4 Model-Order Reduction for Linear PPDEs

In the last chapter we have investigated the (reduced) optimal control problem under consideration, i.e. $\mathbf{O}(\mu)$. We have seen that for the proposed solution algorithm, i.e. Algorithm 3.4.1, we have to perform many evaluations of the reduced cost functional $\mathbf{J}(\mu)$. Moreover, for each evaluation of the latter we need to solve the state equation $\mathrm{S}(\mu)$, which is a linear elliptic parametrized partial differential equation (PPDE). Hence, these evaluations are quite expensive.

In this chapter we investigate the possibility of replacing $\mathrm{S}(\mu)$ by a reduced order model, such that the evaluation of $\mathbf{J}(\mu)$ can be accelerated significantly by deriving an appropriate approximation. The method of choice is the reduced-basis method (RBM), as it is designed for "*rapid*, [*repeated*] and *reliable* evaluation of *input-output relationships* in which the output is expressed as a functional of a field variable that is the solution of an *input-parametrized* partial differential equation" (c.p. [39]).

As we aim in applying RBM, we start by presenting the main idea in the next section. Moreover, we give a short outline of the remaining parts of this chapter. Finally, note that although we aim in using RBM for accelerating the evaluation of $\mathbf{J}(\mu)$, the theory in this chapter is conceived as general as possible and our particular problem mainly serves for the numerical experiments.

Finally, we want to mention that in [42] we already applied RBM for our particular problem at hand, but were lacking a-posteriori error estimators to that point what we want to make up for in this chapter as well as in Chapter 6.

4.1 Main Idea & Outline

Before going into detail, we note that [39] contains an extensive overview of both the historical background of reduced-basis methods and the current state of the art. Moreover, the mentioned article provides many examples for its application within two particular contexts – the real-time context and the many-query context. These contexts are of particular interest, as the methodology is not only motivated by but also optimized for application within. To be more precise, in both contexts many repeated evaluations of input-output relationships have to be performed. Moreover, decoupling the computation of an approximation to the field and the output of interest, respectively, into an expensive *offline* and an inexpensive *online* phase, allows for obtaining these approximations rapidly. Finally, the availability of a-posteriori error estimators does not only guarantee the reliability of these approximations, but the efficiency of the methodology itself, too. This explains the three keywords characterizing reduced-basis methods, i.e. repeated, rapid and reliable, as already emphasized in the introduction to this chapter. Although in both

contexts many repeated evaluations have to be performed, there is one main difference: In the real-time context the complexity of the offline phase can be neglected by definition, whereas in the many-query context it may not. Hence, for the application in the latter context the number of necessary evaluations should be sufficiently large in order to be profitable. One should recall, that we aim in applying the reduced-basis method to speed up solving $\mathbf{O}(\mu)$. Consequently, we are in the many-query context and the number of evaluations roughly equals the number of iterations (times a constant) of the utilized optimization procedure, i.e. Algorithm 3.4.1.

Now, to formalize the input-output relationship, let X^e, $H_0^1(\Omega) \subset X^e \subset H^1(\Omega)$, denote the *exact* function space and $\mu = (\mu_1, \dots, \mu_d)^T \in \mathcal{D}_{\text{ad}}$ the parameter vector (configuration). Moreover, let the linear elliptic PPDE be defined by the bi-linear form $a(\cdot, \cdot; \mu)\colon X^e \times X^e \to \mathbb{R}$ and the linear form $f(\cdot; \mu)\colon X^e \to \mathbb{R}$ and let $o(\cdot; \mu)\colon X^e \to \mathbb{R}$ denote some output functional, additionally. Then, we define:

Definition 4.1.1. *Let $X \subset X^e$.*

1. *For all $\mu \in \mathcal{D}_{\text{ad}}$ finding $u(\mu) \in X$, such that*

$$a(u(\mu), v; \mu) = f(v; \mu), \quad v \in X, \tag{$P(\mu)$}$$

   is called the Primal Problem.

