



ELSEVIER

Contents lists available at ScienceDirect

C. R. Acad. Sci. Paris, Ser. I

www.sciencedirect.com



Partial Differential Equations/Numerical Analysis

A combination between the reduced basis method and the ANOVA expansion: On the computation of sensitivity indices



Une combinaison entre la méthode des bases réduites et l'extension ANOVA : Sur le calcul des indices de sensibilité

Denis Devaud^a, Andrea Manzoni^{b,c}, Gianluigi Rozza^{b,c}

^a EPFL SMA, CH-1015, Lausanne, Switzerland

^b EPFL MATHICSE-CMCS, Station 8-MA, CH-1015, Lausanne, Switzerland

^c SISSA, International School for Advanced Studies, Mathlab, Via Bonomea 265, 34136 Trieste, Italy

ARTICLE INFO

Article history:

Received 9 April 2013

Accepted after revision 29 July 2013

Available online 19 August 2013

Presented by Olivier Pironneau

ABSTRACT

We consider a method to efficiently evaluate in a real-time context an output based on the numerical solution of a partial differential equation depending on a large number of parameters. We state a result allowing to improve the computational performance of a three-step RB–ANOVA–RB method. This is a combination of the reduced basis (RB) method and the analysis of variations (ANOVA) expansion, aiming at compressing the parameter space without affecting the accuracy of the output. The idea of this method is to compute a first (coarse) RB approximation of the output of interest involving all the parameter components, but with a large tolerance on the *a posteriori* error estimate; then, we evaluate the ANOVA expansion of the output and freeze the least important parameter components; finally, considering a restricted model involving just the retained parameter components, we compute a second (fine) RB approximation with a smaller tolerance on the *a posteriori* error estimate. The fine RB approximation entails lower computational costs than the coarse one, because of the reduction of parameter dimensionality. Our result provides a criterion to avoid the computation of those terms in the ANOVA expansion that are related to the interaction between parameters in the bilinear form, thus making the RB–ANOVA–RB procedure computationally more feasible.

© 2013 Académie des sciences. Published by Elsevier Masson SAS. All rights reserved.

R É S U M É

Nous considérons une méthode permettant d'évaluer en temps réel, de manière efficace, une fonctionnelle basée sur la solution numérique d'une équation aux dérivées partielles dépendant d'un grand nombre de paramètres. Nous présentons un résultat qui permet d'améliorer la méthode en trois étapes, RB–ANOVA–RB. Cette dernière est une combinaison de la méthode des bases réduites (RB) et de la méthode d'expansion d'analyse des variations (ANOVA). Le but est de pouvoir compresser l'espace des paramètres sans affecter la précision de notre fonctionnelle. Dans un premier temps, nous calculons une première approximation (grossière) RB de notre problème, en considérant toutes les composantes des paramètres. Ensuite, nous utilisons l'approximation obtenue pour calculer l'expansion ANOVA de la fonctionnelle afin de déterminer l'influence de chacune des composantes de nos paramètres sur ce dernier et fixer les moins influentes. Finalement, une deuxième approximation (fine) RB est faite sur le modèle ne contenant que les composantes les plus

E-mail addresses: denis.devaud@epfl.ch (D. Devaud), amanzoni@sissa.it (A. Manzoni), gianluigi.rozza@epfl.ch (G. Rozza).

importantes. Le résultat que nous présentons ici donne un critère pour éviter le calcul de termes, basé sur l'interaction des composantes des paramètres dans la forme bilinéaire, permettant ainsi de diminuer drastiquement les coûts computationnels liés à l'expansion ANOVA.

© 2013 Académie des sciences. Published by Elsevier Masson SAS. All rights reserved.

Version française abrégée

Ce que nous présentons ici est une amélioration de la méthode en trois étapes RB–ANOVA–RB originalement introduite par Hesthaven et Zhang dans [6] pour évaluer des fonctionnelles linéaires dépendant de la solution d'une équation aux dérivées partielles paramétrique. Dans la méthode ANOVA, le but est de réduire les coûts computationnels liés à la partie dite *online* de la méthode des bases réduites, en étant intéressé à évaluer une fonctionnelle sans affecter la précision de notre approximation. Jusqu'ici, rien n'a été fait pour réduire les calculs de la phase *offline*, qui sont souvent très élevés. Nous présentons alors un résultat qui permet d'éviter le calcul de certains termes de l'expansion ANOVA et, par conséquent, d'en réduire les coûts computationnels, rendant ainsi la méthode beaucoup plus efficace.

