Reduced Basis Approximation and a Posteriori Error Estimation for Affinely Parametrized Elliptic Coercive Partial Differential Equations

Application to Transport and Continuum Mechanics

G. Rozza · D.B.P. Huynh · A.T. Patera

Abstract In this paper we consider (hierarchical, Lagrange) reduced basis approximation and *a posteriori* error estimation for linear functional outputs of affinely parametrized elliptic coercive partial differential equations. The essential ingredients are (primal-dual) Galerkin projection onto a low-dimensional space associated with a smooth “parametric manifold”—dimension reduction; efficient and effective greedy sampling methods for identification of optimal and numerically stable approximations—rapid convergence; *a posteriori* error estimation procedures—rigorous and sharp bounds for the linear-functional outputs of interest; and Offline-Online computational decomposition strategies—minimum *marginal cost* for high performance in the real-time/embedded (e.g., parameter-estimation, con-

trol) and many-query (e.g., design optimization, multi-model/scale) contexts. We present illustrative results for heat conduction and convection-diffusion, inviscid flow, and linear elasticity; outputs include transport rates, added mass, and stress intensity factors.

1 Introduction and Motivation

In this work we describe reduced basis (RB) approximation and *a posteriori* error estimation methods for rapid 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 (PDE). In this particular paper we shall focus on linear output functionals and affinely parametrized linear elliptic coercive PDEs; however the methodology is much more generally applicable, as we discuss in Sect. 2.

We emphasize applications in transport and mechanics: unsteady and steady heat and mass transfer; acoustics; and solid and fluid mechanics. (Of course we do not preclude other domains of inquiry within engineering (e.g., electromagnetics) or even more broadly within the quantitative disciplines (e.g., finance).) The *input-parameter* vector typically characterizes the geometric configuration, the physical properties, and the boundary conditions and sources. The *outputs of interest* might be the maximum system temperature, an added mass coefficient, a crack stress intensity factor, an effective constitutive property, an acoustic waveguide transmission loss, or a channel flowrate or pressure drop. Finally, the *field variables* that connect the input parameters to the outputs can represent a distribution function, temperature or concentration, displacement, pressure, or velocity.

The methodology we describe in this paper is motivated by, optimized for, and applied within two particular con-

This work was supported by DARPA/AFOSR Grants FA9550-05-1-0114 and FA-9550-07-1-0425, the Singapore-MIT Alliance, the Pappalardo MIT Mechanical Engineering Graduate Monograph Fund, and the Progetto Roberto Rocca Politecnico di Milano-MIT. We acknowledge many helpful discussions with Professor Yvon Maday of University of Paris 6 and Luca Dedé of MOX-Politecnico di Milano.

G. Rozza
Mechanical Engineering Department, Massachusetts Institute of Technology, Room 3-264, 77 Mass Avenue, Cambridge, MA 02142-4307, USA
e-mail: rozza@mit.edu

D.B.P. Huynh
Singapore-MIT Alliance, E4-04-10, National University of Singapore, 4 Eng. Drive, Singapore, 117576, Singapore
e-mail: baophuong@nus.edu.sg

A.T. Patera
Massachusetts Institute of Technology, Room 3-266, 77 Mass Avenue, Cambridge, MA 02142-4307, USA
e-mail: patera@mit.edu

texts: the *real-time context* (e.g., parameter-estimation [54, 96, 154] or control [124]); and the *many-query context* (e.g., design optimization [107] or multi-model/scale simulation [26, 49]). Both these contexts are crucial to computational engineering and to more widespread adoption and application of numerical methods for PDEs in engineering practice and education.

We first illustrate the *real-time context*: we can also characterize this context as “deployed” or “in the field” or “embedded.” As an example of a real-time—and in fact, often also many-query—application, we consider a crack in a critical structural component such as a composite-reinforced concrete support (or an aircraft engine). We first pursue Non-Destructive Evaluation (NDE) parameter estimation procedures [17, 96, 154]—say by vibration or thermal transient analysis—to determine the location and configuration of the delamination crack in the support. We then evaluate stress intensity factors to determine the critical load for fracture or the anticipated crack growth due to fatigue. Finally we modify “on site” the installation or subsequent mission profile to prolong life. Safety and economics require rapid and reliable response in the field.

We next illustrate the many-query context. As an example of a general family of many-query applications, we cite multiscale (temporal, spatial) or multiphysics models in which behavior at a larger scale must “invoke” many spatial or temporal realizations of parametrized behavior at a smaller scale. Particular cases (to which RB methods have been applied) include stress intensity factor evaluation [5, 62] within a crack fatigue growth model [63]; calculation of spatially varying cell properties [26, 75] within homogenization theory [25] predictions for macroscale composite properties; assembly and interaction of many similar building blocks [79] in large (e.g., cardio-vascular [137]) biological networks; or molecular dynamics computations based on quantum-derived energies/forces [36]. In all these cases, the number of input-output evaluations is often measured in the tens of thousands.

Both the real-time and many-query contexts present a significant and often unsurmountable challenge to “classical” numerical techniques such as the finite element (FE) method. These contexts are often much better served by the reduced basis approximations and associated *a posteriori* error estimation techniques described in this work. We note, however, that the RB methods we describe do not replace, but rather build upon and are measured (as regards accuracy) relative to, a finite element model [22, 41, 126, 147, 158]: the reduced basis approximates not the exact solution but rather a “given” finite element discretization of (typically) very large dimension N . In short, we pursue an algorithmic collaboration rather than an algorithmic competition.

2 Historical Perspective and Background

The development of the reduced basis (RB) method can perhaps be viewed as a response to the considerations and imperatives described above. In particular, the parametric real-time and many-query contexts represent not only computational challenges, but also computational opportunities. We identify two key opportunities that can be gainfully exploited:

Opportunity I. In the *parametric* setting, we restrict our attention to a typically *smooth* and rather *low-dimensional* parametrically induced manifold: the set of fields engendered as the input varies over the parameter domain; in the case of single parameter, the parametrically induced manifold is a one-dimensional filament within the infinite dimensional space which characterizes *general* solutions to the PDE. Clearly, generic approximation spaces are unnecessarily rich and hence unnecessarily expensive within the parametric framework.

Opportunity II. In the *real-time* or *many-query* contexts, in which the premium is on *marginal cost* (or equivalently asymptotic average cost) per input-output evaluation, we can accept greatly increased pre-processing or “Offline” cost—not tolerable for a single or few evaluations—in exchange for greatly decreased “Online” (or deployed) cost for each new/additional input-output evaluation. Clearly, resource allocation typical for “single-query” investigations will be far from optimal for many-query and real-time exercises.

We shall review the development of RB methods in terms of these two opportunities.

Opportunity I

Reduced Basis discretization is, in brief, (Galerkin) projection on an N -dimensional approximation space that focuses (typically through Taylor expansions or Lagrange “snapshots”) on the parametrically induced manifold identified in Opportunity I. Initial work grew out of two related streams of inquiry: from the need for more effective, and perhaps also more interactive, many-query design evaluation—[48] considers linear structural examples; and from the need for more efficient parameter continuation methods—[4, 98, 99, 101, 104, 105] consider nonlinear structural analysis problems. (Several modal analysis techniques from this era [92] are also closely related to RB notions.)

The ideas present in these early somewhat domain-specific contexts were soon extended to (i) general finite-dimensional systems as well as certain classes of PDEs (and ODEs) [19, 47, 76, 100, 106, 120, 131, 132], and (ii) a variety of different reduced basis approximation spaces—in particular Taylor and Lagrange [119] and more recently Hermite [67] expansions. The next decade(s) saw further

expansion into different applications and classes of equations, such as fluid dynamics and the incompressible Navier-Stokes equations [57, 66–69, 114].

