, the elements of the structure that are not in contact with the external flow.

The constraints of problem (21) can be grouped into two sets, based on how they are evaluated:

- Static aeroelastic constraints that are computed using fluid and structural HDMs that are based on the three-field formulation of fluid-structure interaction problems [33,34]. These constraints govern the weight, lift to drag ratio, and the von Mises stress. They are obtained from the solution of the governing nonlinear PDEs, which are not reduced in this paper.

Table 3
Geometrical and material properties of the ARW-2.

Parameter	Type/Value
Geometry (inches)	
Wingspan	104.9
Root	40.2
Tip	12.5
Material properties	
Skin (except flaps)	Composite materials
Stiffeners	Aluminum
E (psi)	1.03×10^7
ρ (lb·s ² /in ⁴)	2.6×10^{-4}
ν	0.32

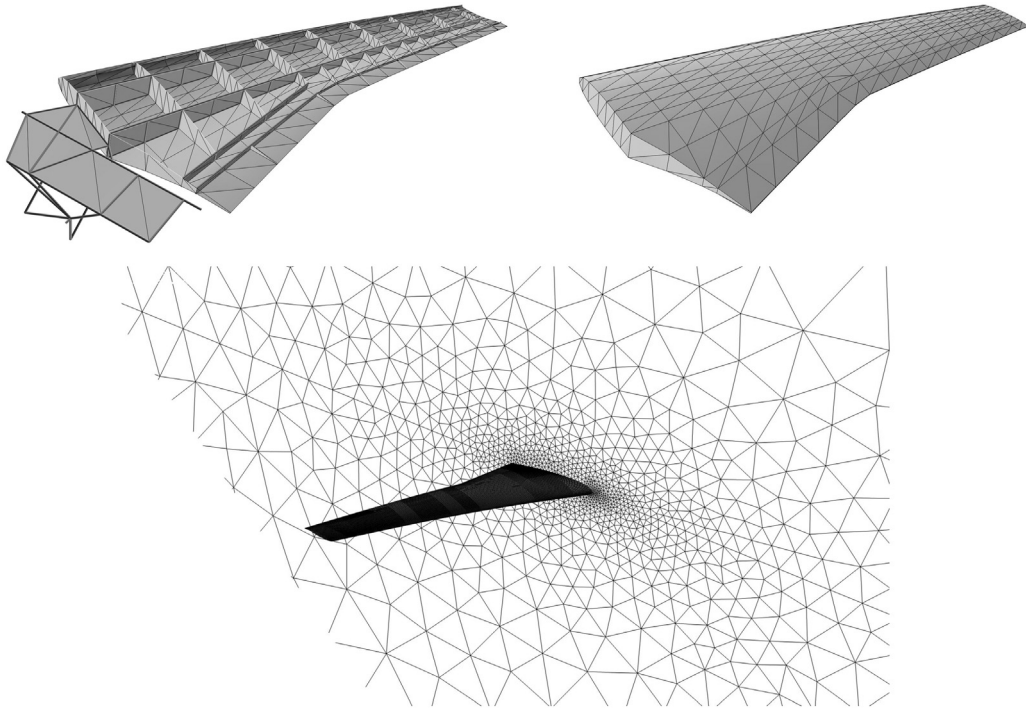


Fig. 3. Computational models for the ARW-2: detailed FE structural model with upper skin hidden (top, left); skin surface of this model (top, right); and CFD mesh (bottom).

- Dynamic, aeroelastic, flutter constraints, which are computed by linearizing the aforementioned HDMs around an equilibrium state [35,36]. Here, the proposed PROM database strategy is exploited to alleviate the large computational cost associated with the computation of these constraints. Precisely, the computation of the vector of damping ratios ζ and its sensitivities $\frac{d\zeta}{d\mu}$ is accelerated using the PROM interpolation approach described in Section 5.

Specifically, the constrained design optimization problem (21) is considered here for a configuration of the Aeroelastic Research Wing (ARW-2) [37]. The physical dimensions and material properties of this wing, which has been built by NASA, are reported in Table 3. A detailed, Finite Element (FE), structural model of this wing is shown in Fig. 3: it includes, among others, FE representations of the spars, ribs, hinges, and control surfaces of this wing. The model has a total of 2,556 degrees of freedom (dofs). Air is assumed to behave as a perfect gas, and its flow around this wing is assumed to be inviscid. Hence, the fluid flow is modeled here by the Euler equations. The computational fluid domain is discretized by a three-dimensional, unstructured, CFD mesh (Fig. 3) with 63,484 nodes. The focus is set on a cruise condition defined by: the altitude $h = 4,000$ ft where the atmospheric density is $\rho_\infty = 1.0193 \times 10^{-7}$ lb·s²/in⁴ and the atmospheric pressure is $P_\infty = 12.7$ lb/in²; the free-stream Mach number $M_\infty = 0.8$; and the zero angle of attack.

For this wing, the upper bound for the von Mises stress constraint, σ_{upper} , is set to 2.5×10^4 psi, and that for the weight, W_{upper} , is set to 400 lbs. Several values of the lower bound for the damping ratios, ζ_{lower} , are considered in the interval $[4.1 \times 10^{-4}, 4.5 \times 10^{-4}]$ in order to study how the optimal aeroelastic design varies with this flutter threshold.

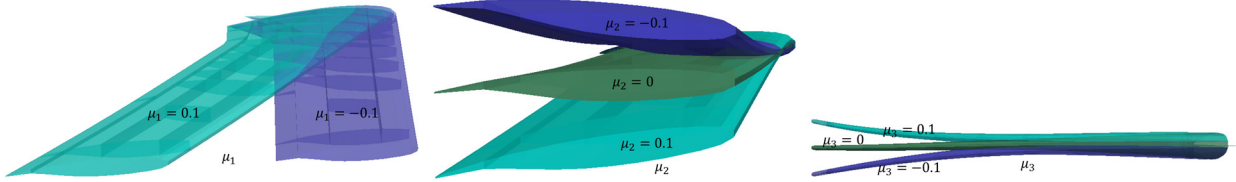


Fig. 4. Aerodynamic shape design parameters.

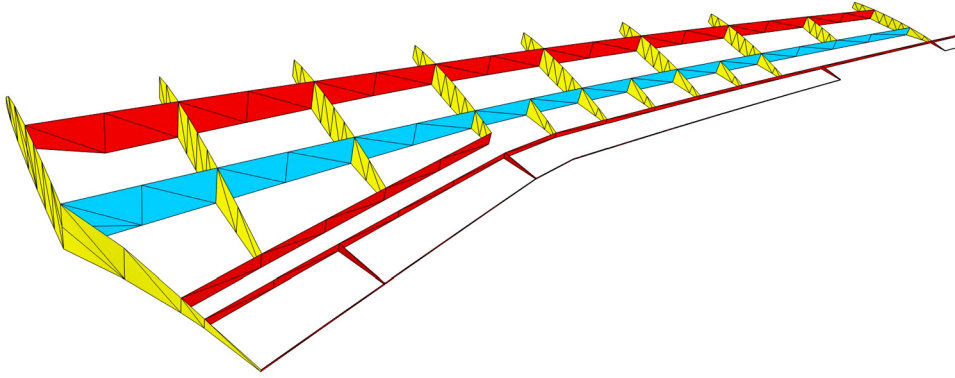


Fig. 5. Organization of the ARW-2 stiffeners in three different groups.