Le problème considéré est le suivant : étant donné $\boldsymbol{\mu} \in \mathcal{D} = \bigcup_{i=1}^d [a_i, b_i]$ avec $a_i, b_i \in \mathbb{R}$, évaluer la fonctionnelle $s(\boldsymbol{\mu}) = l(u(\boldsymbol{\mu}))$, où $u(\boldsymbol{\mu}) \in X(\Omega(\boldsymbol{\mu}))$ est la solution du problème $a(u(\boldsymbol{\mu}), v; \boldsymbol{\mu}) = f(v)$, $\forall v \in X(\Omega(\boldsymbol{\mu}))$. Ici, la forme bilinéaire $a(\cdot, \cdot; \boldsymbol{\mu})$ dépend de manière affine de $\boldsymbol{\mu}$ (voir (1)). De plus, nous imposons que $a(\cdot, \cdot; \boldsymbol{\mu})$ soit symétrique et que $l = f$. Dans notre cas, $a(\cdot, \cdot; \boldsymbol{\mu})$ est associée à un opérateur aux dérivées partielles paramétrique.

Afin de résoudre ce problème, nous considérons une méthode en trois étapes, RB–ANOVA–RB, qui est une combinaison de la méthode des bases réduites (RB) et de la méthode d'analyse des variations dite ANOVA [6]. Le but est de pouvoir compresser l'espace des paramètres sans réduire la précision sur la fonctionnelle. L'idée est d'appliquer une première fois la méthode des bases réduites à notre problème, ce qui donne une approximation $s_N^{\mathcal{N}}(\boldsymbol{\mu})$ de notre fonctionnelle. Ici, \mathcal{N} est la dimension de l'espace éléments finis et N celui de l'espace bases réduites. Ensuite, nous calculons une expansion ANOVA de notre fonctionnelle (voir (2), Section 1) afin d'évaluer l'influence de chaque composante des paramètres, et fixons les moins influentes. Finalement, une approximation fine en bases réduites de notre problème, dont l'espace des paramètres a été compressé, est calculée, donnant lieu à des coûts computationnels beaucoup plus faibles que ceux liés au problème original. La raison pour laquelle une telle méthode est utilisée est que lorsque le nombre de paramètres augmente, les coûts computationnels augmentent considérablement. Le but est donc de réduire le nombre de paramètres sans affecter la précision de notre fonctionnelle. L'intuition de cette méthode vient du fait que, dans le cas où les paramètres sont géométriques, on peut s'attendre à trouver des colinéarités entre les composantes, tandis que dans le cas de paramètres physiques, l'influence sur la fonctionnelle peut être, dans certains cas, négligeable.

Au cours de cette procédure, le calcul de certains termes de l'expansion ANOVA peut être très coûteux. Nous donnons alors un critère permettant de déterminer a priori si certains de ces termes sont nuls. Étant donnée l'expansion affine (définie en (1)) de notre forme bilinéaire en termes $\Theta_a^q(\boldsymbol{\mu})a^q(\cdot, \cdot; \boldsymbol{\mu})$, $1 \leq q \leq Q$, et un ensemble τ d'indices (correspondant à certaines composantes de nos vecteurs de paramètres), si pour chaque $1 \leq q \leq Q$ il existe un indice $i_q \in \tau$ tel que $\Theta_a^q(\cdot)$ soit indépendant de $\boldsymbol{\mu}_{i_q}$, alors $S_\tau = 0$, où S_τ est l'indice de sensibilité – défini en (3) – utilisé pour déterminer l'influence des composantes indicées par τ sur la fonctionnelle. Ce résultat est très utile, car il permet de gagner un temps de calcul considérable en étudiant les termes de la décomposition affine de la forme bilinéaire $a(\cdot, \cdot; \boldsymbol{\mu})$. Ceci est d'autant plus avantageux que, lorsque $|\tau|$ augmente, les coûts computationnels liés au calcul de S_τ augmentent considérablement. Dans le cas où les paramètres sont physiques, les termes $\Theta_a^q(\cdot)$ dépendent souvent d'une ou deux composantes de $\boldsymbol{\mu}$ (car liés à la propriété physique d'une région du domaine) et donc certains termes pour lesquels $|\tau| = 2$ et tous ceux qui sont tels que $|\tau| \geq 3$ n'ont pas besoin d'être calculés. Par ailleurs, dans le cas de paramètres géométriques, il est parfois possible de découper le domaine de manière à ce que le critère donné soit respecté.