However, in these early methods, the approximation spaces tended to be rather local and typically rather low-dimensional in parameter (often a single parameter). In part, this was due to the nature of the applications—parametric continuation. But it was also due to the absence of *a posteriori* error estimators and effective sampling procedures. It is clear that in more global, higher-dimensional parameter domains the *ad hoc* reduced basis predictions “far” from any sample points can not necessarily be trusted, and hence *a posteriori* error estimators are crucial to reliability (and ultimately, safe engineering interventions *in particular in the real-time context*). It is equally clear that in more global, higher-dimensional parameter domains simple tensor-product/factorial “designs” are not practicable, and hence sophisticated sampling strategies are crucial to convergence and computational efficiency.¹

Much current effort is thus devoted to development of (i) *a posteriori* error estimation procedures and in particular rigorous error bounds for outputs of interest [121], and (ii) effective sampling strategies in particular for higher (than one) dimensional parameter domains [32, 33, 97, 136, 153]. The *a posteriori* error bounds are of course indispensable for rigorous certification of any particular reduced basis (Online) output prediction. However, the error estimators can also play an important role in efficient and effective (greedy) sampling procedures: the inexpensive error bounds permit us first, to explore much larger subsets of the parameter domain in search of most representative or best “snapshots,” and second, to determine when we have *just enough* basis functions. Just as in the finite element context [12], the simultaneous emergence of error estimation and adaptive sampling/approximation capabilities is certainly not a coincidence.

We note here that greedy sampling methods are similar in objective to, but very different in approach from, more well-known Proper Orthogonal Decomposition (POD) methods [8, 24, 58, 73, 77, 93, 127–129, 144, 145, 157]. For reasons that we shall explore, the former are applied in the (*multi*-dimensional) parameter domain, while the latter are most often applied in the (*one*-dimensional) temporal domain. However, POD economization techniques can be, and have successfully been, applied within the parametric RB context [31, 40, 42, 49, 59, 83, 154]. (We shall conduct a brief comparison of greedy and POD approaches—computational cost and performance—in Sect. 8.1.4.)

¹ Several early papers [102–104] did indeed discuss *a posteriori* error estimation and even adaptive improvement/sampling of the RB space; however, the approach could not be efficiently or rigorously applied to partial differential equations due to the computational requirements, the residual norms employed, and the absence of any stability considerations.

Opportunity II

Early work on the reduced basis method certainly exploited Opportunity II—but not fully. In particular, and perhaps at least partially because of the difficult nonlinear nature of the initial applications, early RB approaches did not fully decouple the underlying FE approximation—of very high dimension $\mathcal{N}$ —from the subsequent reduced basis projection and evaluation—of very low dimension N . More precisely, most often the Galerkin stiffness equations for the reduced basis system were generated by direct appeal to the high-dimensional FE representation: in nuts and bolts terms, pre- and post-multiplication of the FE stiffness system by rectangular basis matrices. As a result of this expensive projection the computational savings provided by RB treatment (relative to classical FE evaluation) were typically rather modest [98, 119, 120] *despite* the very small size of the ultimate reduced basis stiffness system.

Much current work is thus devoted to *full* decoupling of the FE and RB spaces through Offline-Online procedures: the complexity of the Offline stage depends on $\mathcal{N}$ (the dimension of the FE space); the complexity of the Online stage—in which we respond to a new value of the input parameter—depends only on N (the dimension of the reduced basis space) and Q (which measures the parametric complexity of the operator and data, as defined below). In essence, in the Online stage we are guaranteed the accuracy of a high-fidelity finite element model but at the very low cost of a reduced-order model.

In the context of *affine parameter dependence*, in which the operator is expressible as the sum of Q products of parameter-dependent functions and parameter-independent operators, the Offline-Online idea is quite self-apparent and indeed has been re-invented often [15, 66, 70, 114]; however, application of the concept to *a posteriori* error estimation—note the Online complexity of both the output *and* the output error bound calculation must be independent of $\mathcal{N}$ —is more involved and more recent [64, 121, 122]. In the case of *nonaffine parameter dependence* the development of Offline-Online strategies is much less transparent, and only in the last few years have effective procedures—in effect, efficient methods for approximate reduction to affine form—been established [18, 53, 138]. Clearly, Offline-Online procedures are a crucial ingredient in the real-time and many-query contexts.

We note that historically [47] and in this paper RB methods are built upon, and measured (as regards accuracy) relative to, underlying *finite element* discretizations (or related spectral element approaches [79–82, 111]): the variational framework provides a very convenient setting for approximation and error estimation. However there are certainly many good reasons to consider alternative settings: a systematic finite volume framework for RB approximation and

a posteriori error estimation is proposed and developed in [60]. We do note that boundary and integral approximations are less amenable to RB treatment or at least Offline-Online decompositions, as the inverse operator will typically not be affine in the parameter.

3 Current Status of Reduced Basis Methods

3.1 Affinely Parametrized Elliptic Coercive Problems

In this paper we shall consider the case of linear functional outputs of affinely parametrized linear elliptic coercive partial differential equations. This class of problems—relatively simple, yet relevant to many important applications in transport (e.g., conduction and convection-diffusion) and continuum mechanics (e.g., linear elasticity)—proves a convenient expository vehicle for the methodology. We provide here a brief “table of contents” for the remainder of this review.

In Sect. 4 (compliant problems) and at the conclusion of the paper in Sect. 11 (non-compliant problems) we describe the affine linear elliptic coercive setting; in Sect. 5 we consider admissible classes of piecewise-affine geometry and coefficient parametric variation; in Sect. 6 we introduce several “working examples” which shall serve to illustrate the formulation and methodology.

In Sect. 7.1 for compliant problems and subsequently Sect. 11.2 for non-compliant problems we discuss (primal-dual [117]) RB Galerkin projection [121] and optimality; in Sect. 7.2 we describe (briefly) POD methods [8, 24, 58, 73] and (more extensively) greedy sampling procedures [32, 33, 153] for optimal space identification; in Sect. 8.1 for one parameter and Sect. 8.2 for many parameters we investigate the critical role of parametric smoothness [47, 85] in convergence theory and practice.

In Sect. 9 we present rigorous and relatively sharp *a posteriori* output error bounds [3, 23, 108] for RB approximations [121, 142]; in Sect. 10 we develop the coercivity-constant lower bounds [64] required by the *a posteriori* error estimation procedures.

In Sect. 7.1 for the output prediction and Sect. 9.4 for the output error bounds [84] we present the Offline-Online computational strategies crucial to rapid real-time/many-query (Online) response. In Sect. 8.2 we also provide a quantitative comparison between RB (Offline and Online) and FE computational performance.

Although this paper focuses on the affine linear elliptic coercive case, the reduced basis approximation and *a posteriori* error estimation methodology is much more general. Furthermore, most of the basic concepts introduced in the affine linear elliptic coercive case are equally crucial—with suitable extension—to more general equations. In the next section we briefly review the current landscape and provide references for further inquiry.

3.2 Extensions and Generalizations

First, the reduced basis approach can also be readily applied to the more general case of affine linear elliptic *non-coercive* problems.² The canonical example is the ubiquitous Helmholtz (reduced-wave equation) relevant to time-harmonic acoustics [141], elasticity [95], and electromagnetics. (We also note that a special formulation for *quadratic* output functionals [62, 141]—important in such applications as acoustics (transmission loss outputs) and linear elastic fracture theory (stress intensity factor outputs)—is perforce non-coercive.) With respect to the elements we consider in the current paper on *coercive* problems, the key new methodological challenges for *non-coercive* problems are the development of (i) *discretely* stable primal-dual RB approximations [86, 134], and (ii) efficient Offline-Online computational procedures [64, 142] for the construction of lower bounds for the (no longer *coercivity*, but rather) *inf-sup* constant [9] required by the *a posteriori* error estimators. The possibility of resonances and near-resonances can adversely affect the efficiency of both the RB approximation and the RB error bounds, often limiting the dimensionality or extent of the parameter domain.