Three shape and three structural parameters are considered – that is, $\boldsymbol{\mu}_s \in \mathbb{R}^3$ and $\boldsymbol{\mu}_m \in \mathbb{R}^3$. The shape parameters are the back sweep angle μ_s^1 , the twist angle μ_s^2 , and the dihedral angle μ_s^3 . These parameters, which are graphically depicted in Fig. 4, are constrained to have a magnitude less than 0.1 radian.

For the purpose of optimization, the stiffeners of the ARW-2 are organized in three different groups: the outer spars, the central spar, and the ribs (see Fig. 5). Each of these groups contains elements with different thicknesses, but is attributed a single, collective thickness increment as an optimization parameter. Hence, the structural parameters stored in $\boldsymbol{\mu}_m$ are the three thickness increments for the outer spars (μ_m^1), the central spar (μ_m^2), and the ribs (μ_m^3). The magnitude of each of these increments is constrained to be less than 0.01 in – which corresponds on average to 10% of the thickness of a stiffening element of the considered initial configuration of the ARW-2.

Specifically, two different instances of the design optimization problem (21) are considered in the remainder of this section for the ARW-2:

1. A first instance where only the aerodynamic shape parameters are activated as optimization parameters – that is, $\boldsymbol{\mu} = \boldsymbol{\mu}_s$.
2. A second instance where both the aerodynamic shape and structural parameters are activated as optimization parameters – that is, $\boldsymbol{\mu} = (\boldsymbol{\mu}_s, \boldsymbol{\mu}_m)$.

For each instance outlined above, the performance of the model reduction approach adopted here for computing the flutter constraints and its impact on the overall performance of the NAND solution algorithm is measured relative to the performance results obtained using the HDM-based solution strategies.

In order to keep this paper as self-contained as possible, the three-field formulation of fluid-structure interaction problems and its linearization, which are used here to formulate the flutter constraint $\boldsymbol{\zeta}(\boldsymbol{\mu}) \geq \boldsymbol{\zeta}_{\text{lower}}$ in (21), are next overviewed (for more details, the reader is referred to [33–36]).

6.1.2. Linearized CFD-based computational aeroelasticity and model order reduction

A high-fidelity, nonlinear, CFD-based fluid-structure (or aeroelastic) system can be described by a three-field Arbitrary Lagrangian Eulerian (ALE) formulation [33,34] that treats the moving CFD mesh as a pseudo-structural system [38]. After semi-discretization, linearization of the governing semi-discrete equations about a fluid-structure equilibrium state [35,36], and elimination from these equations of the CFD mesh position vector, this formulation leads to the following system of linear Ordinary Differential Equations (ODEs)

$$\mathbf{C}_v(\boldsymbol{\mu})\dot{\mathbf{w}} + \mathbf{H}(\boldsymbol{\mu})\mathbf{w} + \mathbf{R}(\boldsymbol{\mu})\dot{\mathbf{u}} + \mathbf{G}(\boldsymbol{\mu})\mathbf{u} = \mathbf{0} \quad (22)$$

$$\mathbf{M}(\boldsymbol{\mu})\ddot{\mathbf{u}} + \mathbf{D}(\boldsymbol{\mu})\dot{\mathbf{u}} + \mathbf{K}(\boldsymbol{\mu})\mathbf{u} = \mathbf{P}(\boldsymbol{\mu})\mathbf{w} \quad (23)$$

Algorithm 4 Construction of a fluid ROB for an aeroelastic system.

Input: Parameter vector $\boldsymbol{\mu} = (\boldsymbol{\mu}_s, \boldsymbol{\mu}_m) \in \mathbb{R}^{p_s+p_m}$, frequency sampling $\{\xi_j\}_{j=1}^{N_\xi}$, structural eigenmodes $\mathbf{X}(\boldsymbol{\mu}) \in \mathbb{R}^{N_s \times k_s}$, desired ROB dimension k_f

Output: ROB $\mathbf{V}(\boldsymbol{\mu})$

- 1: Compute $\mathbf{C}_v = \mathbf{C}_v(\boldsymbol{\mu})$, $\mathbf{H} = \mathbf{H}(\boldsymbol{\mu})$, $\mathbf{R} = \mathbf{R}(\boldsymbol{\mu})$, and $\mathbf{G} = \mathbf{G}(\boldsymbol{\mu})$
- 2: **for** $i = 1, \dots, k_s$ **do**
- 3: **for** $l = 1, \dots, N_\xi$ **do**
- 4: Solve the linear system $(j\xi_l \mathbf{C}_v + \mathbf{H})\mathbf{w}_{i,l} = -(j\xi_l \mathbf{R} + \mathbf{G})\mathbf{x}_i$, where \mathbf{x}_i is the i -th vector in \mathbf{X}
- 5: **end for**
- 6: **end for**
- 7: Construct the complex-valued snapshot matrix $\mathbf{W} = [\mathbf{w}_{1,1}, \dots, \mathbf{w}_{k_s, N_\xi}]$
- 8: Compute the SVD $\mathbf{C}_v^{-\frac{1}{2}} [\text{Re}(\mathbf{W}), \text{Im}(\mathbf{W})] = \mathbf{U}\mathbf{\Sigma}\mathbf{Z}^T$
- 9: Retain the first k_f left singular vectors $\mathbf{V}(\boldsymbol{\mu}) = \mathbf{C}_v^{-\frac{1}{2}} \mathbf{U}_{k_f}$

where: a dot denotes a derivative with respect to time t ; $\mathbf{C}_v \in \mathbb{R}^{N_f \times N_f}$ is the diagonal matrix of cell volumes in the CFD mesh, and N_f is the dimension of the semi-discretized fluid subsystem; $\mathbf{w}(t) \in \mathbb{R}^{N_f}$ denotes the conservative fluid state vector; $\mathbf{u}(t) \in \mathbb{R}^{N_s}$ is the vector of structural displacements of dimension N_s ; $\mathbf{H} \in \mathbb{R}^{N_f \times N_f}$ and $\mathbf{G} \in \mathbb{R}^{N_f \times N_x}$ denote the Jacobians of the numerical fluid fluxes with respect to \mathbf{w} and \mathbf{x} , respectively; $\mathbf{R} \in \mathbb{R}^{N_f \times N_x}$ is the Jacobian of the numerical fluid fluxes with respect to the fluid mesh velocity; $\mathbf{M} \in \mathbb{R}^{N_s \times N_s}$ is the mass matrix associated with the FE representation of the structural subsystem; $\mathbf{D} \in \mathbb{R}^{N_s \times N_s}$ and $\mathbf{K} \in \mathbb{R}^{N_s \times N_s}$ are the FE damping and stiffness matrices, respectively; and $\mathbf{P} \in \mathbb{R}^{N_s \times N_f}$ denotes the Jacobian of the aerodynamic forces with respect to \mathbf{w} and acts on the “wet” surface of the structure – that is, the fluid-structure interface.