L'exemple numérique considéré est un cas de conductivité thermique pure. On divise un carré en quatre bandes, chacune ayant son propre coefficient de conductivité. On impose des conditions de type Dirichlet homogènes sur le bord supérieur, Neumann unitaires sur le bord inférieur et Neumann homogènes sur les bords latéraux. Dans ce cas, la fonctionnelle est donnée par l'intégrale de la chaleur sur le bord inférieur (car $l = f$). Ainsi, chaque Θ_a^q , $q = 1, \dots, 4$, ne dépend que d'une seule composante du paramètre, puisque $\Theta_a^q(\boldsymbol{\mu}) = 100^{2\mu_q - 1}$ si $q = 1$ et $\Theta_a^q(\boldsymbol{\mu}) = 1.1^{2\mu_q - 1}$ sinon. On calcule tous les termes nuls a priori et on montre qu'ils sont bien nuls et que leur évaluation représente environ 90% du temps total de calcul.

Par conséquent, dans le cas où le problème considéré comprend un grand nombre de paramètres, il est important d'utiliser un tel résultat afin d'éviter que les coûts computationnels ne soient trop élevés.

1. Introduction

Let us consider the following problem: given $\boldsymbol{\mu} \in \mathcal{D} = \bigcup_{i=1}^p [a_i, b_i]$ with $a_i, b_i \in \mathbb{R}$, evaluate the output $s(\boldsymbol{\mu}) = l(u(\boldsymbol{\mu}))$, where $u(\boldsymbol{\mu}) \in X(\Omega(\boldsymbol{\mu}))$ is the solution of a parameterized elliptic PDE: $a(u(\boldsymbol{\mu}), v; \boldsymbol{\mu}) = f(v)$, $\forall v \in X(\Omega)$. Here $\mathcal{D} \subseteq \mathbb{R}^p$, $p \geq 1$, is the parameter domain; $\Omega \subseteq \mathbb{R}^d$, $d = 2$ or 3 , is a (smooth enough) spatial domain, independent of $\boldsymbol{\mu}$ (otherwise, a transformation needs to be considered to a reference domain, which introduces a Jacobian in the bilinear form); X is a

Hilbert space on Ω , such that $(H_0^1(\Omega))^v \subseteq X \subseteq (H^1(\Omega))^v$; $a(\cdot, \cdot; \boldsymbol{\mu}) : X \times X \rightarrow \mathbb{R}$ is a continuous and coercive bilinear form; $f : X \rightarrow \mathbb{R}$ is a continuous linear form; $l : \mathcal{D} \rightarrow \mathbb{R}$ is the linear output and u is the field variable. A crucial hypothesis for the offline–online procedure of the RB method used to solve this problem is that the form $a(\cdot, \cdot; \boldsymbol{\mu})$ is linear in parameter with respect to $\boldsymbol{\mu}$ (affinity assumption), so that:

$$a(u, v; \boldsymbol{\mu}) = \sum_{q=1}^Q \Theta_a^q(\boldsymbol{\mu}) a^q(u, v) \quad \forall u, v \in X, \forall \boldsymbol{\mu} \in \mathcal{D}, \tag{1}$$

where $a^q(\cdot, \cdot) : X \times X \rightarrow \mathbb{R}$ are continuous $\boldsymbol{\mu}$ -independent bilinear forms and $\Theta_a^q : \mathcal{D} \rightarrow \mathbb{R}$ are $\boldsymbol{\mu}$ -dependent functions for $1 \leq q \leq Q$. The same kind of assumption holds also for f . In what follows, we assume that f is $\boldsymbol{\mu}$ -independent for the sake of simplicity. Moreover, we assume that $a(\cdot, \cdot; \boldsymbol{\mu})$ and $a^q(\cdot, \cdot; \boldsymbol{\mu})$ are symmetric for $1 \leq q \leq Q$, $\forall \boldsymbol{\mu} \in \mathcal{D}$, and that $l = f$; this is the so-called *compliant* case (see, e.g., [8] for more details).

We introduce a finite element (FE) space $\mathcal{X}^{\mathcal{N}}$ of dimension $\dim(\mathcal{X}^{\mathcal{N}}) = \mathcal{N} < \infty$ and consider that the FE solution $u^{\mathcal{N}}(\boldsymbol{\mu}) \in \mathcal{X}^{\mathcal{N}}$ is the truth one. The problem then becomes to evaluate the output $s^{\mathcal{N}}(\boldsymbol{\mu}) = f(u^{\mathcal{N}}(\boldsymbol{\mu}))$ for $\boldsymbol{\mu} \in \mathcal{D}$, where $u^{\mathcal{N}}(\boldsymbol{\mu}) \in X^{\mathcal{N}}$ is the solution of $a(u^{\mathcal{N}}(\boldsymbol{\mu}), v; \boldsymbol{\mu}) = f(v)$, $\forall v \in X^{\mathcal{N}}$. Very often, the dimension \mathcal{N} is too large to perform real-time evaluations, for any parameter value $\boldsymbol{\mu} \in \mathcal{D}$. To make it possible, we introduce an RB approximation: we consider a collection of parameters $\boldsymbol{\mu}^1, \dots, \boldsymbol{\mu}^N$ properly selected by a greedy algorithm based on *a posteriori* residual-based error bounds (see, e.g., [8]). Then, we define the RB approximation space as $X_N^{\mathcal{N}} := \text{span}\{u^{\mathcal{N}}(\boldsymbol{\mu}^j) : 1 \leq j \leq N\}$, where each $u^{\mathcal{N}}(\boldsymbol{\mu}^j)$ is the FE solution corresponding to the input parameter $\boldsymbol{\mu}^j$.