RB-like “snapshot” ideas are also found in some of the many Reduced Order Model (ROM) approaches in the temporal domain [14, 38, 39, 89, 116, 130, 145, 155, 156]: POD sampling procedures are often invoked, and more recently greedy sampling approaches have also been considered [21]. Thus, *combination* of “parameter + time” approaches—essentially the marriage of ROM in time with RB in parameter, sometimes referred to as PROM (Parametric ROM)—is quite natural [31, 40, 42, 46, 59, 83, 143]. The exploration of the “parameter + time” framework in the important context of affine linear (stable) *parabolic* PDEs—such as the heat equation and the passive scalar convection-diffusion equation (also the Black-Scholes equation of derivative theory [118])—is carried out in [52, 55, 60, 133]; many of the primal-dual approximations, greedy (or, better yet, greedy + POD) sampling strategies, *a posteriori* error estimation concepts, and Offline-Online computational strategies described here for elliptic PDEs admit ready extension to the parabolic case.

The reduced basis methodology, in both the elliptic and parabolic cases, can also be extended to problems with non-affine parametric variation. The strategy is ostensibly simple: reduce the nonaffine operator and data to approximate affine form, and then apply the methods developed for affine operators described in this paper. However, this reduction

²The special issues associated with *saddle problems* [28, 29], in particular the Stokes equations of incompressible flow, are addressed for divergence-free spaces in [57, 67, 114] and non-divergence-free spaces in [135, 139].

must be done efficiently in order to avoid a proliferation of parametric functions and a corresponding degradation of Online response time. This extension is based on the Empirical Interpolation Method (EIM) [18]: a collateral RB space for the offending nonaffine coefficient functions; an interpolation system that avoids costly ($\mathcal{N}$ -dependent) projections; and several (from less rigorous/simple to completely rigorous/very cumbersome) *a posteriori* error estimators. The EIM within the context of RB treatment of elliptic and parabolic PDEs with nonaffine coefficient functions is considered in [52, 53, 95, 138, 148]; the resulting approximations preserve the usual Offline-Online efficiency—the complexity of the Online stage is independent of $\mathcal{N}$.

The reduced basis approach and associated Offline-Online procedures can be applied without serious computational difficulties to quadratic (and arguably cubic [34, 153]) nonlinearities. Much work focuses on the stationary incompressible (quadratically nonlinear) Navier-Stokes equations [29, 50, 57] of incompressible fluid flow: suitable stable approximations are considered in [57, 67, 114, 123, 137, 139]; rigorous *a posteriori* error estimation—within the general Brezzi-Rappaz-Raviart (“BRR”) *a posteriori* framework [30, 34]—is considered in [45, 97, 151, 152]. The latter is admittedly quite complicated, and presently limited to very few parameters—a Reynolds number and perhaps a Prandtl number or aspect ratio.

Symmetric eigenproblems associated with (say) the Laplacian [10] or linear elasticity operator are another important example of quadratic nonlinearities. Reduced basis formulations for one or two lowest eigenvalues (as relevant in structural mechanics) and for the first “many” eigenvalues (as relevant in quantum chemistry) are developed in [84] and [35, 36, 113], respectively. Here, implicitly, the interpretation of the BRR theory is unfortunately less compelling due to the (guaranteed) multiplicity of often nearby solutions; hence the *a posteriori* error estimators for eigenvalue problems [84, 113] are currently less than satisfactory.

Nonpolynomial nonlinearities (in the operator and also output functional) for both elliptic and parabolic PDEs may be considered. The Empirical Interpolation Method can be extended to address this important class of problems [36, 53, 113]: the nonlinearity is treated in a collateral reduced basis expansion, the coefficients of which are then obtained by interpolation relative to the reduced basis approximation of the field variable; the usual Offline-Online efficiency can be maintained—Online evaluation of the output is independent of $\mathcal{N}$. (For alternative approaches to nonlinearities in the ROM context, see [16, 39, 115].) Unfortunately, for this difficult class of problems we can not yet cite either rigorous *a posteriori* error estimators or particularly efficient sampling procedures. (It perhaps not surprising that initial work in RB methods [98, 101], which focused on highly nonlinear problems, attempted neither complete Offline-Online decoupling nor rigorous error estimation.)

Finally, we mention two other topics of current research interest. First, the “reduced basis element method” [79–82] is a marriage of reduced basis and domain decomposition concepts that permits much greater geometric complexity and also provides a framework for the integration of multiple models. Second, (at least linear) hyperbolic problems are also ripe for further development: although there are many issues related to smoothness and stability, there are also proofs-of-concept in both the first order [60, 111] and second order [74] contexts which demonstrate that RB approximation and *a posteriori* error estimation can be gainfully applied to hyperbolic equations.

4 Elliptic Coercive Parametric PDEs

We consider the following problem: Given $\mu \in \mathcal{D} \subset \mathbb{R}^P$, evaluate

$$s^e(\mu) = \ell(u^e(\mu)),$$

where $u^e(\mu) \in X^e(\Omega)$ satisfies

$$a(u^e(\mu), v; \mu) = f(v), \quad \forall v \in X^e. \quad (1)$$

The superscript e refers to “exact.” Here μ is the input parameter—a P -tuple; $\mathcal{D}$ is the parameter domain—a subset of $\mathbb{R}^P$; s^e is the scalar output; ℓ is the linear output functional; u^e is the field variable; Ω is a suitably regular bounded spatial domain in R^d (for $d = 2$ or 3) with boundary $\partial\Omega$; X^e is a Hilbert space; and a and f are the bilinear and linear forms, respectively, associated with our PDE.

We shall exclusively consider second-order partial differential equations, and hence $(H_0^1(\Omega))^v \subset X^e \subset (H^1(\Omega))^v$, where $v = 1$ (respectively, $v = d$) for a scalar (respectively, vector) field. Here $H^1(\Omega) = \{v \in L^2(\Omega) \mid \nabla v \in (L^2(\Omega))^d\}$, $H_0^1(\Omega) = \{v \in H^1(\Omega) \mid v|_{\partial\Omega} = 0\}$, and $L^2(\Omega) = \{v \text{ measurable} \mid \int_{\Omega} v^2 \text{ finite}\}$. We associate to X^e an inner product and induced norm (equivalent to the $H^1(\Omega)$ norm), the choice of which shall be described below.

We shall assume that the bilinear form $a(\cdot, \cdot; \mu): X^e \times X^e \rightarrow \mathbb{R}$ is *continuous* and *coercive* over X^e for all μ in $\mathcal{D}$. (We provide precise definitions of our continuity and coercivity constants and conditions below.) We further assume that f is a *bounded linear* functional over X^e . Under these standard hypotheses on a and f , (1) admits a unique solution.

We shall further presume for most of this paper that we are “in compliance” [112]. In particular, we assume that (i) a is symmetric— $a(w, v; \mu) = a(v, w; \mu)$, $\forall w, v \in X^e$, $\forall \mu \in \mathcal{D}$, and furthermore (ii) $\ell = f$. This assumption will greatly simplify the presentation while still exercising most of the important RB concepts; furthermore, many important engineering problems are in fact “compliant” (see Sect. 6).

At the conclusion of the paper, Sect. 11, we indicate the rather straightforward (primal-dual) extension to the “non-compliant” case in which now a may be non-symmetric and ℓ may be any bounded linear functional over X^e .