2. *For all $\mu \in \mathcal{D}_{\text{ad}}$*

$$\mathbf{s}(\mu) := o(u(\mu); \mu)$$

   is called the Output of Interest*, where $u(\mu)$ solves $P(\mu)$.*

For solving $P(\mu)$ numerically, we have to use a discrete function space $X^{\mathcal{N}} \subset X^e$ for X, e.g. a finite element space or a spectral element space, specified by $\mathcal{N}$ degrees of freedom. In the context of RBM $X^{\mathcal{N}}$ is called *truth approximation* space, as it is assumed to be rich enough, such that $u^e(\mu)$, i.e. the solution of $P(\mu)$ using $X = X^e$, and $u^{\mathcal{N}}(\mu)$, i.e. the solution of $P(\mu)$ using $X = X^{\mathcal{N}}$, are indistinguishable. Hence, by this assumption we can restrict all our investigations to this *truth approximation* space, i.e. we always use $X := X^{\mathcal{N}}$ while referring to $P(\mu)$. Next, in order to ensure the well-posedness of $P(\mu)$, we have to suppose the following assumption to hold true, where for the inner product in X, i.e. $(w, v)_X$, $w, v \in X$, we consider the usual $H^1(\Omega)$ inner product (recall Definition 2.2.2), and by $\|\cdot\|_X := \sqrt{(\cdot, \cdot)_X}$ we denote its induced norm:

Assumption 4.1.2. *For all $\mu \in \mathcal{D}_{\text{ad}}$*

1. *$a(\cdot, \cdot; \mu)$ is continuous, i.e.*

$$\sup_{w \in X} \sup_{v \in X} \frac{a(w, v; \mu)}{\|w\|_X \|v\|_X} =: \gamma(\mu) < \infty, \tag{4.1.1}$$

2. $a(\cdot,\cdot;\mu)$ *satisfies the inf-sup stability condition, i.e.*

$$\inf_{w\in X}\sup_{v\in X}\frac{a(w,v;\mu)}{\|w\|_X\|v\|_X} =: \beta(\mu) > 0, \tag{4.1.2}$$

3. $f(\cdot;\mu)\in X'$*, i.e.* $f(\cdot;\mu)$ *is linear and bounded.*

This assumption ensures unique solvability (c.p. [2]) of $P(\mu)$ for all $\mu\in\mathcal{D}_{\text{ad}}$. For the successful application of RBM, however, we need to assume further, that $u(\mu)$ evolves on a low dimensional manifold $\mathcal{M} := \{u(\mu), \mu\in\mathcal{D}_{\text{ad}}\}\subset X$ rather than being an arbitrary member of X, which is visualized schematically in Figure 4.1. Note, that the same assumption is also crucial for other techniques of model order reduction, e.g. the proper orthogonal decomposition (POD), which we detail in Section 4.4.1.

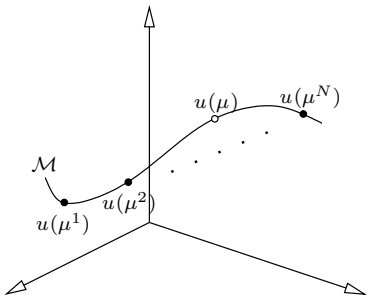

Figure 4.1: Manifold $\mathcal{M}$ and snapshots $u(\mu^1),\dots,u(\mu^N)$.

If this assumption holds true, a reduced-order model can be obtained by Galerkin projection of $P(\mu)$ onto a small subspace of X spanned by solutions of $P(\mu)$. To be more precise, the reduced-order approximation $u^N(\mu)\in X^N := \text{span}\{\xi^1,\dots,\xi^N\}\subset X$, to $u(\mu)\in X$ satisfies

$$a(u^N(\mu),v;\mu) = f(v;\mu), \quad v\in X^N. \tag{4.1.3}$$

Obviously, if $N \ll \mathcal{N}$, the complexity for solving (4.1.3) is significantly reduced compared to the complexity for solving $P(\mu)$. On the other hand, the complexity for assembling (4.1.3) still depends on $\mathcal{N}$, such that the solution $u^N(\mu)$ can be obtained more rapidly, but not rapidly enough for the application within the two contexts described above. Hence, despite the projection of $P(\mu)$ onto a low dimensional subspace of X, the second important ingredient is the full decoupling into an offline-/online phase, such that the latter becomes independent of $\mathcal{N}$. However, the possibility for performing this decoupling directly relates to the parametric dependency of the forms a and f on the parameter μ. The usual assumption is, that these forms allow for an affine

decomposition, i.e. for $w, v \in X$ and $\mu \in \mathcal{D}_{\text{ad}}$ we have

$$a(w, v; \mu) = \sum_{m=1}^{M_a} \vartheta_m^a(\mu)\, a^m(w, v), \tag{4.1.4a}$$

$$f(v; \mu) = \sum_{m=1}^{M_f} \vartheta_m^f(\mu)\, f^m(v), \tag{4.1.4b}$$

where the forms a^m, $1 \leq m \leq M_a$, and f^m, $1 \leq m \leq M_f$, are independent of the parameter $\mu \in \mathcal{D}_{\text{ad}}$ and the corresponding coefficients ϑ_m only depend on the latter.