The RB formulation is then the following one: given $\boldsymbol{\mu} \in \mathcal{D}$, evaluate the output $s_N^{\mathcal{N}}(\boldsymbol{\mu}) = l(u_N^{\mathcal{N}}(\boldsymbol{\mu}))$, where $u_N^{\mathcal{N}}(\boldsymbol{\mu}) \in X_N^{\mathcal{N}}$ is the solution of $a(u_N^{\mathcal{N}}(\boldsymbol{\mu}), v; \boldsymbol{\mu}) = f(v)$, $\forall v \in X_N^{\mathcal{N}}$. Using the assumption (1), the RB method is divided in two steps. In the first *offline* stage, we compute the $\boldsymbol{\mu}$ -independent matrices A^q and the right-hand side vector F defined by $(A^q)_{i,j} = a^q(\phi_i, \phi_j)$ and $(F)_j = f(\phi_j)$, respectively, where ϕ_j are the basis functions of the reduced space $X_N^{\mathcal{N}}$. In the *online* stage, for any new parameter $\boldsymbol{\mu}$, we compute the coefficients $\Theta_a^q(\boldsymbol{\mu})$, we assemble the matrix $A_N = \sum_{q=1}^Q \Theta_a^q(\boldsymbol{\mu}) A^q$. This linear system is very small (since $N \ll \mathcal{N}$) and independent of the FE space dimension \mathcal{N} ; thus it can be solved in a very fast way [8,7]. Moreover, we are able to compute independently of \mathcal{N} a measure error $\Delta_N(\boldsymbol{\mu})$ such that $\|u^{\mathcal{N}}(\boldsymbol{\mu}) - u_N^{\mathcal{N}}(\boldsymbol{\mu})\| \leq \Delta_N(\boldsymbol{\mu})$. The evaluation of this *a posteriori* error estimator is based on the residual of the RB approximation and a lower bound $\alpha_{\text{LB}}(\boldsymbol{\mu})$ of the coercivity constant of $a(\cdot, \cdot; \boldsymbol{\mu})$, see, e.g., [8,7].

Nevertheless, when the number of parameters $\boldsymbol{\mu}$ grows, the *offline* – and, if N increases too much, also the *online* – computational costs of the RB method become unaffordable if we aim at keeping a small tolerance on the error bound of the RB approximation with respect to the truth solution.¹ This makes the RB method not very well-suited for real-time evaluations in case of parameters showing a large dimensionality. However, since in several cases some parameter components have a negligible influence on the output, we can freeze them while keeping a good accuracy on the output. This happens for instance if there exists some colinearity between geometrical parameters. In order to translate this idea into practice, a combination of the RB method [8,7] and the ANOVA expansion [4,5,1,9] has been proposed by Hesthaven and Zhang [6], yielding a three-step RB–ANOVA–RB procedure. An alternative to the RB–ANOVA–RB approach is the use of *hp* reduced basis methods to create a local basis by decomposing the parameter space and create a local reduced basis [3] (see, e.g., [6,2] for recent developments).

The main improvement presented in this work deals with the possibility to avoid the computation of several sensitivity indices used in the ANOVA expansion based on the interaction of components of our parameters in the affine expansion (1) of $a(\cdot, \cdot; \boldsymbol{\mu})$. Before stating our main result, we sketch the main features of the RB–ANOVA–RB procedure.

First, we choose a tolerance $\varepsilon_c > 0$ and compute a *coarse* RB approximation of our problem, yielding $u_{N_c}^{\mathcal{N}}(\boldsymbol{\mu})$, $s_{N_c}^{\mathcal{N}}(\boldsymbol{\mu})$ $\forall \boldsymbol{\mu} \in \mathcal{D}$, being $N = N_c$ the minimum N for which $\Delta_N(\boldsymbol{\mu}) \leq \varepsilon_c$ $\forall \boldsymbol{\mu} \in \mathcal{D}$.