We shall make one last assumption, crucial to Offline-Online procedures and hence computational performance. In particular, we assume that the parametric bilinear form a is “affine” in the parameter μ , by which we shall mean

$$a(w, v; \mu) = \sum_{q=1}^Q \Theta^q(\mu) a^q(w, v). \quad (2)$$

Here, for $q = 1, \dots, Q$, the $\Theta^q: \mathcal{D} \rightarrow \mathbb{R}$ are (typically very smooth) μ -dependent functions, and the $a^q: X^e \times X^e \rightarrow \mathbb{R}$ are μ -independent X^e -continuous bilinear forms. (In the compliant case the a^q , $1 \leq q \leq Q$, are additionally symmetric; however, in neither the compliant nor non-compliant cases do we place any restrictions on the “sign” of the a^q , $1 \leq q \leq Q$.) In actual practice, f may also depend affinely on the parameter: in this case, $f(v; \mu)$ may be expressed as a sum of Q^f products of parameter-dependent functions and parameter-independent X^e -bounded linear forms.

As we shall see in Sect. 5, the assumption of affine parameter dependence is *broadly relevant* to many instances of both property *and* geometry parametric variation. Furthermore, as already described in Sect. 3.2, the assumption of affine parameter dependence may be relaxed; however, even in the non-affine case, the (now *approximate*) representation (2) shall still play a crucial computational role in the Offline-Online computational procedures.

We next proceed to the finite element (FE) approximation to the problem (1): Given $\mu \in \mathcal{D} \subset \mathbb{R}^P$, evaluate

$$s^{\mathcal{N}}(\mu) = f(u^{\mathcal{N}}(\mu))$$

(recall our compliance assumption: $\ell = f$), where $u^{\mathcal{N}}(\mu) \in X^{\mathcal{N}} \subset X^e$ satisfies

$$a(u^{\mathcal{N}}(\mu), v; \mu) = f(v), \quad \forall v \in X^{\mathcal{N}}. \quad (3)$$

Here $X^{\mathcal{N}} \subset X^e$ is a sequence of (conforming) FE approximation spaces indexed by $\dim(X^{\mathcal{N}}) \equiv \mathcal{N}$. It follows directly from our assumptions on a , f , and $X^{\mathcal{N}}$ that (3) admits a unique solution. (Note in actual practice we must often replace a with $a^{\mathcal{N}}$ and f with $f^{\mathcal{N}}$ to reflect numerical quadrature and domain approximation “variational crimes.”) We recall that our RB field and RB output shall approximate, for given $\mathcal{N}$, the FE field $u^{\mathcal{N}}(\mu)$ and FE output $s^{\mathcal{N}}(\mu)$ (and not $u^e(\mu)$ and $s^e(\mu)$).

We can now define our inner products and norms for members of $X^{\mathcal{N}}$ and X^e . First,

$$((w, v))_{\mu} \equiv a(w, v; \mu), \quad \forall w, v \in X^e, \quad (4)$$

$$\|w\|_{\mu} \equiv ((w, w))_{\mu}^{1/2}, \quad \forall w \in X^e, \quad (5)$$

shall define our energy inner product and energy norm, respectively. (Recall that a is coercive and symmetric.) Next, for given $\bar{\mu} \in \mathcal{D}$ and (non-negative) real τ ,

$$(w, v)_X \equiv ((w, v))_{\bar{\mu}} + \tau(w, v)_{L^2(\Omega)}, \quad \forall w, v \in X^e, \quad (6)$$

$$\|w\|_X \equiv (w, w)_X^{1/2}, \quad \forall w, v \in X^e,$$

shall define our $X^{\mathcal{N}}$ and X^e inner product and norm, respectively; here $(w, v)_{L^2(\Omega)} \equiv \int_{\Omega} wv$. Corresponding dual norms and dual spaces—for which we must now distinguish between $X^{\mathcal{N}}$ and X^e —will be introduced as needed.³ It is imperative to observe that although our RB approximation shall be built relative to a finite-dimensional FE approximation, we must insist on $(H^1(\Omega))$ -equivalent norms that are stable as $\mathcal{N} \rightarrow \infty$.

Finally, we can now define more precisely our coercivity and continuity constants (and coercivity and continuity conditions). In particular, we define the exact and FE coercivity constants as

$$\alpha^e(\mu) = \inf_{w \in X^e} \frac{a(w, w; \mu)}{\|w\|_X^2}, \quad (7)$$

and

$$\alpha^{\mathcal{N}}(\mu) = \inf_{w \in X^{\mathcal{N}}} \frac{a(w, w; \mu)}{\|w\|_X^2}, \quad (8)$$

respectively. It follows from our coercivity hypothesis that $\alpha^e(\mu) \geq \alpha_0 > 0$, $\forall \mu \in \mathcal{D}$, and from our conforming hypothesis (ignoring any variational crimes) that $\alpha^{\mathcal{N}}(\mu) \geq \alpha^e(\mu)$, $\forall \mu \in \mathcal{D}$. Similarly, we define the exact and FE continuity constants as

$$\gamma^e(\mu) = \sup_{w \in X^e} \sup_{v \in X^e} \frac{a(w, v; \mu)}{\|w\|_X \|v\|_X}, \quad (9)$$

and

$$\gamma^{\mathcal{N}}(\mu) = \sup_{w \in X^{\mathcal{N}}} \sup_{v \in X^{\mathcal{N}}} \frac{a(w, v; \mu)}{\|w\|_X \|v\|_X}, \quad (10)$$

respectively. It follows from our continuity hypothesis that $\gamma^e(\mu)$ is finite $\forall \mu \in \mathcal{D}$, and from our conforming hypothesis (ignoring any variational crimes) that $\gamma^{\mathcal{N}}(\mu) \leq \gamma^e(\mu)$, $\forall \mu \in \mathcal{D}$.

5 Problem “Scope”

We begin in Sect. 5 with perhaps the most difficult issue: the family of geometric parametric variations consistent with

³We note that the choice of $\bar{\mu}$ and τ will affect the quality and efficiency of our RB *a posteriori* error estimators, but will not directly affect our RB output predictions. We return to this point in Sects. 10 and 11.

our affine restriction, (2). Then, in Sect. 5.2 we describe the general class of scalar problems that admit the abstract statement of Sect. 4. (For simplicity, we consider only the scalar case; the vector case [61] permits an analogous treatment.)

5.1 Geometry

5.1.1 Affine Geometry Precondition

The RB recipe, in effect, requires that Ω be a *parameter-independent* domain: if we wish to consider linear combinations of snapshots, these snapshots must be defined relative to a common spatial configuration. Thus to permit geometric variation we must interpret Ω , our parameter-independent reference domain, as the pre-image of $\Omega_o(\mu)$, the parameter-dependent “actual” or “original” domain of interest. The geometric transformation will yield variable (parameter-dependent) coefficients in the reference-domain linear and bilinear forms that, under suitable hypotheses to be discussed below, will take the requisite affine form, (2).

We shall assume that, for all μ in $\mathcal{D}$, $\Omega_o(\mu)$ is expressed as

$$\overline{\Omega}_o(\mu) = \bigcup_{\ell=1}^{L_{\text{reg}}} \overline{\mathcal{R}}_{o,\ell}(\mu), \quad (11)$$

where the $\mathcal{R}_{o,\ell}(\mu)$, $1 \leq \ell \leq L_{\text{reg}}$, are mutually non-overlapping open “regions,”

$$\mathcal{R}_{o,\ell}(\mu) \cap \mathcal{R}_{o,\ell'}(\mu) = \emptyset, \quad 1 \leq \ell < \ell' \leq L_{\text{reg}}, \quad (12)$$

the integrity of which must be respected/preserved in subsequent representations and discretizations. Typically the different regions correspond to different materials and hence material properties, or more generally different (discontinuously varying in space) PDE coefficients; however the regions may also be introduced for algorithmic purposes to ensure well-behaved mappings, as discussed in the next section. We shall refer to the boundaries of $\mathcal{R}_{o,\ell}(\mu)$, $1 \leq \ell \leq L_{\text{reg}}$, that do *not* reside on the boundary of $\Omega_o(\mu)$ as *internal interfaces*.