At this point we note, that if this assumption holds true and if M_a and M_f are of moderate size, this case is very well known (c.p. [39] and references therein for linear outputs, i.e. $o(v; \mu) = \ell(v; \mu)$, $\ell(\cdot; \mu) \in X'$ linear and bounded, and [22, 23] for quadratic outputs, i.e. $o(v; \mu) = b(v, v; \mu)$, $b(\cdot, \cdot; \mu) \colon X \times X \to \mathbb{R}$ bi-linear and continuous) and we could basically stop our investigations.

However, if it does not hold true we have to replace these forms by affine approximations, say a^ϵ and f^ϵ, in a first step and build the reduced-order model (ROM) based on these approximations, afterwards, what is formalized in Chapter 4.2. A possible way of deriving such affine approximations in an economic way, i.e. with almost as few terms as possible, is presented in Chapter 4.3. Although this new method enables very economic affine approximations, for small approximation tolerances ϵ the number of terms M_a^ϵ needed for approximating a up to ϵ is still large in our particular application, which leads us to the next issue: a-posteriori error estimation.

Generally speaking, the availability of effective and rigorous[1] a-posteriori error estimators is key for the success of the RBM, as rigorousity implies reliability of the derived approximations on the one hand and effectivity implies efficiency of the method itself on the other hand. In other words, these two properties ensure that we can choose N, the number of used basis functions for X^N, as large as necessary, but not unnecessarily large at the same time. Following this argument for the approximation tolerance ϵ from above, it is obviously desirable to (be able to) choose it as small as necessary, but not unnecessarily small, too. The derivation of an a-posteriori error estimator that takes both approximation errors (affine approximation of non-affine forms as well as Galerkin projection onto X^N) into account and remains rigorous and effective is subject of Section 4.5. The benefit of such an estimator is, that we can extent the standard greedy procedure, which is usually used to generate "good" reduced-basis functions for X^N and is presented along with the proper orthogonal decomposition in Section 4.4, such that the approximation tolerance ϵ can be adjusted adaptively, too, what is detailed in Section 4.5.1.

Before we turn to the a-posteriori error estimation for the output of interest, we anticipate that the computation of the a-posteriori error estimator for the field variable basically consists of two parts, namely the computation of a lower bound for the stability factor $\beta(\mu)$, which is subject of Section 4.7, as well as the computation of $\|\cdot\|_{X'}$ for some residuals. As this is the most expensive part, we investigate in Section 4.5.2 in which situations it can pay off to replace

[1] The terms "effective" and "rigorous" are defined in the further course of this chapter.

it by the computation of some localized quantities.

In Section 4.6 we turn to the output of interest and derive rigorous error estimators for both linear and quadratic outputs. Moreover, we again carefully analyze which approximation accuracy is needed at which point in order to achieve a prescribed approximation tolerance for the output while the overall complexity is not unnecessarily large at the same time. The latter is all the more important as we do not want to use any kind of postprocessing step.

Finally, we present many numerical experiments underlining the impact of the proposed way of action for our motivating application.

Remark 4.1.3. *In this chapter we focus on the application of RBM for solving the state equation* $\mathrm{S}(\mu)$ *and the evaluation of the reduced cost functional* $\mathbf{J}(\mu)$*. However, for accelerating Algorithm 3.4.1 we also have to apply RBM for solving the adjoint equation* $\mathrm{A}(\mu)$ *as well as the evaluation of the reduced gradient* $\nabla\mathbf{J}(\mu)$ *(c.p.* (3.3.12)*). Hence, in Definition 4.1.1 we called* $P(\mu)$ *the Primal Problem, which already indicates that we are going to apply the results from this chapter to both the state and the adjoint equation later on, i.e. in Chapter 5. Moreover, for the output approximation investigated in Section 4.6 we again have to take advantage of some adjoint problems, which are called Dual Problems in this context in order to avoid confusion with the adjoint equation* $\mathrm{A}(\mu)$*.*

4.2 Reduced-Order Model (ROM)

In this section we introduce the reduced-order model. As already mentioned in Section 4.1, if the forms a and/or f are non-affine separable w.r.t. the parameter $\mu \in \mathcal{D}_{\mathrm{ad}}$ or if the corresponding numbers of summands in (4.1.4), i.e. M_a and/or M_f, are too large, we need to approximate $P(\mu)$ in a first step, what is formalized by the definition to follow.