First, the dimensionality of the structural subsystem is reduced using modal truncation. For this purpose, a ROB $\mathbf{X}(\boldsymbol{\mu}) \in \mathbb{R}^{N_s \times k_s}$, where \mathbf{X} depends on the parameter vector $\boldsymbol{\mu} = (\boldsymbol{\mu}_s, \boldsymbol{\mu}_m) \in \mathbb{R}^{p_s+p_m}$, is constructed using the first k_s modes of the structural subsystem. Then, the subspace approximation

$$\mathbf{u}(t) \approx \mathbf{X}(\boldsymbol{\mu})\mathbf{u}_r(t)$$

where $\mathbf{u}_r \in \mathbb{R}^{k_s}$ denotes the vector of k_s generalized coordinates associated with the modal approximation, is performed. Finally, Equation (23) is reduced by Galerkin projection onto $\mathbf{X}(\boldsymbol{\mu})$ to obtain

$$\ddot{\mathbf{u}}_r + \mathbf{D}_r(\boldsymbol{\mu})\dot{\mathbf{u}}_r + \boldsymbol{\Omega}_r^2(\boldsymbol{\mu})\mathbf{u}_r = \mathbf{X}(\boldsymbol{\mu})^T \mathbf{P}(\boldsymbol{\mu})\mathbf{w}$$

where $\mathbf{D}_r(\boldsymbol{\mu}) = \mathbf{X}(\boldsymbol{\mu})^T \mathbf{D}(\boldsymbol{\mu})\mathbf{X}(\boldsymbol{\mu}) \in \mathbb{R}^{k_s \times k_s}$ and $\boldsymbol{\Omega}_r^2 \in \mathbb{R}^{k_s \times k_s}$ is the diagonal matrix of eigenvalues associated with the k_s eigenmodes.

The dimensionality of the fluid subsystem is reduced using the POD method in the frequency domain [39]. Specifically, a ROB $\mathbf{V}(\boldsymbol{\mu}) \in \mathbb{R}^{N_f \times k_f}$ is constructed using Algorithm 4. Note that by construction, this global ROB satisfies the orthogonality condition $\mathbf{V}(\boldsymbol{\mu})^T \mathbf{C}_v(\boldsymbol{\mu})\mathbf{V}(\boldsymbol{\mu}) = \mathbf{I}_{k_f}$. Next, the state vector \mathbf{w} is approximated as

$$\mathbf{w}(t) \approx \mathbf{V}(\boldsymbol{\mu})\mathbf{w}_r(t)$$

where \mathbf{w}_r denotes the vector of k_f generalized coordinates associated with $\mathbf{V}(\boldsymbol{\mu})$, and Galerkin projection is performed to transform (22) into the set of reduced-order coupled linear ODEs

$$\dot{\mathbf{w}}_r + \mathbf{H}_r(\boldsymbol{\mu})\mathbf{w}_r + \mathbf{R}_r(\boldsymbol{\mu})\dot{\mathbf{u}}_r + \mathbf{G}_r(\boldsymbol{\mu})\mathbf{u}_r = \mathbf{0} \quad (24)$$

$$\ddot{\mathbf{u}}_r + \mathbf{D}_r(\boldsymbol{\mu})\dot{\mathbf{u}}_r + \boldsymbol{\Omega}_r^2(\boldsymbol{\mu})\mathbf{u}_r - \mathbf{P}_r(\boldsymbol{\mu})\mathbf{w}_r = \mathbf{0} \quad (25)$$

where $\mathbf{H}_r(\boldsymbol{\mu}) = \mathbf{V}(\boldsymbol{\mu})^T \mathbf{H}(\boldsymbol{\mu}_s)\mathbf{V}(\boldsymbol{\mu})$, $\mathbf{R}_r(\boldsymbol{\mu}) = \mathbf{V}(\boldsymbol{\mu})^T \mathbf{R}(\boldsymbol{\mu}_s)\mathbf{X}(\boldsymbol{\mu})$, and $\mathbf{P}_r(\boldsymbol{\mu}) = \mathbf{X}(\boldsymbol{\mu})^T \mathbf{P}\mathbf{V}(\boldsymbol{\mu})$. If needed, the stability of this resulting fluid-structure PROM can be enforced using the fast algorithm described in [40].

The coupled PROM (24)–(25) can be re-written in compact form as

$$\dot{\mathbf{q}}_r = \mathbf{N}_r(\boldsymbol{\mu})\mathbf{q}_r$$

where

$$\mathbf{q}_r(t) = \begin{bmatrix} \mathbf{w}_r(t) \\ \dot{\mathbf{u}}_r(t) \\ \mathbf{u}_r(t) \end{bmatrix} \in \mathbb{R}^{k_f+2k_s} \quad (26)$$

and

$$\mathbf{N}_r(\boldsymbol{\mu}) = \begin{bmatrix} -\mathbf{H}_r(\boldsymbol{\mu}) & -\mathbf{R}_r(\boldsymbol{\mu}) & -\mathbf{G}_r(\boldsymbol{\mu}) \\ \mathbf{P}_r(\boldsymbol{\mu}) & -\mathbf{D}_r(\boldsymbol{\mu}) & -\boldsymbol{\Omega}_r^2(\boldsymbol{\mu}) \\ \mathbf{0} & \mathbf{I}_{k_s} & \mathbf{0} \end{bmatrix} \in \mathbb{R}^{(k_f+2k_s) \times (k_f+2k_s)} \quad (27)$$

Table 4
Performance comparison of three different instances of the greedy procedure.

Greedy procedure	Standard	Random	Saturation
N_{Ξ}	125	125	125
N_{Π}	-	20	20
Number of sampled HDMs	16	15	15
Maximum Relative Indicated Error (%)	4.9	3.9	4.6
Total wall-clock time (hrs)	60.9	12.1	8.4
Speed-up with respect to standard greedy procedure	1	5.1	7.3

For simplicity, and without any loss of generality, structural damping is neglected throughout the remainder of this section – that is $\mathbf{D}_r(\boldsymbol{\mu}) = \mathbf{0}$.

Given that $k_s \ll N_s$ and $k_f \ll N_f$, the reduced-order eigenvalue problem

$$\mathbf{N}_r(\boldsymbol{\mu})\mathbf{q}_r(\boldsymbol{\mu}) = \lambda(\boldsymbol{\mu})\mathbf{q}_r(\boldsymbol{\mu}) \quad (28)$$

can be solved in real-time, for any specified value of $\boldsymbol{\mu}$, to obtain the $2k_s + k_f$ eigenpairs $\{\lambda_j, \mathbf{q}_{rj}\}_{j=1}^{2k_s+k_f}$. For each computed eigenvalue λ_j , which can be complex-valued, the corresponding damping ratio ζ_j is defined as

$$\zeta_j = -\frac{\lambda_j^R}{\sqrt{(\lambda_j^R)^2 + (\lambda_j^I)^2}} \quad (29)$$

where λ_j^R and λ_j^I are the real and imaginary part of λ_j , respectively.