Then, we proceed to a sensitivity analysis of $s_{N_c}^{\mathcal{N}}(\boldsymbol{\mu})$ using the so-called ANOVA expansion. In general, given a function $f \in L^2(\mathcal{H}^p) := \{g : \mathcal{H}^p \rightarrow \mathbb{R} \mid \int_{\mathcal{H}^p} |g(x)|^2 dx < \infty\}$, its ANOVA expansion is defined as:

$$f(x) = f_0 + \sum_{\tau \subsetneq \mathcal{P}} f_{\tau}(x_{\tau}), \quad \forall x \in \mathcal{H}^p, \tag{2}$$

i.e. as a sum of functions depending on combination of its variables, being the coefficient f_0 defined as $f_0 = \int_{\mathcal{H}^p} f(x) dx$, where $\mathcal{P} := \{1, \dots, n\}$ and $\mathcal{H}^p = [0, 1]^p$. The sum is defined over all subsets $\tau \subsetneq \mathcal{P}$ and x_{τ} corresponds to the restriction of x on the components indexed by τ . By imposing this $\forall \tau \in \mathcal{P}$ and $i \in \tau$, we have that $\int_0^1 f(x_{\tau}) d\mu(x_i) = 0$ defines uniquely all terms in (2); see [9] for more details. We then define:

$$V_{\tau} = \int_{\mathcal{H}^{|\tau|}} |f_{\tau}(x_{\tau})|^2 d\mu(x_{\tau}), \quad V = \int_{\mathcal{H}^p} |f(x)|^2 d\mu(x) - f_0^2, \quad S_{\tau} = \frac{V_{\tau}}{V}. \tag{3}$$

¹ Note that this approach strongly improves the offline computing time, whereas it slightly impacts on the online evaluation (as a matter of fact, it does not change the behavior of the Kolmogorov N -width decay). In fact, the RB–ANOVA–RB procedure serves at avoiding to scan the entire high-dimensional parameter space to find the next maximum in the greedy algorithm.

The coefficients S_τ are called *sensitivity indices* and are used to determine the influence of x_τ on f . In particular, we have $\sum_{\tau \in \mathcal{P}} S_\tau = 1$. Denoting the truncated ANOVA expansion by:

$$\hat{f}(x) = f_0 + \sum_{\tau \subseteq \mathcal{P}_k} f_\tau(x_\tau), \quad \text{where } \mathcal{P}_k = \{\tau \subseteq \mathcal{P}: |\tau| \leq k\}, \quad k < n, \tag{4}$$

we define:

$$\varepsilon(f, \hat{f}) = \|f - \hat{f}\|_{L^2}^2 = 1 - \sum_{\tau \in \mathcal{P}_k} S_\tau \tag{5}$$

as the measure error between f and \hat{f} , where $\|\cdot\|_{L^2}$ denotes the norm on $L^2(\mathcal{H}^p)$. Hence, the coefficients S_τ can be used to determine the influence of the components indexed by τ on f . Note that, according to (5), we do not need to compute all the terms in the expansion (4), but only those necessary to reach a certain tolerance $\bar{\varepsilon}$ such that $\varepsilon(f, \hat{f}) \leq \bar{\varepsilon}$. Thus, following the procedure above, we compute the sensitivity indices S_τ associated with the coarse approximation $s_{N_c}^N$ of the RB output.

Then, based on the results of the ANOVA expansion, we freeze to an arbitrarily chosen value – say, $\bar{\mu}$ – those parameter components that show a very small influence. Finally, we compute a *fine* RB approximation solution $u_{N_f}^N(\bar{\mu})$ and evaluate the output $s_{N_f}^N(\bar{\mu})$ by taking into account just the most relevant components of the parameter vector, that is, by considering $\bar{\mu}_i = \bar{\mu}$ for any index i corresponding to a *frozen* component. Here N_f is the minimum N for which $\Delta_N(\bar{\mu}) \leq \varepsilon_f \forall \bar{\mu} \in \mathcal{D}$, ε_f being a very small tolerance. Even if $\varepsilon_f \ll \varepsilon_c$, the number of basis functions N_f needed to reach this tolerance shall be smaller than N_c because the parameter has now a reduced dimensionality. Note that it is possible to improve the results by freezing the components to suitably chosen values, as shown, e.g., in [4].