Definition 4.2.1. *Let* $a^\epsilon(\cdot,\cdot;\mu)\colon X \times X \to \mathbb{R}$ *and* $f^\epsilon(\cdot;\mu) \in X'$*, such that for all* $\mu \in \mathcal{D}_{\mathrm{ad}}$

$$a(w,v;\mu) - a^\epsilon(w,v;\mu) \leq \epsilon \|w\|_X \|v\|_X, \quad w,v \in X, \tag{4.2.1a}$$

$$f(v;\mu) - f^\epsilon(v;\mu) \leq \epsilon \|v\|_X, \qquad v \in X. \tag{4.2.1b}$$

Moreover, let $a^\epsilon(\cdot,\cdot;\mu)$ *and* $f^\epsilon(\cdot,\cdot;\mu)$ *be affine separable w.r.t. the parameter* $\mu \in \mathcal{D}_{\mathrm{ad}}$*, i.e. for* $w,v \in X$ *and* $\mu \in \mathcal{D}_{\mathrm{ad}}$ *we have*

$$a^\epsilon(w,v;\mu) = \sum_{m=1}^{M_a^\epsilon} \vartheta_m^a(\mu)\, a^m(w,v), \tag{4.2.2a}$$

$$f^\epsilon(v;\mu) = \sum_{m=1}^{M_f^\epsilon} \vartheta_m^f(\mu)\, f^m(v), \tag{4.2.2b}$$

where $a^m\colon X \times X \to \mathbb{R}$*,* $\vartheta_m^a\colon \mathcal{D}_{\mathrm{ad}} \to \mathbb{R}$*,* $1 \leq m \leq M_a^\epsilon$*, and* $f^m \in X'$*,* $\vartheta_m^f\colon \mathcal{D}_{\mathrm{ad}} \to \mathbb{R}$*,* $1 \leq m \leq M_f^\epsilon$*, and for given* $\epsilon > 0$ *the number of summands in* (4.2.2)*, i.e.* $M_a^\epsilon, M_f^\epsilon \in \mathbb{N}$*, is chosen as small as*

possible, such that (4.2.1) *holds true. Then, finding* $\hat{u}(\mu) \in X$*, such that*

$$a^\epsilon(\hat{u}(\mu), v; \mu) = f^\epsilon(v; \mu), \quad v \in X, \tag{$\hat{P}(\mu)$}$$

is called Approximated Primal Problem.

Note, that $\hat{P}(\mu)$ is set up for theoretical purposes, only. Moreover, we want to emphasize that regardless of the way μ enters the forms a and f as well as how the approximations a^ϵ and f^ϵ are obtained, the a-posteriori error estimator to be derived in Section 4.5 remains valid as it only relies on (4.2.1) and (4.2.2). One particular way of obtaining these approximations is presented in Section 4.3, where μ is assumed to enter the coefficients of the PDE. Another possibility is to consider a truncated Karhunen-Loève expansion in case of stochasticity in the coefficients of the PDE.

Having $\hat{P}(\mu)$ at hand (although not by construction but at least by definition), we can introduce the reduced-order model by Galerkin projection of $\hat{P}(\mu)$ onto a low-dimensional subspace of X as already announced in Section 4.1.

Definition 4.2.2. *Let* a^ϵ *and* f^ϵ *be as in Definition 4.2.1 and* $X^N := \operatorname{span}\{\xi^1, \dots, \xi^N\} \subset X$*. Then, finding* $\hat{u}^N(\mu) \in X^N$*, such that*

$$a^\epsilon(\hat{u}^N(\mu), v; \mu) = f^\epsilon(v; \mu), \quad v \in X^N, \tag{$\hat{P}^N(\mu)$}$$

is called Primal Reduced-Order Model.

At this point it is not clear how to choose the basis functions ξ^n, $1 \leq n \leq N$, in order to obtain a good reduced-basis approximation space X^N. This issue, however, will be dealt with in Section 4.4, where we describe POD as well as a prototype of the standard greedy algorithm usually used in the context of RBM.