During the solution of the optimization problem (21), many values of the parameter vector $\boldsymbol{\mu}$ will be queried. For any value $\tilde{\boldsymbol{\mu}}$ for which the blocks of the matrix $\mathbf{N}_r(\boldsymbol{\mu})$ (27) are not available in the PROM database \mathcal{DB} constructed offline, the PROM interpolation method described in Section 5 is applied to construct $\mathbf{N}_r(\tilde{\boldsymbol{\mu}})$ and its sensitivities with respect to the optimization parameters block-by-block. For this purpose, Equation (28) can be identified with Equation (9) – or more specifically, the reduced-order matrix \mathbf{N}_r can be identified with the reduced-order matrix \mathbf{A}_r – even if (1) is a linear system of equations but (28) is a linear eigenvalue problem.

The solution of the optimization problem (21) also requires the sensitivity of the vector of damping ratios with respect to the optimization parameters, $\frac{\partial \zeta}{\partial \mu_i}$, $i = 1, \dots, N_{\mathcal{D}}$. Appendix 8 describes a real-time algorithm for computing online these sensitivities.

6.2. Offline database construction

Three different instances of the greedy procedure are applied offline to construct three different databases of coupled PROMs of the form (27) and of reduced dimensions $k_s = 6$ and $k_f = 100$ and compare their performances. In all three instances: the set of candidate parameter points Ξ is constructed using FF design and five points along each dimension of the parameter space, which leads to $N_{\Xi} = 125$; the convergence tolerance is set to $\epsilon_{tol} = 5 \times 10^{-2}$; and the error indicator is chosen to be the residual-based error indicator (13) with \mathbf{R}_L set – for this purpose, and only for this purpose – to the residual associated with the first eigenpair of the PROM matrix (27).

The three aforementioned instances differ as follows:

- *Standard*. In this case the residual-based error indicator is evaluated at each iteration at every candidate parameter point in Ξ .
- *Random*. In this case, the residual-based error indicator is evaluated at each iteration at every candidate parameter point in a random subset $\Pi_m \subset \Xi$ of size $N_{\Pi} = 20$.
- *Saturation*. In this case, the greedy procedure is that described in Algorithm 12 with $N_{\Pi} = 20$ and $N'_{\Pi} = 50$. The saturation constant is set to $\tau_s = 2$ in order to capture a fluctuation of the error at some parameter points as more of these points are sampled to refine the database of PROMs.

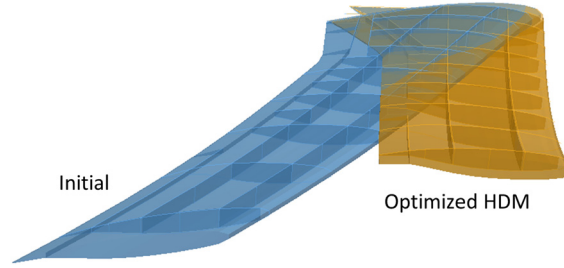
The performances of these databases are contrasted in Table 4 for $\boldsymbol{\mu} = \boldsymbol{\mu}_s$. The reader can observe that the accelerated parameter sampling procedure based on a saturation assumption described in Section 4.1 and Algorithm 12 reduces the computational cost of the standard greedy procedure by a factor greater than 7.

All performance results reported in the remainder of this paper are for the database of coupled PROMs constructed using Algorithm 12.

Table 5

Performance results for the evaluation of the flutter constraint and its sensitivities for a single queried parameter vector.

Computational model	$\mu = \mu_s (N_D = 3)$			$\mu = (\mu_s, \mu_m) (N_D = 6)$		
	# of CPUs	Wall-clock time	CPU time	# of CPUs	Wall-clock time	CPU time
HDM	32	0.342	10.934	32	0.492	15.744
Database of PROMs	1	0.019	0.019	1	0.017	0.017
Speed-up factor		18.0	575.5		28.9	926.1

**Fig. 6.** Initial and optimized designs of the ARW-2.

6.3. Speed-up factors for one online evaluation of the flutter constraint

Table 5 compares the performance of the proposed methodology for accelerating the evaluation of a discretized, parametric, linear PDE constraint with that of the HDM-based evaluation of this constraint, for a single computation of the vector of damping ratios in problem (21). For this purpose, this computation (ζ and $\frac{d\zeta}{d\mu}$) is carried out using the fixed-point iteration algorithm described in [41]. For the case where $\mu = \mu_s$, the PROM database approach is shown to deliver a speed-up factor roughly equal to 17 for the wall-clock time, and 535 for the CPU time. For the case $\mu = (\mu_s, \mu_m)$, the speed-up factor for the wall-clock time is roughly equal to 29, and that for the CPU time is slightly larger than 926.

6.4. Speed-up factors for the solution of the optimization problem with $N_D = 3$

Table 6 summarizes the results obtained for the solution of the optimization problem (21) with $\mu = \mu_s$ and $\zeta_{\text{lower}} = 4.2 \times 10^{-4}$. Given the same initial shape defined by $\mu = (0.1, 0.1, 0.1)$, both computational approaches based on the HDM and PROM database are found to deliver almost the same optimal design, after almost the same number of iterations. (The initial and optimized shapes of the ARW-2 are shown in Fig. 6). However, Fig. 7 shows that the optimization paths followed by both computational models are different. In both cases, the flutter constraint is violated initially, but satisfied at convergence to the optimal shape. After optimization, the lift to drag ratio is increased by 15.1%, and the weight and maximum von Mises stresses are increased by 4.4% and 18.2%, respectively, without violating their respective constraints. Fig. 8 shows the iteration-histories of the lift to drag ratio, minimum damping ratio, maximum von Mises stress, and weight. Again, this figure shows that the two computational models contrasted here converge to the same results, but follow different paths in general.

Table 6 also reports the CPU timing results obtained for each considered computational model. The CPU timing result reported for the offline phase of model reduction corresponds to the CPU time elapsed in constructing the PROM database using the accelerated parameter sampling procedure described in Section 4.1. The “static” constraints refer to all of the weight, von Mises, and box constraints, as well as the objective function. The flutter constraint refers as before to the constraint on the vector of damping ratios. The following observations are noteworthy:

- The CPU-based speed-up factor achieved for the online evaluation of the flutter constraint is consistent with that reported in Table 5 for a single such evaluation.
- The smaller speed-up factor achieved for the entire online phase is due to the fact that the so-called static constraints are not reduced.
- Because it requires an offline investment, the PROM database approach can be expected to deliver a (significant) speed-up factor only if this investment is amortized. As this table shows, such an investment is rarely amortized by the solution of a single problem – in this case, an optimization problem. However, it will be shown next that such an investment can be amortized if multiple optimizations are performed – for example, as in multi-start global optimization strategies.

Table 6

Performance results for the solution of the ARW-2 shape optimization problem with $\zeta_{\text{lower}} = 4.2 \times 10^{-4}$ and $\mu = \mu_s$ ($N_D = 3$).