2. On the computation of sensitivity indices

In the algorithm introduced, above the sensitivity analysis procedure, it is the computation of S_τ coefficients that absorbs a big part of the computational costs. Moreover, these costs grow very fast with $|\tau|$. It is therefore important to be able to avoid the computation of those terms appearing in (2) which are meaningless. According to the result presented below, we can drastically reduce the necessary resources for the computation by simply analyzing the behavior of the functions Θ_q^a , $1 \leq q \leq Q$. The goal is to be able to define a priori if some of these terms are vanishing. To show our result, we need a preliminary lemma.

Lemma 2.1. *Let $\tau \subseteq \mathcal{P} = \{1, \dots, n\}$ and $i \in \tau$. Then $\nu = \tau \setminus \{i\}$ is the only set such that $\nu \subset \tau$, $|\nu| = |\tau| - 1$ and $i \notin \nu$.*

Proof. Each subset $\theta \subseteq \tau$ such that $|\theta| = |\tau| - 1$ and $\theta \subset \tau$ is of the form $\theta = \tau \setminus \{j\}$ for $j \in \tau$. Since i appears only once in τ , it results that ν is the only set with these properties.

In what follows, we say that Θ_q^a is independent of μ_k , if $\Theta_q^a(\mu^1) = \Theta_q^a(\mu^2), \forall \mu^1, \mu^2 \in \mathcal{D}$, such that $\mu_j^1 = \mu_j^2, j \in \mathcal{P} \setminus \{k\}$.

Proposition 2.2. *Let $\tau \subset \mathcal{P}$ and assume that for every $q = 1, \dots, Q$, there exists $i_q \in \tau$ such that Θ_q^a is independent of μ_{i_q} . Then $S_\tau = 0$.*

Proof. First, note that since we are in the compliant case [8], it is possible to write $s(\mu) = f(u(\mu)) = \sum_{q=1}^Q \alpha_q \Theta_q^a(\mu)$, with $\alpha_q = a^q(u, u)$; note that everything keeps holding if f depends on μ , thanks to the compliance assumption. Hence, by linearity of the integral:

$$s_\tau(\mu_\tau) = \sum_{q=1}^Q \alpha_q \left[\int_{\mathcal{H}^{n-|\tau|}} \Theta_q^a(\mu) d\mu(\mu_{\mathcal{P} \setminus \tau}) - \sum_{\nu \subsetneq \tau} \Theta_{q,\nu}^a(\mu_\nu) - \Theta_{q,0}^a \right] := \sum_{q=1}^Q \alpha_q \Theta_{q,\tau}(\mu_\tau),$$

where $\Theta_{q,\tau}^a$ is the term corresponding to the subscript τ in the ANOVA expansion of Θ_q^a . If we show that $\Theta_{q,\tau} = 0$ for $q = 1, \dots, Q$, then the result follows since $S_\tau = \frac{1}{V} \int_{\mathcal{H}^{|\tau|}} s_\tau^2(\mu_\tau) d\mu(\mu_\tau)$. For $q \in \{1, \dots, Q\}$, we denote $\nu_q = \tau \setminus \{i_q\}$. We show that $\Theta_{q,\tau} = 0$ by induction on $|\tau|$. If $|\tau| = 1$, then $\tau = \{i\}$, where $i_q = i$ for $q = 1, \dots, Q$ and so:

$$\Theta_{q,\tau}(\mu_\tau) = \int_{\mathcal{H}^{n-1}} \Theta_q^a(\mu) d\mu(\mu_{\mathcal{P} \setminus \{i_q\}}) - \Theta_{q,0}^a \stackrel{\text{indep.}}{=} \int_{\mathcal{H}^n} \Theta_q^a(\mu) d\mu(\mu_{\mathcal{P}}) - \Theta_{q,0}^a = 0.$$

For $|\tau| > 1$, we have:

$$\begin{aligned}
\Theta_{q,\tau}(\boldsymbol{\mu}_\tau) &= \int_{\mathcal{H}^{n-|\tau|}} \Theta_q^a(\boldsymbol{\mu}) d\boldsymbol{\mu}(\boldsymbol{\mu}_{\mathcal{P}\setminus\tau}) - \sum_{\nu \subsetneq \tau} \Theta_{q,\nu}^a(\boldsymbol{\mu}_\nu) - \Theta_{q,0}^a \\
&\stackrel{\text{indep.}}{=} \int_{\mathcal{H}^{n+1-|\tau|}} \Theta_q^a(\boldsymbol{\mu}) d\boldsymbol{\mu}(\boldsymbol{\mu}_{\mathcal{P}\setminus\nu_q}) - \sum_{\nu \subsetneq \tau} \Theta_{q,\nu}^a(\boldsymbol{\mu}_\nu) - \Theta_{q,0}^a \\
&= \Theta_{q,\nu_q}(\boldsymbol{\mu}_{\nu_q}) + \sum_{\nu \subsetneq \nu_q} \Theta_{q,\nu}^a(\boldsymbol{\mu}_\nu) + \Theta_{q,0}^a - \sum_{\nu \subsetneq \tau} \Theta_{q,\nu}^a(\boldsymbol{\mu}_\nu) - \Theta_{q,0}^a \\
&= \Theta_{q,\nu_q}(\boldsymbol{\mu}_{\nu_q}) - \sum_{\nu \subsetneq \tau, |\nu|=|\tau|-1} \Theta_{q,\nu}^a(\boldsymbol{\mu}_\nu) - \sum_{\nu \subsetneq \tau, \nu \not\subset \nu_q, |\nu|<|\tau|-1} \Theta_{q,\nu}^a(\boldsymbol{\mu}_\nu) \\
&= - \sum_{\nu \subsetneq \tau, |\nu|=|\tau|-1, \nu \neq \nu_q} \Theta_{q,\nu}^a(\boldsymbol{\mu}_\nu) - \sum_{\nu \subsetneq \tau, \nu \not\subset \nu_q, |\nu|<|\tau|-1} \Theta_{q,\nu}^a(\boldsymbol{\mu}_\nu).
\end{aligned}$$

Thanks to [Lemma 2.1](#), ν_q is the only subset of τ such that $|\nu_q| = |\tau| - 1$, $i_q \notin \nu_q$ and $\nu_q \subset \tau$. Moreover, every ν in the second sum on the last row contains i_q , since if $i_q \notin \nu$ with $|\nu| < |\tau| - 1$ and $\nu \subsetneq \tau$, then $\nu \subset \nu_q$. Hence each subset ν in the last sums contains i_q and satisfies $|\nu| \leq |\tau| - 1$. The result follows by induction. \square

This result is very useful since it allows us to avoid the computation of terms that can be costly, but do not have any influence on the ANOVA expansion. An important consideration is related to physical parameters. In this case, each Θ_q^a is in general related only to the physical properties of a region of the domain, thus it depends only on some specific components of the parameter vector. Hence, thanks to the previous proposition, we do not need to compute those terms depending on two or more variables in the ANOVA expansion. Concerning the case of geometrical parameters, it is sometimes possible to fulfill the hypothesis of [Proposition 2.2](#) by considering a *clever* subdivision of our domain.

The result we have shown allows us to improve the three-step method RB-ANOVA-RB [\[6\]](#), by decreasing the very demanding *offline* computational costs entailed by the ANOVA expansion. Even if the original method enables us to speed up the online computation – in order to make it suitable in a real-time context – no efforts have been produced so far to speed up the offline stage of this procedure. In this sense, our result makes the RB-ANOVA-RB method more attractive and computationally feasible.

3. Numerical results

The case treated here is the so-called *thermal block* problem [\[8\]](#). We consider heat conduction in a square domain $\Omega = [0, 2]^2$, given by the superposition of four horizontal layers \mathcal{R}_q , $q = 1, \dots, 4$ of equal sizes, each one corresponding to a different thermal conductivity. We impose homogeneous Dirichlet conditions on the top wall Γ_{top} of the square, homogeneous Neumann conditions on both sides and unitary Neumann conditions on the bottom wall Γ_{bot} . Our output of interest is the average temperature on Γ_{bot} , so that:

$$a(u, v; \boldsymbol{\mu}) = \int_{\Omega} v(\boldsymbol{\mu}) \nabla u \cdot \nabla v = \sum_{q=1}^4 \Theta_q^q(\boldsymbol{\mu}) \int_{\mathcal{R}_q} \nabla u \cdot \nabla v, \quad f(v) = \int_{\Gamma_{\text{bot}}} v, \quad s(\boldsymbol{\mu}) = \int_{\Gamma_{\text{bot}}} u(\boldsymbol{\mu}),$$

where $\nu(\boldsymbol{\mu})$ is the thermal conductivity. Our parameter domain is $\mathcal{D} = [0, 1]^4$. Each function Θ_q^q corresponds to a region (the bottom region being given by $q = 1$) of our computational domain and is defined as follows:

$$\Theta_q^1(\boldsymbol{\mu}) = 100^2 \mu_{q-1}, \quad \Theta_q^q(\boldsymbol{\mu}) = 1.1^2 \mu_{q-1}, \quad q = 2, 3, 4.$$