We next introduce, for all μ in $\mathcal{D}$, a domain decomposition of $\Omega_o(\mu)$,

$$\overline{\Omega}_o(\mu) = \bigcup_{k=1}^{K_{\text{dom}}} \overline{\Omega}_o^k(\mu), \quad (13)$$

where the $\Omega_o^k(\mu)$, $1 \leq k \leq K_{\text{dom}}$, are mutually non-overlapping open subdomains,

$$\Omega_o^k(\mu) \cap \Omega_o^{k'}(\mu) = \emptyset, \quad 1 \leq k < k' \leq K_{\text{dom}}, \quad (14)$$

that “honor” the regions in the sense that

$$\overline{\mathcal{R}}_{o,\ell}(\mu) = \bigcup_{k \in \mathbb{K}_\ell} \overline{\Omega}_o^k(\mu), \quad (15)$$

where the $\mathbb{K}_\ell$, $1 \leq \ell \leq L_{\text{reg}}$, are mutually exclusive subsets of $\{1, \dots, K_{\text{dom}}\}$.

We now choose a value $\mu_{\text{ref}} \in \mathcal{D}$ and define our reference domain as $\Omega \equiv \Omega_o(\mu_{\text{ref}})$. It immediately follows from (13), (14), and (15) that

$$\overline{\Omega} = \bigcup_{k=1}^{K_{\text{dom}}} \overline{\Omega}^k, \quad (16)$$

$$\Omega^k \cap \Omega^{k'} = \emptyset, \quad 1 \leq k < k' \leq K_{\text{dom}}, \quad (17)$$

and

$$\overline{\mathcal{R}}_\ell = \bigcup_{k \in \mathbb{K}_\ell} \overline{\Omega}^k, \quad (18)$$

for $\Omega^k = \Omega_o^k(\mu_{\text{ref}})$, $1 \leq k \leq K_{\text{dom}}$, and $\mathcal{R}_\ell = \mathcal{R}_{o,\ell}(\mu_{\text{ref}})$, $1 \leq \ell \leq L_{\text{reg}}$.

We will build our FE approximation on a very fine “ $\mathcal{N}$ ” FE subtriangulation of the coarse “ K_{dom} ” domain decomposition—which we shall denote our “RB triangulation”—of Ω . (Recall that our FE and RB approximations are defined over the reference domain.) This FE subtriangulation ensures that the FE approximation accurately treats the perhaps discontinuous coefficients (arising from property and geometry variation) associated with the different regions; the subtriangulation also plays an important role in the generation of our affine representation, (2). We emphasize that μ_{ref} only affects the accuracy of the underlying FE approximation upon which the RB discretization and a *a posteriori* error estimator is built: typically a value of μ_{ref} at the “center” of $\mathcal{D}$ minimizes distortion and reduces the requisite $\mathcal{N}$ (for a given acceptable FE error over $\mathcal{D}$).

We now state our Affine Geometry Precondition. We can treat any original domain $\Omega_o(\mu)$ and associated regions (11) that admits a domain decomposition (13)–(15) for which (there exists a $\mu_{\text{ref}} \in \mathcal{D}$ such that), $\forall \mu \in \mathcal{D}$,

$$\overline{\Omega}_o^k(\mu) = \mathcal{T}^{\text{aff},k}(\overline{\Omega}^k; \mu), \quad 1 \leq k \leq K_{\text{dom}}, \quad (19)$$

for affine mappings $\mathcal{T}^{\text{aff},k}(\cdot; \mu): \Omega^k \rightarrow \Omega_o^k(\mu)$, $1 \leq k \leq K_{\text{dom}}$, that are (i) individually *bijective*, and (ii) collectively *continuous*,

$$\begin{aligned} \mathcal{T}^{\text{aff},k}(\mathbf{x}; \mu) &= \mathcal{T}^{\text{aff},k'}(\mathbf{x}; \mu), \\ \forall \mathbf{x} \in \overline{\Omega}^k \cap \overline{\Omega}^{k'}, \quad 1 \leq k < k' \leq K_{\text{dom}}. \end{aligned} \quad (20)$$

The Affine Geometry Precondition is a necessary condition for affine parameter dependence as defined in (2). Note that we purposely define K_{dom} with respect to the *exact* problem, rather than the FE approximation: K_{dom} can not depend on $\mathcal{N}$ (to be meaningful).

We now define our (bijective) affine mappings more explicitly: for $1 \leq k \leq K_{\text{dom}}$, any μ in $\mathcal{D}$, and any $x \in \Omega^k$,

$$\mathcal{T}_i^{\text{aff},k}(x; \mu) = C_i^{\text{aff},k}(\mu) + \sum_{j=1}^d G_{ij}^{\text{aff},k}(\mu)x_j, \quad 1 \leq i \leq d, \quad (21)$$

for given $C^{\text{aff},k}: \mathcal{D} \rightarrow \mathbb{R}^d$ and $G^{\text{aff},k}: \mathcal{D} \rightarrow \mathbb{R}^{d \times d}$. We can then define the associated Jacobians

$$J^{\text{aff},k}(\mu) = |\det(G^{\text{aff},k}(\mu))|, \quad 1 \leq k \leq K_{\text{dom}}, \quad (22)$$

where $\det$ denotes determinant; note the Jacobian is constant in space over each subdomain. We further define, for any $\mu \in \mathcal{D}$,

$$D^{\text{aff},k}(\mu) = (G^{\text{aff},k}(\mu))^{-1}, \quad 1 \leq k \leq K_{\text{dom}}; \quad (23)$$

this matrix shall prove convenient in subsequent derivative transformations.

We may interpret our local mappings in terms of a global transformation. In particular, for any $\mu \in \mathcal{D}$, the local mappings (19) induce a global bijective *piecewise-affine* transformation $\mathcal{T}^{\text{aff}}(\cdot; \mu): \Omega \rightarrow \Omega_o(\mu)$: for any $\mu \in \mathcal{D}$,

$$\mathcal{T}^{\text{aff}}(x; \mu) = \mathcal{T}^{\text{aff},k}(x; \mu), \quad k = \min_{k' \in \{1, \dots, K_{\text{dom}}\} | x \in \overline{\Omega}^{k'}} k'; \quad (24)$$

note the one-to-one property of this mapping (and, hence the arbitrariness of our “min” choice in (24)) is ensured by the interface condition (20). We can further demonstrate that these global *continuous* mappings are compatible with our second-order PDE variational formulation: for any $\mu \in \mathcal{D}$, given any $w_o \in H^1(\Omega_o(\mu))$, $w \equiv w_o \circ \mathcal{T}^{\text{aff}} \in H^1(\Omega)$; this ensures that our mapped problem on the reference domain is of the classical “conforming” variety.

Although this largely concludes the formal exposition of admissible geometry variations, the relevance and ultimately application of these conditions requires familiarity with the scope of the affine mappings (21). We first consider, in Sect. 5.1.2, a single subdomain (and hence necessarily a single region). We next consider, in Sect. 5.1.3, the case of multiple subdomains: we give a prescriptive definition of a family of admissible domains, and we briefly summarize an algorithm to identify the associated domain decomposition and affine mappings. Finally, in Sects. 5.2 and 5.3, we discuss the incorporation of these affine mappings into our weak form and provide several illustrative computational results.

5.1.2 Affine Mappings: Single Subdomain

As we consider a single subdomain in this section, we shall suppress the subdomain superscript for clarity of exposition.

We shall focus on the two-dimensional case ($d = 2$); the extension to three dimensions ($d = 3$) is, although certainly possible, not trivial. Note that some of the elementary material presented here is available in linear algebra or computer graphics texts [90, 146]; we emphasize application within our particular parametric PDE context.