Design	Initial	Optimized	Optimized
Computational model		HDM	PROM database
# of CPUs		32	variable
μ_s	(0.1,0.1,0.1)	(−0.1,−0.0814,0.1)	(−0.1,−0.0807,0.1)
L/D	10.6	12.2	12.2
Weight (lbs)	342.6	357.7	357.7
Min. ζ	2.24×10^{-4}	4.22×10^{-4}	4.25×10^{-4}
Max. σ_{VM} (psi)	18,731.9	22,139.5	22,139.0
# of iterations		5	6
# of function evaluations		24	22
CPU time (hours)			
Offline phase (32 CPUs)		0	268.8
Online evaluation of the static constraints (32 CPUs)		14.2	14.1
Online evaluation of the flutter constraint (1 CPU)		240.5	0.43
Total (online)		254.7	14.5
Total (offline + online)		254.7	283.3
Speed-up factors based on CPU time			
Speed-up factor for flutter constraint	1		559.3
Speed-up factor for online phase	1		17.6
Total speed-up factor	1		0.7

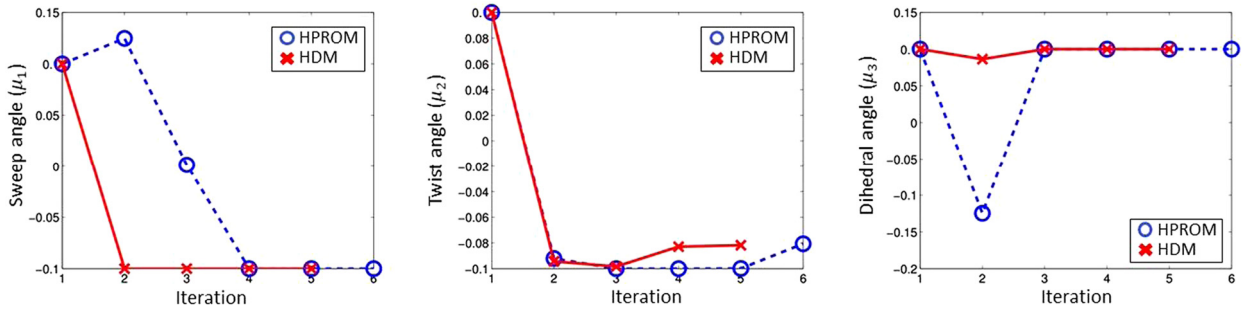
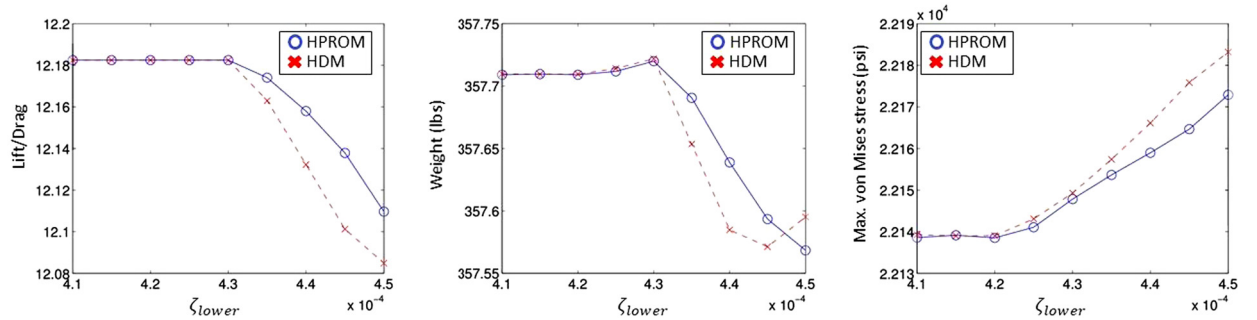
**Fig. 7.** Iteration-histories of the shape parameters.**Fig. 8.** Iteration-histories of various QoIs.

Table 7 reports the CPU timing results obtained for multiple instances of the flutter optimization problem (21), corresponding to the following nine different values of the lower bound for the damping ratios, $\zeta_{\text{lower}} = \{4.1, 4.15, 4.2, 4.25, 4.3, 4.35, 4.4, 4.45, 4.5\} \times 10^{-4}$, for the purpose of studying the influence of this flutter threshold on the optimized design. Furthermore for each damping ratio lower bound, a multi-start strategy with ten different initial points is applied in order to provide a more robust search for a global optimizer. Therefore, the performance results reported in Table 7 correspond to 90 independent optimization problems that are solved using both of the HDM and PROM database computational models. In this case, the reader can observe that:

- The proposed PROM database approach delivers a speed-up factor based on the CPU time that is equal to 1,459.3 for the flutter constraint, and a speed-up factor based on the CPU time that is equal to 43.6 for the overall online phase.

Table 7

Performance results for the solution of nine instances of the ARW-2 shape optimization problem with $\mu = \mu_s$ ($N_D = 3$), but different values of ζ_{lower} .

Computational model	HDM	PROM Database
# of CPUs	32	variable
CPU time (hours)		
Offline phase (32 CPUs)	0	268.8
Online evaluation of the static constraints (32 CPUs)	2,065.1	755.7
Online evaluation of the flutter constraints (1 CPU)	31,813.5	21.8
Online miscellaneous	14.2	0.06
Total (online)	33,892.8	777.6
Total (offline + online)	33,892.8	1,046.4
Speed-up factors based on CPU time		
Speed-up factor for flutter constraints	1	1,459.3
Speed-up factor for online phase	1	43.6
Total speed-up factor	1	32.4

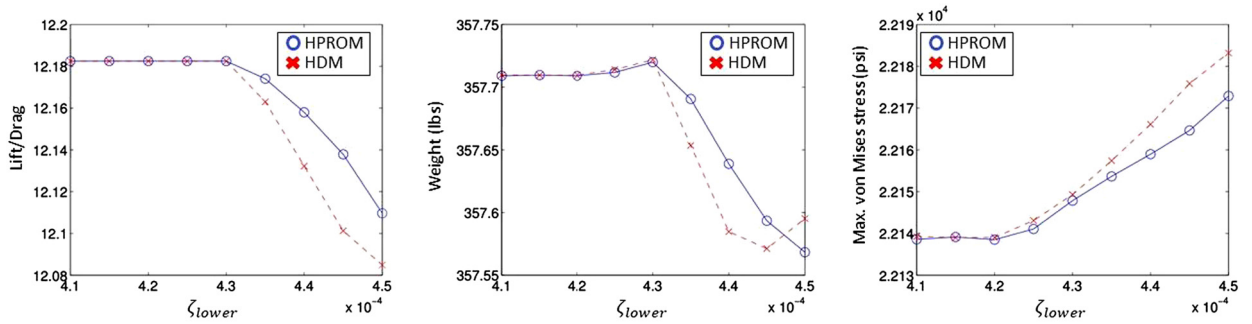


Fig. 9. Influence of the lower bound of the flutter constraint on various QoIs associated with the computed optimal design.