According to [Proposition 2.2](#), we do not need to compute S_τ for $|\tau| \geq 2$, because every Θ_q^q depends only on one variable. Applying the procedure described in [Section 1](#), we carried out a first RB approximation. In this case, the FE approximation dimension is given by $\mathcal{N} = 3120$ and the coarse RB by $N_c = 4$. We then compute the ANOVA expansion of our output of interest, which yields $S_1 = 9.9997 \cdot 10^{-1}$, $S_2 = S_3 = S_4 = 7.1416 \cdot 10^{-6}$, whereas we find that all other computed terms are equal to zero – this is consistent with our proposition. In view of the results obtained, we freeze the parameter components $q = 2, 3, 4$ to $\bar{\mu}_q = 0.5$, i.e. $\bar{\boldsymbol{\mu}} = (\mu_1, 0.5, 0.5, 0.5)$. The fine RB approximation obtained with this restrained model is built on the same FE spaces, but requires only an $N_f = 2$ basis. Note that the RB dimension has been divided by a factor 2. We now compute the relative error between the truth output and the fine RB approximation, by evaluating:

$$\text{err}_{\text{rel}}(\boldsymbol{\mu}^i) = \frac{|s^{\mathcal{N}}(\boldsymbol{\mu}^i) - s_{N_{\text{fine}}}^{\mathcal{N}}(\bar{\boldsymbol{\mu}}^i)|}{|s^{\mathcal{N}}(\boldsymbol{\mu}^i)|}, \quad \text{err}_{\text{max}} = \max_{i=1, \dots, 1000} \text{err}_{\text{rel}}(\boldsymbol{\mu}^i), \quad \text{err}_{\text{av}} = \frac{\sum_{i=1}^{1000} \text{err}_{\text{rel}}(\boldsymbol{\mu}^i)}{1000},$$

for 1000 parameter values randomly sampled on \mathcal{D} , being $\bar{\boldsymbol{\mu}}^i = (\mu_1^i, 0.5, 0.5, 0.5)$. We obtain $\text{err}_{\text{max}} = 8.71 \cdot 10^{-2}$ and $\text{err}_{\text{av}} = 1.53 \cdot 10^{-2}$. We remark that, by evaluating the relative error between the coarse RB output $s_{N_{\text{coarse}}}^{\mathcal{N}}(\boldsymbol{\mu}^i)$ and the fine

RB approximation, we would obtain $err_{\max} = 8.18 \cdot 10^{-2}$ and $err_{\text{av}} = 1.59 \cdot 10^{-2}$, thus making the difference between the coarse and the fine RB approximation of the output rather negligible.

Turning now to the computation costs, we obtain that terms depending on one component (non-zero ones) entail 9.898% of the whole CPU time to be evaluated. Due to our proposition, we have then been able to save about 90% of the necessary resources to compute the ANOVA expansion. Moreover, the sensitivity analysis entails in this case 99% of the whole computation time.

References

- [1] Y. Cao, Z. Chen, M. Gunzburger, ANOVA expansions and efficient sampling methods for parameter dependent nonlinear PDEs, *Int. J. Numer. Anal. Model.* 6 (2009) 256–273.
- [2] J.L. Eftang, B. Stamm, Parameter multi-domain hp empirical interpolation, *Int. J. Numer. Methods Eng.* 90 (4) (2012) 412–428.
- [3] J.L. Eftang, A.T. Patera, E.M. Rønquist, An “hp” certified reduced basis method for parametrized elliptic partial differential equations, *SIAM J. Sci. Comput.* 32 (6) (2010) 3170–3200.
- [4] Z. Gao, J.S. Hesthaven, On ANOVA expansions and strategies for choosing the anchor point, *Appl. Math. Comput.* 217 (7) (2010) 3274–3285.
- [5] M. Gunzburger, A. Labovsky, An efficient and accurate numerical method for high-dimensional stochastic PDEs, *SIAM J. Sci. Comput.* (2013), submitted for publication.
- [6] J.S. Hesthaven, S. Zhang, On the use of ANOVA expansions in reduced basis methods for high-dimensional parametric partial differential equations, Technical Report 2011-31, Scientific Computing Group, Brown University, Providence, RI, USA, 2011, submitted for publication.
- [7] A. Quarteroni, G. Rozza, A. Manzoni, Certified reduced basis approximation for parametrized partial differential equations in industrial applications, *J. Math. Ind.* 1 (3) (2011).
- [8] G. Rozza, D.B.P. Huynh, A.T. Patera, Reduced basis approximation and a posteriori error estimation for affinely parametrized elliptic coercive partial differential equations, *Arch. Comput. Methods Eng.* 15 (2008) 229–275.
- [9] I.M. Sobol, Theorems and examples on high dimensional model representation, *Reliab. Eng. Syst. Saf.* 79 (2003) 187–193.