Finally, we consider the *offline-/online decomposition*. By the bi-linearity of $a^\epsilon(\cdot, \cdot; \mu)$ and the linearity of $f^\epsilon(\cdot; \mu)$ the coefficients of the solution $\hat{u}^N(\mu) = \sum_{n=1}^N \hat{u}_n^N(\mu)\xi^n \in X^N$ of $\hat{P}^N(\mu)$, i.e. $\underline{\hat{u}}^N(\mu) := \big(\hat{u}_n^N(\mu)\big)_{n=1}^N \in \mathbb{R}^N$, can be computed by solving

$$A^\epsilon(\mu)\underline{\hat{u}}^N(\mu) = F^\epsilon(\mu), \tag{4.2.3}$$

where

$$A^\epsilon(\mu) := \Big(a^\epsilon(\xi^{n'}, \xi^n; \mu)\Big)_{1 \leq n, n' \leq N} \in \mathbb{R}^{N \times N},$$
$$F^\epsilon(\mu) := \Big(f^\epsilon(\xi^n; \mu)\Big)_{1 \leq n \leq N} \in \mathbb{R}^N,$$

which is of complexity $\mathcal{O}(N^3)$ as $A^\epsilon(\mu)$ usually is a dense matrix. Moreover, from (4.2.2) we

infer that in the *online phase* we can assemble (4.2.3) in $\mathcal{O}(M_f^\epsilon N + M_a^\epsilon N^2)$ as

$$A^\epsilon(\mu) = \sum_{m=1}^{M_a^\epsilon} \vartheta_m^a(\mu)\, A^m,$$
$$F^\epsilon(\mu) = \sum_{m=1}^{M_f^\epsilon} \vartheta_m^f(\mu)\, F^m,$$

where

$$A^m := \Big(a^m(\xi^{n'}, \xi^n)\Big)_{1\le n,n'\le N} \in \mathbb{R}^{N\times N}, \quad 1 \le m \le M_a^\epsilon,$$
$$F^m := \Big(f^m(\xi^n)\Big)_{1\le n\le N} \in \mathbb{R}^N, \qquad\qquad 1 \le m \le M_f^\epsilon,$$

can be precomputed in the *offline phase* in $\mathcal{O}((M_f^\epsilon N + M_a^\epsilon N^2)\mathcal{N})$. Hence, the complexity of the online phase is $\mathcal{N}$-independent and $\hat{\underline{u}}^N(\mu)$ can be computed rapidly.

4.3 Non-Affine Function Approximation

In this section we focus on obtaining approximations for the forms a and f used in Definition 4.2.1, i.e. a^ϵ and f^ϵ satisfying (4.2.1) and (4.2.2), where we assume that the parametric dependency of a and f on μ takes the form

$$a(w, v; \mu) = \tilde{a}(w, v; \underline{\underline{A}}(\cdot; \mu)),$$
$$f(v; \mu) = \tilde{f}(v; \underline{F}(\cdot; \mu)),$$

where $\tilde{a}(\cdot, \cdot; \underline{\underline{A}})\colon X \times X \to \mathbb{R}$ and $\tilde{f}(\cdot; \underline{F}) \in X'$ are linear in their last argument, i.e. $\underline{\underline{A}}$ and $\underline{F}$, respectively, too. Moreover, $\underline{\underline{A}}(x; \mu)$ and $\underline{F}(x; \mu)$ may consist of several components, where each component is in $L_\infty(\Omega)$ for fixed $\mu \in \mathcal{D}_{\text{ad}}$ and sufficiently smooth in $\mathcal{D}_{\text{ad}}$ for fixed $x \in \Omega$. We will subsequently give a particular example. Now, we assume that we can derive constants $\gamma_{\tilde{a}}, \gamma_{\tilde{f}} < \infty$, such that for $w, v \in X$

$$\tilde{a}(w, v; \underline{\underline{A}}) \le \gamma_{\tilde{a}} \|\underline{\underline{A}}\|_{L_\infty(\Omega)} \|w\|_X \|v\|_X,$$
$$\tilde{f}(v; \underline{F}) \le \gamma_{\tilde{f}} \|\underline{F}\|_{L_\infty(\Omega)} \|v\|_X,$$

where e.g. $\|\underline{\underline{A}}\|_{L_\infty(\Omega)}$ denotes the maximum value of $\|\cdot\|_{L_\infty(\Omega)}$ applied to each component of $\underline{\underline{A}}$, see also (4.3.17). Then, the problem of deriving the desired approximations a^ϵ and f^ϵ reduces to (due to the linearity of $\tilde{a}$ in $\underline{\underline{A}}$ and $\tilde{f}$ in $\underline{F}$) deriving approximations for $\underline{\underline{A}}(x; \mu)$ and $\underline{F}(x; \mu)$