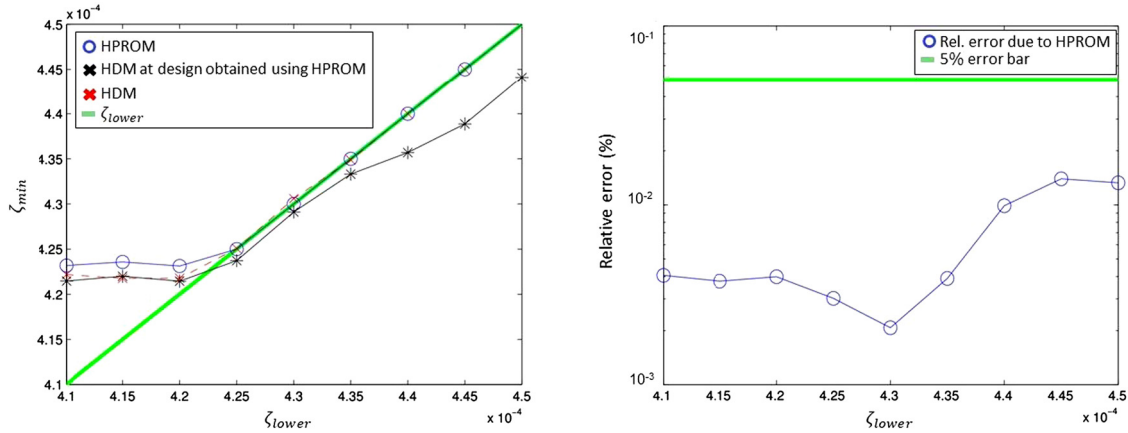


Fig. 10. Influence of the lower bound of the flutter constraint on: the minimum damping ratio of the computed optimal design (left); and the relative error associated with the computation of this QoI using the PROM database approach.

- This approach achieves a speed-up factor based on the CPU time that is equal to 32.4 for the entire computations, including those incurred by the construction offline of the database of PROMs, which confirms the observation made above.

Fig. 9 graphically depicts the influence of ζ_{lower} on the optimal lift to drag ratio, the optimal weight, and the maximum von Mises stress associated in each case with the optimal design. Overall, this influence appears to be minimal. Similarly, Fig. 10-left shows the influence of ζ_{lower} on the minimum damping ratio associated with an optimal design.

Fig. 10-right reports on the relative error associated with the computation of the minimum damping ratio using the PROM database approach, for all considered values of ζ_{lower} . It shows that in all cases, this error is significantly below the convergence threshold of 5% adopted during the construction offline of the PROM database.

Table 8

Performance results for the solution of the ARW-2 shape optimization problem with $\zeta_{\text{lower}} = 4.2 \times 10^{-4}$ and $\mu = (\mu_s, \mu_m)$ ($N_D = 6$).

Design	Initial	Optimized	Optimized
Computational model		HDM	PROM Database
# of CPUs		32	variable
μ_s	(0,0,0)	(−0.1,−0.1,0.1)	(−0.1,−0.1,0.1)
μ_m	(0,0,0)	(0.1,0.1,−0.015)	(0.1,0.1,0.013)
L/D	11.3	12.2	12.2
Weight (lbs)	349.9	349.8	366.3
Min. ζ	3.7×10^{-4}	4.2×10^{-4}	4.2×10^{-4}
Max. σ_{VM} (psi)	20,297.1	18,991.9	18,953.3
# of iterations		6	5
# of function evaluations		11	9
CPU time (hours)			
Offline phase (32 CPUs)		0	8,800
Online evaluation of the static constraints (32 CPUs)		11.0	12.8
Online evaluation of the flutter constraint (1 CPU)		173.2	0.15
Total (online)		184.2	13.0
Total (offline + online)		184.2	8,813
Speed-up factors based on CPU time			
Speed-up factor for flutter constraint		1	1,154.7
Speed-up factor for online phase		1	14.2
Total speed-up factor		1	0.02

6.5. Speed-up factors for the solution of the optimization problem with $N_D = 6$

Table 8 presents the performance results obtained for the solution of the optimization problem (21) with $\mu = (\mu_s, \mu_m)$ and $\zeta_{\text{lower}} = 4.2 \times 10^{-4}$. As in the case of three shape parameters ($\mu = \mu_s$), both computational models based on the HDM and PROM database approaches are found to deliver almost the same optimal shape, given the same initial design. However, the two different computational models lead in this case to two slightly different optimal structural parameters. In both cases, the flutter constraint is violated initially but satisfied at convergence, and the lift to drag ratio is increased by 8.0%. However, the weight is increased by 4.7% in the optimal design found using the PROM database as a computational model, but is maintained almost constant in the counterpart design found by the HDM-based optimization process. The maximum von Mises stress is decreased by 6.6% in the optimal design found using the PROM database computational model, and by 6.4% in the counterpart design found using the HDM-based approach. Although the optimal shape parameters obtained in this case – $\mu_s = (-0.1, -0.1, 0.1)$ – are similar to those obtained in the case of optimization using only the three shape parameters (see Table 6), the maximum von Mises stress experienced in this case by the optimal design is smaller than its counterpart experienced by the optimal design obtained when $\mu = \mu_s$. This lower value is achieved by increasing the thickness of the stiffeners without violating the weight constraint. The reader can also observe that for both considered computational models, similar iteration counts and similar numbers of function evaluations are performed.

Table 8 also reports the CPU time results and associated speed-up factors achieved for this instance of the optimization problem (21). Due to the larger computational expense induced by the computation of a larger number of sensitivities, the CPU time-based speed-up factor achieved by the PROM database approach for a single evaluation of the flutter constraint and its sensitivities is substantially higher (speed-up factor = 1,154.7) than for the case where $N_D = 3$ (speed-up factor = 559.3). The speed-up factor for the overall online phase of model reduction is similar however in magnitude. Again, the total speed-up factor for the combined offline and online phases is less than 1, due to the computational expense incurred by building the database. As demonstrated however for the case where $N_D = 3$, this computational overhead can be amortized and significant overall speed-up factors can be achieved when many optimization problems are solved using the same PROM database.

7. Conclusions

A novel methodology is introduced in this paper for accelerating the solution of Partial Differential Equation (PDE)-constrained optimization problems where at least one PDE-based constraint is linear. This methodology consists in constructing offline a database of pointwise, linear, Projection-based Reduced-Order Models (PROMs), and equipping it with two essential computational technologies: an offline pre-processing algorithm for enforcing mathematical consistency among the pre-computed PROMs that are stored in the database; and an online algorithm for interpolating on matrix manifolds the pre-computed PROMs and their sensitivities at any queried but unsampled parameter point in the design parameter space. To reduce the computational overhead incurred by the construction offline of the PROM database, a parameter sampling procedure based on an appropriate saturation assumption is proposed to minimize the number of pointwise PROMs to be pre-computed and stored in the database, while maximizing the accuracy of the online interpolation of these PROMs at

any queried but unsampled parameter point. The interpolation on matrix manifolds, which preserves some desirable properties of the PROMs to be interpolated and their sensitivities, incurs the standard interpolation in the tangent space to a considered manifold – which is a vector space – of some parametric quantities. It is shown in this paper that when the design parameter space is high-dimensional, such an interpolation can be practically and effectively performed using radial basis functions. A new and rigorous approach for interpolating on matrix manifolds matrix sensitivities is also presented. The PROM database approach and its peripherals are illustrated with the solution of a number of different instances of a PDE-constrained, aeroelastic, optimization problem associated with NASA's ARW-2 wing and involving a linear, PDE-based flutter constraint. Using this problem as a background problem, it is also shown in this paper that when multiple optimizations are performed – for example, as in multi-start global optimization strategies – the computational overhead incurred by the construction offline of a PROM database can be amortized and a significant overall, CPU time-based, speed-up factor can be achieved. Finally, it is noted that given the recent work published in [42], the methodology presented in this paper is readily extendible to the case where the PDE-based optimization problem contains a combination of linear and nonlinear PDE-based constraints.