We first rewrite our affine transformation (21), for simplicity, without the subdomain superscript: for any given $\mu \in \mathcal{D}$, the reference domain Ω induces the parameter-dependent geometry of interest, $\Omega_o(\mu)$, through the affine mapping

$$\mathcal{T}_i^{\text{aff}}(x; \mu) = C_i^{\text{aff}}(\mu) + \sum_{j=1}^d G_{ij}^{\text{aff}}(\mu)x_j, \quad 1 \leq i \leq d; \quad (25)$$

we shall refer to $C^{\text{aff}}(\mu) \in \mathbb{R}^d$ and $G^{\text{aff}}(\mu) \in \mathbb{R}^{d \times d}$ as the “mapping coefficients.” Under our assumption that the mapping is invertible we know that our Jacobian, $J^{\text{aff}}(\mu)$ of (22), is strictly positive, and that the derivative transformation matrix, $D^{\text{aff}}(\mu) = (G^{\text{aff}}(\mu))^{-1}$ of (23), is well defined.

We recall that, in two dimensions, an affine transformation maps *straight lines* to *straight lines* and in fact parallel lines to parallel lines and indeed parallel lines of equal length to parallel lines of equal length: it follows that a triangle maps onto a triangle (or a tetrahedron onto a tetrahedron in three dimensions), that a parallelogram maps onto a parallelogram (or a parallelepiped onto a parallelepiped in three dimensions), and that an n -gon maps onto an n -gon (or an n -hedron onto an n -hedron in three dimensions). We also recall that an affine transformation maps *ellipses* to *ellipses* (and ellipsoids to ellipsoids). These “straight line” and “ellipse” properties will be crucial for the descriptions of domains relevant in the engineering context.

Basic Technology

Our affine transformation (25) is completely defined, for $d = 2$, by the $(d(d+1)) = 6$ mapping coefficients $C^{\text{aff}}(\mu) \in \mathbb{R}^{d=2}$ and $G^{\text{aff}}(\mu) \in \mathbb{R}^{d \times d=2 \times 2}$. It immediately follows that, for any $\mu \in \mathcal{D}$, we can uniquely identify $C^{\text{aff}}(\mu)$ and $G^{\text{aff}}(\mu)$ from the relationship between 3 non-colinear pre-image points—or “nodes”—in Ω , $(\bar{z}^1, \bar{z}^2, \bar{z}^3) \equiv ((\bar{z}_1^1, \bar{z}_2^1), (\bar{z}_1^2, \bar{z}_2^2), (\bar{z}_1^3, \bar{z}_2^3))$, and 3 parametrized image nodes in $\Omega_o(\mu)$, $(\bar{z}_o^1(\mu), \bar{z}_o^2(\mu), \bar{z}_o^3(\mu)) \equiv ((\bar{z}_{o1}^1, \bar{z}_{o2}^1), (\bar{z}_{o1}^2, \bar{z}_{o2}^2), (\bar{z}_{o1}^3, \bar{z}_{o2}^3))(\mu)$. (Note that, from our assumption that the affine transformation is bijective, the image nodes are perforce also non-colinear.) In particular, for given $\mu \in \mathcal{D}$, application of (25) to the selected “nodes” yields

$$\bar{z}_{oi}^m(\mu) = C_i^{\text{aff}}(\mu) + \sum_{j=1}^2 G_{ij}^{\text{aff}}(\mu)\bar{z}_j^m, \quad 1 \leq i \leq 2, \quad 1 \leq m \leq 3; \quad (26)$$

(26) constitutes 6 independent equations by which to determine the 6 mapping coefficients. (In three space dimensions, we must follow the “trajectories” of the 3 coordinates of 4 pre-image/image points: this yields 12 equations for the 12 mapping coefficients.)

To be more explicit in our construction, we first form the matrix $\mathbb{B}^{\text{aff}} \in \mathbb{R}^{6 \times 6}$ (more generally, $\mathbb{R}^{(d^2+d) \times (d^2+d)}$),

$$\mathbb{B}^{\text{aff}} = \begin{bmatrix} 1 & 0 & \bar{z}_1^1 & \bar{z}_2^1 & 0 & 0 \\ 0 & 1 & 0 & 0 & \bar{z}_1^1 & \bar{z}_2^1 \\ 1 & 0 & \bar{z}_1^2 & \bar{z}_2^2 & 0 & 0 \\ 0 & 1 & 0 & 0 & \bar{z}_1^2 & \bar{z}_2^2 \\ 1 & 0 & \bar{z}_1^3 & \bar{z}_2^3 & 0 & 0 \\ 0 & 1 & 0 & 0 & \bar{z}_1^3 & \bar{z}_2^3 \end{bmatrix}. \quad (27)$$

We further introduce the vector $V^{\text{aff}}(\boldsymbol{\mu})$ of image nodal locations,

$$V^{\text{aff}}(\boldsymbol{\mu}) = \begin{bmatrix} \bar{z}_{o1}^1(\boldsymbol{\mu}) \\ \bar{z}_{o2}^1(\boldsymbol{\mu}) \\ \bar{z}_{o1}^2(\boldsymbol{\mu}) \\ \bar{z}_{o2}^2(\boldsymbol{\mu}) \\ \bar{z}_{o1}^3(\boldsymbol{\mu}) \\ \bar{z}_{o2}^3(\boldsymbol{\mu}) \end{bmatrix}. \quad (28)$$

The solution of the linear system (26) can then be succinctly expressed as

$$\begin{bmatrix} C_1^{\text{aff}}(\boldsymbol{\mu}) \\ C_2^{\text{aff}}(\boldsymbol{\mu}) \\ G_{11}^{\text{aff}}(\boldsymbol{\mu}) \\ G_{12}^{\text{aff}}(\boldsymbol{\mu}) \\ G_{21}^{\text{aff}}(\boldsymbol{\mu}) \\ G_{22}^{\text{aff}}(\boldsymbol{\mu}) \end{bmatrix} = (\mathbb{B}^{\text{aff}})^{-1} V^{\text{aff}}(\boldsymbol{\mu}); \quad (29)$$

note that $\mathbb{B}^{\text{aff}}$ is non-singular under our hypothesis of non-colinear pre-image nodes.

The matrix $\mathbb{B}^{\text{aff}}$ is independent of $\boldsymbol{\mu}$; the parametric dependence derives from $V^{\text{aff}}(\boldsymbol{\mu})$. In particular, the $\boldsymbol{\mu}$ dependence of the geometry “enters” through the parametrized locations of the image nodes $(\bar{z}_o^1(\boldsymbol{\mu}), \bar{z}_o^2(\boldsymbol{\mu}), \bar{z}_o^3(\boldsymbol{\mu}))$ as represented in $V^{\text{aff}}(\boldsymbol{\mu})$. To illustrate how the parametric dependence propagates from the (desired) parametrized domain to

the mapping coefficients—and to exercise the general affine technology presented above—we now consider several different domains. (In Sect. 5.2 we consider how the parametric dependence further propagates from the mapping coefficients to the affine expansion of the bilinear form, (2), associated with our PDE.)

We note that parallelograms and (in three space dimensions) parallelepipeds are the most intuitive subdomains by which to effect transformations “by hand”—invoking the usual translation, dilation, rotation, and shear primitives; we consider such a case in the three-dimensional linear elasticity example of Sect. 6. However, it is (Standard) triangles, Elliptical Triangles, and more general “Curvy” Triangles which admit symbolic and numerical automation, and which are thus the building blocks of choice in general multi-subdomain software (e.g., [1]). We shall thus focus on triangular building blocks in our discussion here.