Credit authorship contribution statement

All authors have contributed to all technical aspects of this paper. Prof. Farhat was also responsible for the Funding acquisition and for Supervision.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgements

Youngsoo Choi acknowledges partial support by the Office of Naval Research under Grant N00014-11-1-0707 and partial support by the Army High Performance Computing Research Center under Cooperative Agreement W911NF-07-2-0027, while in residence at Stanford University. Gabriele Boncoraglio and Charbel Farhat acknowledge partial support by the Office of Naval Research under Grant N00014-17-1-2749, partial support by the Boeing Company under Contract Sponsor Ref. 134824, and partial support by a research grant from the King Abdulaziz City for Science and Technology (KACST). Spenser Anderson acknowledges the support by the Department of Defense through the National Defense Science & Engineering Graduate Fellowship (NDSEG) Program. This document does not necessarily reflect the position of these institutions, and no official endorsement should be inferred.

Appendix 8

Real-time computation of the sensitivity of the vector of damping ratios with respect to the optimization parameters, $\frac{\partial \xi}{\partial \mu_i}$, $i = 1, \dots, N_D$ For each eigenpair $(\lambda_j, \mathbf{q}_{rj})$, $j = 1, \dots, k_f + 2k_s$, the differentiation of (28) with respect to μ_i , $i = 1, \dots, N_D$, leads to

$$\frac{\partial \mathbf{N}_r}{\partial \mu_i} \mathbf{q}_{rj} - \frac{\partial \lambda_j}{\partial \mu_i} \mathbf{q}_{rj} + (\mathbf{N}_r - \lambda_j \mathbf{I}) \frac{\partial \mathbf{q}_{rj}}{\partial \mu_i} = \mathbf{0}$$

where the reduced-order matrix \mathbf{N}_r and the associated reduced-order vector \mathbf{q}_r are defined in (27) and (26), respectively, \mathbf{q}_{rj} is the j -th right eigenvector of \mathbf{N}_r , λ_j is the corresponding eigenvalue, \mathbf{I} is the identity matrix, and μ_i is the i -th component of the parameter vector $\boldsymbol{\mu} \in \mathcal{D} \subset \mathbb{R}^{N_D}$. Multiplying (8) from the left by \mathbf{p}_{rj}^H , where \mathbf{p}_{rj} denotes the j -th left eigenvector of \mathbf{N}_r and the superscript H designates the transpose of the conjugate of a complex-valued quantity, and noting that $\mathbf{p}_{rj}^T (\mathbf{N}_r - \lambda_j \mathbf{I}) = \mathbf{0}$ gives

$$\frac{\partial \lambda_j}{\partial \mu_i} = \frac{\mathbf{p}_{rj}^H \frac{\partial \mathbf{N}_r}{\partial \mu_i} \mathbf{q}_{rj}}{\mathbf{p}_{rj}^H \mathbf{q}_{rj}}$$

In the above result, the sensitivity matrix $\frac{\partial \mathbf{N}_r}{\partial \mu_i}$ can be computed at any queried but unsampled parameter point in real-time, block-by-block, using interpolation on matrix manifolds as described in Section 6.1.2 (see also Table 1 and Table 2). Then, from Equation (29) and the chain rule, it follows that the sensitivity of each damping ratio ξ_j with respect to any parameter component μ_i can be computed in real-time as follows

$$\frac{\partial \xi_j}{\partial \mu_i} = \frac{\partial \lambda_j^R}{\partial \mu_i} \left(\frac{(\lambda_j^R)^2}{|\lambda_j|^3} - \frac{1}{|\lambda_j|} \right) + \frac{\partial \lambda_j^I}{\partial \mu_i} \frac{\lambda_j^R \lambda_j^I}{|\lambda_j|^3}, \quad j = 1, \dots, k_f + 2k_s, \quad i = 1, \dots, N_D \quad (30)$$

References

- [1] B. Moore, Principal component analysis in linear systems: controllability, observability, and model reduction, *IEEE Trans. Autom. Control* 26 (1981) 17–32.
- [2] L. Sirovich, Turbulence and the dynamics of coherent structures. Part I: coherent structures, *Q. Appl. Math.* 45 (1987) 561–571.
- [3] B. Epureanu, A parametric analysis of reduced order models of viscous flows in turbomachinery, *J. Fluids Struct.* 17 (2003) 971–982.
- [4] T. Lieu, C. Farhat, Adaptation of aeroelastic reduced-order models and application to an F-16 configuration, *AIAA J.* 45 (2007) 1244–1269.
- [5] P.A. LeGresley, J. Alonso, Airfoil design optimization using reduced order models based on proper orthogonal decomposition, in: *AIAA Paper 2000-2545 Fluids 2000 Conference and Exhibit*, Denver, CO, American Institute of Aeronautics and Astronautics, Reston, Virginia, 2000, pp. 1–14.
- [6] G. Weickum, M. Eldred, K. Maute, A multi-point reduced-order modeling approach of transient structural dynamics with application to robust design optimization, *Struct. Multidiscip. Optim.* 38 (2009) 599–611.
- [7] D. Amsallem, M. Zahr, Y. Choi, C. Farhat, Design optimization using hyper-reduced-order models, *Struct. Multidiscip. Optim.* 51 (2015) 919–940.
- [8] D. Amsallem, M.J. Zahr, C. Farhat, Nonlinear model order reduction based on local reduced-order bases, *Int. J. Numer. Methods Eng.* 92 (2012) 891–916.
- [9] M. Fahl, E.W. Sachs, Reduced order modelling approaches to PDE-constrained optimization based on proper orthogonal decomposition, in: *Large-Scale PDE-Constrained Optimization*, Springer, 2003, pp. 268–280.
- [10] Y. Yue, K. Meerbergen, Accelerating optimization of parametric linear systems by model order reduction, *SIAM J. Optim.* 23 (2013) 1344–1370.
- [11] M.J. Zahr, D. Amsallem, C. Farhat, Construction of parametrically-robust CFD-based reduced-order models for PDE-constrained optimization, in: *AIAA Paper 2013-2845, 21st AIAA Computational Fluid Dynamics Conference*, San Diego, CA, June 26–29, 2013, American Institute of Aeronautics and Astronautics, Reston, Virginia, 2013, pp. 1–11.
- [12] M.J. Zahr, C. Farhat, Progressive construction of a parametric reduced-order model for PDE-constrained optimization, *Int. J. Numer. Methods Eng.* 102 (2015) 1111–1135.
- [13] D. Amsallem, J. Cortial, K. Carlberg, C. Farhat, A method for interpolating on manifolds structural dynamics reduced-order models, *Int. J. Numer. Methods Eng.* 80 (2009) 1241–1258.
- [14] D. Amsallem, C. Farhat, An online method for interpolating linear parametric reduced-order models, *SIAM J. Sci. Comput.* 33 (2011) 2169–2198.
- [15] D. Amsallem, R. Tezaur, C. Farhat, Real-time solution of linear computational problems using databases of parametric reduced-order models with arbitrary underlying meshes, *J. Comput. Phys.* 326 (2016) 373–397.
- [16] P.E. Gill, W. Murray, M.A. Saunders, Snopt: an SQP algorithm for large-scale constrained optimization, *SIAM J. Optim.* 12 (2002) 979–1006.
- [17] A.R. Conn, N.I. Gould, P.L. Toint, *Trust Region Methods*, MOS-SIAM Series on Optimization, vol. 1, SIAM, 2000.
- [18] A. Wächter, L.T. Biegler, On the implementation of an interior-point filter line-search algorithm for large-scale nonlinear programming, *Math. Program.* 106 (2006) 25–57.
- [19] D. Amsallem, J. Cortial, C. Farhat, Towards real-time computational-fluid-dynamics-based aeroelastic computations using a database of reduced-order information, *AIAA J.* 48 (2010) 2029–2037.
- [20] D. Amsallem, C. Farhat, Interpolation method for adapting reduced-order models and application to aeroelasticity, *AIAA J.* 46 (2008) 1803–1813.
- [21] T. Bui-Thanh, K. Willcox, O. Ghattas, Model reduction for large-scale systems with high-dimensional parametric input space, *SIAM J. Sci. Comput.* 30 (2008) 3270–3288.
- [22] T. Bui-Thanh, K. Willcox, O. Ghattas, Parametric reduced-order models for probabilistic analysis of unsteady aerodynamic applications, *AIAA J.* 46 (2008) 2520–2529.
- [23] M.A. Grepl, A.T. Patera, A posteriori error bounds for reduced-basis approximations of parametrized parabolic partial differential equations, *ESAIM Math. Model. Numer. Anal.* 39 (2005) 157–181.
- [24] J.S. Hesthaven, B. Stamm, S. Zhang, Efficient greedy algorithms for high-dimensional parameter spaces with applications to empirical interpolation and reduced basis methods, *ESAIM Math. Model. Numer. Anal.* 48 (2014) 259–283.
- [25] A. Paul-Dubois-Taine, D. Amsallem, An adaptive and efficient greedy procedure for the optimal training of parametric reduced-order models, *Int. J. Numer. Methods Eng.* 102 (2015) 1262–1292.
- [26] K. Veroy, C. Prud'homme, D.V. Rovas, A.T. Patera, A posteriori error bounds for reduced-basis approximation of parametrized noncoercive and nonlinear elliptic partial differential equations, in: *Proceedings of the 16th AIAA Computational Fluid Dynamics Conference*, vol. 3847, 2012, pp. 23–26.
- [27] C. Prud'homme, D.V. Rovas, K. Veroy, L. Machiels, Y. Maday, A.T. Patera, G. Turinici, Reliable real-time solution of parametrized partial differential equations: reduced-basis output bound methods, *J. Fluids Eng.* 124 (2002) 70–80.
- [28] B. Haasdonk, M. Ohlberger, Reduced basis method for finite volume approximations of parametrized linear evolution equations, *ESAIM Math. Model. Numer. Anal.* 42 (2008) 277–302.
- [29] B. Haasdonk, M. Ohlberger, Efficient reduced models and a posteriori error estimation for parametrized dynamical systems by offline/online decomposition, *Math. Comput. Model. Dyn. Syst.* 17 (2011) 145–161.
- [30] D. Amsallem, U. Hetmaniuk, A posteriori error estimators for linear reduced-order models using Krylov-based integrators, *Int. J. Numer. Methods Eng.* 102 (2015) 1238–1261.
- [31] I. Najfeld, T.F. Havel, Derivatives of the matrix exponential and their computation, *Adv. Appl. Math.* 16 (1995) 321–375.
- [32] G. Fasshauer, Adaptive least squares fitting with radial basis functions on the sphere, in: *Mathematical Methods for Curves and Surfaces*, 1995, pp. 141–150.
- [33] C. Farhat, M. Lesoinne, N. Maman, Mixed explicit/implicit time integration of coupled aeroelastic problems: three-field formulation, geometric conservation and distributed solution, *Int. J. Numer. Methods Fluids* 21 (1995) 807–835.
- [34] C. Farhat, P. Peuzaine, G. Brown, Application of a three-field nonlinear fluid-structure formulation to the prediction of the aeroelastic parameters of an F-16 fighter, *Comput. Fluids* 32 (2003) 3–29.
- [35] M. Lesoinne, M. Sarkis, U. Hetmaniuk, C. Farhat, A linearized method for the frequency analysis of three-dimensional fluid/structure interaction problems in all flow regimes, *Comput. Methods Appl. Mech. Eng.* 190 (2001) 3121–3146.
- [36] M. Lesoinne, C. Farhat, CFD-based aeroelastic eigensolver for the subsonic, transonic, and supersonic regimes, *J. Aircr.* 38 (2001) 628–635.
- [37] M.C. Sandford, D.A. Seidel, C.V. Spain, Geometrical and structural properties of an aeroelastic research wing (arw-2), *NASA Technical Memorandum* 4110, 1989.
- [38] C. Farhat, C. Degand, B. Koobus, M. Lesoinne, Torsional springs for two-dimensional dynamic unstructured fluid meshes, *Comput. Methods Appl. Mech. Eng.* 163 (1998) 231–245.
- [39] T. Kim, Frequency-domain Karhunen-Loeve method and its application to linear dynamic systems, *AIAA J.* 36 (1998) 2117–2123.
- [40] D. Amsallem, C. Farhat, On the stability of reduced-order linearized computational fluid dynamics models based on POD and Galerkin projection: descriptor vs non-descriptor forms, in: *Reduced Order Methods for Modeling and Computational Reduction*, Springer, 2014, pp. 215–233.
- [41] D. Amsallem, D. Neumann, Y. Choi, C. Farhat, Linearized aeroelastic computation in the frequency domain based on computational fluid dynamics, *arXiv:1506.07441 [physics.flu-dyn]*, 2015, pp. 1–23.
- [42] K. Washabaugh, Z.J. Matthew, C. Farhat, On the use of discrete nonlinear reduced-order models for the prediction of steady-state flows past parametrically deformed complex geometries, in: *AIAA Paper 2016-1814, AIAA SciTech Forum Conference*, San Diego, CA, January 4–8, 2013, American Institute of Aeronautics and Astronautics, Reston, Virginia, 2016.