$\Omega_o(\boldsymbol{\mu})$: (Standard) Triangles

Triangles will thus be the “workhorses” in our geometric decompositions: the fundamental building blocks. To demonstrate the application of the technology to triangular domains, we consider a single parameter, $\boldsymbol{\mu} \equiv \mu_1 \in \mathcal{D} \equiv [0.5, 2]$. We take for Ω the triangle with vertices (counter-clockwise) $(0, 0)$, $(1, 0)$, $(1, 1)$; these vertices shall also serve as the pre-image nodes, and hence $\bar{z}^1 = (0, 0)$, $\bar{z}^2 = (1, 0)$, $\bar{z}^3 = (1, 1)$. We take for $\Omega_o(\boldsymbol{\mu})$ the triangle with vertices (counter-clockwise) $(0, 0)$, $(\mu_1, 0)$, $(1, 1)$; these vertices shall also serve as the image nodes, and hence $\bar{z}_o^1(\mu_1) = (0, 0)$, $\bar{z}_o^2(\mu_1) = (\mu_1, 0)$, $\bar{z}_o^3(\mu_1) = (1, 1)$. (Note for triangles, our three points uniquely define not only the transformation but also the reference and parametrized domains.) As already noted in Sect. 5.1.1, the pre-image nodes correspond to the image nodes for a particular value of the parameter: in our example here, $\mu_{\text{ref}} = 1$ such that $(\bar{z}^1, \bar{z}^2, \bar{z}^3) = (\bar{z}_o^1(\mu_{\text{ref}}), \bar{z}_o^2(\mu_{\text{ref}}), \bar{z}_o^3(\mu_{\text{ref}}))$. The domains Ω and $\Omega_o(\boldsymbol{\mu})$ are depicted in Figs. 1a and b, respectively.

Fig. 1 (a) Reference domain Ω , and (b) actual (or original) domain $\Omega_o(\mu_1)$

If we turn the crank, we find from (29) that

$$\begin{bmatrix} C_1^{\text{aff}}(\mu_1) \\ C_2^{\text{aff}}(\mu_1) \\ G_{11}^{\text{aff}}(\mu_1) \\ G_{12}^{\text{aff}}(\mu_1) \\ G_{21}^{\text{aff}}(\mu_1) \\ G_{22}^{\text{aff}}(\mu_1) \end{bmatrix} = \begin{bmatrix} 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 \\ -1 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & -1 & 0 & 1 & 0 \\ 0 & -1 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & -1 & 0 & 1 \end{bmatrix} \begin{bmatrix} 0 \\ 0 \\ \mu_1 \\ 0 \\ 1 \\ 1 \end{bmatrix}, \quad (30)$$

and hence ($C^{\text{aff}}(\mu_1) = 0$ —no translation—and)

$$G^{\text{aff}}(\mu_1) = \begin{bmatrix} \mu_1 & 1 - \mu_1 \\ 0 & 1 \end{bmatrix}. \quad (31)$$

It directly follows from (22) and (23) that

$$J^{\text{aff}}(\mu_1) = \mu_1, \quad (32)$$

and

$$D^{\text{aff}}(\mu_1) = \begin{bmatrix} \frac{1}{\mu_1} & -\frac{1-\mu_1}{\mu_1} \\ 0 & 1 \end{bmatrix}. \quad (33)$$

Note that D^{aff} is *not* linear in μ_1 . (We also note that a translation would affect $C^{\text{aff}}(\mu_1)$ but not $G^{\text{aff}}(\mu_1)$, $J^{\text{aff}}(\mu_1)$, or $D^{\text{aff}}(\mu_1)$.)

We can readily construct an affine map from *any* reference triangle in $\mathbb{R}^2$ (simplex in d space dimensions), $\Omega \equiv \text{SIM}$, onto *any* desired triangle in $\mathbb{R}^2$, $\Omega_o(\mu) \equiv \text{SIM}_o(\mu)$: it will prove most convenient to choose for our nodes (say, in two dimensions) $(\bar{z}^1, \bar{z}^2, \bar{z}^3)$ and $(\bar{z}_o^1(\mu), \bar{z}_o^2(\mu), \bar{z}_o^3(\mu))$ the vertices of SIM and $\text{SIM}_o(\mu)$. (We could also choose the “barycentric” co-ordinates of the FE context [27, 41].)

$\Omega_o(\mu)$: “Elliptical Triangles”

It is important to note that the Geometric Affine Precondition places no restriction on the shape of the (sub)domains: although an affine mapping transforms straight lines to straight lines, we are not constrained (say, for $d = 2$) to straight lines in the definition of the boundary of the (sub)domain. For example, if Ω is the unit circle then translation creates a shifted circle, isotropic dilation creates a larger or smaller circle, and anisotropic dilation creates an ellipse.

In fact, an “Elliptical Triangle” $\Omega_o(\mu)$ and even a more general “Curvy Triangle” $\Omega_o(\mu)$ shall be the basic building blocks in our multidomain framework of Sect. 5.2. We consider in this section the former. We depict in Fig. 2 the two

Fig. 2 (a) “Inwards” (convex domain) Elliptical Triangle, and (b) “Outwards” (non-convex domain) Elliptical Triangle

types of Elliptical Triangles, “Inwards” and “Outwards.” In both cases, the Elliptical Triangle $\Omega_o(\mu)$ is defined by the three nodes (vertices) $\bar{z}_o^1(\mu), \bar{z}_o^2(\mu), \bar{z}_o^3(\mu)$, by the two straight edges $\bar{z}_o^1(\mu)\bar{z}_o^2(\mu)$ and $\bar{z}_o^1(\mu)\bar{z}_o^3(\mu)$, and by an elliptical arc $\bar{z}_o^2(\mu)\bar{z}_o^3(\mu)_{\text{arc}}$. We now define each of these components more precisely, and identify constraints that must be honored in order to obtain admissible/controlled geometric descriptions.

We shall describe a point on a prescribed parametrized ellipse

$$(\mathbf{x}_o - O(\mu))^T Q_{\text{rot}}(\mu) S^{-2}(\mu) Q_{\text{rot}}(\mu)^T (\mathbf{x}_o - O(\mu)) = 1 \quad (34)$$

as

$$\mathbf{x}_o \equiv \begin{pmatrix} x_{o1} \\ x_{o2} \end{pmatrix} = O(\mu) + Q_{\text{rot}}(\mu) S(\mu) \begin{pmatrix} \cos t \\ \sin t \end{pmatrix} \quad (35)$$

for given $t \in \mathbb{R}$. Here $O(\mu): \mathcal{D} \rightarrow \mathbb{R}^2$ is the center of the ellipse; $\rho_1: \mathcal{D} \rightarrow \mathbb{R}_+$ and $\rho_2: \mathcal{D} \rightarrow \mathbb{R}_+$ are semi-axes that define the diagonal dilation or “scaling” matrix

$$S(\mu) \equiv \begin{pmatrix} \rho_1(\mu) & 0 \\ 0 & \rho_2(\mu) \end{pmatrix}; \quad (36)$$

and $\phi(\mu): \mathcal{D} \rightarrow \mathbb{R}$ is an angle of inclination that defines the rotation matrix

$$Q_{\text{rot}}(\mu) = \begin{pmatrix} \cos \phi(\mu) & -\sin \phi(\mu) \\ \sin \phi(\mu) & \cos \phi(\mu) \end{pmatrix}. \quad (37)$$

We depict the geometry in Fig. 3.

We now take our points $\bar{z}_o^2(\mu), \bar{z}_o^3(\mu)$ as

$$\bar{z}_o^m(\mu) = O(\mu) + Q_{\text{rot}}(\mu) S(\mu) \begin{pmatrix} \cos t^m \\ \sin t^m \end{pmatrix}, \quad m = 1, 2, \quad (38)$$

Fig. 3 Definition of a point on a prescribed parametrized ellipse

for prescribed $t^2 \in \mathbb{R}$, $t^3 \in \mathbb{R}$, such that $0 \leq t^3 - t^2 < \pi$; the desired elliptical arc is then given by

$$\overline{\bar{z}_0^2(\mu) \bar{z}_0^3(\mu)}^{\text{arc}} = \left\{ O(\mu) + Q_{\text{rot}}(\mu) S(\mu) \begin{pmatrix} \cos t \\ \sin t \end{pmatrix} \mid t^2 \leq t \leq t^3 \right\}. \quad (39)$$

It remains to choose the third point (actually, the first point, $\bar{z}_0^1$) to achieve our ends.

In particular, for an “Elliptical Triangle,” and unlike a Standard Triangle, the choice of the first point is not arbitrary: $\bar{z}_0^1(\mu)$ must be selected such that, under affine transformation, we generate the desired arc (39); if the first point is not selected appropriately not only will we not control the arc but, in the multidomain context, we risk a *discontinuous* global mapping. The first point must also be chosen to ensure a “proper” (Elliptical) Triangle: the internal angle condition $0 < \theta^{12} < \pi$, $0 < \theta^{23} < \pi$, $0 < \theta^{31} < \pi$, must be satisfied. (We shall initially assume that the internal angle condition is satisfied, and subsequently derive the relevant prescriptive criteria.)

We shall choose the first point as

$$\bar{z}_0^1(\mu) = O(\mu) + \omega Q_{\text{rot}}(\mu) S(\mu) \begin{pmatrix} \cos t^1 \\ \sin t^1 \end{pmatrix}, \quad (40)$$

for given $\omega \in \mathbb{R}$ and $t^1 \in [t^2, t^3]$. We can hence express our three image points $\bar{z}_0^m(\mu)$, $1 \leq m \leq 3$, as

$$\bar{z}_0^m(\mu) = O(\mu) + \omega^m Q_{\text{rot}}(\mu) S(\mu) \begin{pmatrix} \cos t^m \\ \sin t^m \end{pmatrix}, \quad (41)$$

for $\omega^1 = \omega$, $\omega^2 = \omega^3 = 1$. It follows that our pre-image points are then given by

$$\bar{z}^m = O(\mu_{\text{ref}}) + \omega^m Q_{\text{rot}}(\mu_{\text{ref}}) S(\mu_{\text{ref}}) \begin{pmatrix} \cos t^m \\ \sin t^m \end{pmatrix}, \quad (42)$$

for $1 \leq m \leq 3$.

We conclude from (41),(42) that (recall $Q_{\text{rot}}(\mu)$ is orthogonal)

$$\bar{z}_0^m(\mu) = O(\mu)$$

$$\begin{aligned} & + Q_{\text{rot}}(\mu) S(\mu) S(\mu_{\text{ref}})^{-1} Q_{\text{rot}}(\mu_{\text{ref}})^T \\ & \times (\bar{z}^m - O(\mu_{\text{ref}})) \\ & = O(\mu) - Q_{\text{rot}}(\mu) S(\mu) S(\mu_{\text{ref}})^{-1} \\ & \times Q_{\text{rot}}(\mu_{\text{ref}})^T O(\mu_{\text{ref}}) \\ & + Q_{\text{rot}}(\mu) S(\mu) S(\mu_{\text{ref}})^{-1} Q_{\text{rot}}(\mu_{\text{ref}})^T \bar{z}^m, \end{aligned} \quad (43)$$

for $1 \leq m \leq 3$. We can then directly identify our affine mapping coefficients from (43) as

$$\begin{aligned} C^{\text{aff}}(\mu) & = O(\mu) - Q_{\text{rot}}(\mu) S(\mu) S(\mu_{\text{ref}})^{-1} \\ & \times Q_{\text{rot}}(\mu_{\text{ref}})^T O(\mu_{\text{ref}}), \end{aligned} \quad (44)$$

and

$$G^{\text{aff}}(\mu) = Q_{\text{rot}}(\mu) S(\mu) S(\mu_{\text{ref}})^{-1} Q_{\text{rot}}(\mu_{\text{ref}})^T; \quad (45)$$

$\mathcal{T}^{\text{aff}}$ is then given by (25). (Of course, if we were to apply the formal procedure of (27)–(29) we would arrive at the same result.)

We note from (44)–(45) that

$$\begin{aligned} x_o - O(\mu) & = Q_{\text{rot}}(\mu) S(\mu) S(\mu_{\text{ref}})^{-1} \\ & \times Q_{\text{rot}}(\mu_{\text{ref}})^T (x - O(\mu_{\text{ref}})), \end{aligned} \quad (46)$$

and hence

$$\begin{aligned} x - O(\mu_{\text{ref}}) & = Q_{\text{rot}}(\mu_{\text{ref}}) S(\mu_{\text{ref}}) S(\mu)^{-1} Q_{\text{rot}}(\mu)^T \\ & \times (x_o - O(\mu)). \end{aligned} \quad (47)$$

We observe, not unexpectedly, that the mapping is a translation combined with the product of a rotation and dilation.

We can now define our Elliptical Triangles. In particular, for both the Inwards and Outwards cases, we can write (since Ω is perforce star-shaped with respect to $\bar{z}^1$)

$$\bar{\Omega} = \bigcup_{t \in [t_2, t_3]} \overline{\bar{z}^1 \left(O(\mu_{\text{ref}}) + Q_{\text{rot}}(\mu_{\text{ref}}) S(\mu_{\text{ref}}) \begin{pmatrix} \cos t \\ \sin t \end{pmatrix} \right)}. \quad (48)$$

It immediately follows from (46)—recall straight lines map to straight lines—that

$$\begin{aligned} \Omega_o(\mu) & = \mathcal{T}^{\text{aff}} \bigcup_{t \in [t_2, t_3]} \overline{\bar{z}^1 \left(O(\mu_{\text{ref}}) + Q_{\text{rot}}(\mu_{\text{ref}}) S(\mu_{\text{ref}}) \begin{pmatrix} \cos t \\ \sin t \end{pmatrix} \right)} \\ & = \bigcup_{t \in [t_2, t_3]} \mathcal{T}^{\text{aff}} \bar{z}^1 \left(O(\mu_{\text{ref}}) + Q_{\text{rot}}(\mu_{\text{ref}}) S(\mu_{\text{ref}}) \begin{pmatrix} \cos t \\ \sin t \end{pmatrix} \right) \\ & = \bigcup_{t \in [t_2, t_3]} \overline{\bar{z}_0^1 \left(O(\mu) + Q_{\text{rot}}(\mu) S(\mu) \begin{pmatrix} \cos t \\ \sin t \end{pmatrix} \right)}, \end{aligned} \quad (49)$$

Fig. 4 Regions in which $\bar{z}_0^1(\mu)$ must reside for an Inwards Elliptical Triangle ($\mathcal{R}_{\text{In}}(\mu)$) and an Outwards Elliptical Triangle ($\mathcal{R}_{\text{Out}}(\mu)$) in order to ensure the internal angle condition

as required by (39). (It might appear that to obtain the requisite elliptical arc we must sacrifice free selection of $\bar{z}^1(\mu)$: this is true. However, in the multidomain context, the $\bar{z}^1(\mu)$ shall always be internal nodes; we shall thus retain control of the (boundary of the) actual domain $\Omega_0(\mu)$.)

It remains to obtain conditions on ω such that the three internal angles, θ^{12} , θ^{23} , θ^{31} , are all bounded between 0 and π : the angle conditions ensure a well-defined domain and subsequent finite element triangulation. It is also important to confirm, as we shall, that the condition can be satisfied for all $\mu \in \mathcal{D}$. To begin, we consider the diagram in Fig. 4. Clearly, a necessary and sufficient condition to ensure the angle condition for Inwards Elliptical Triangles (respectively, Outward Elliptical Triangles) is $\bar{z}_0^1(\mu) \in \mathcal{R}_{\text{In}}(\mu)$ (respectively, $\bar{z}_0^1(\mu) \in \mathcal{R}_{\text{Out}}(\mu)$), where

$$\begin{aligned} \mathcal{R}_{\text{In}}(\mu) = \{ \bar{z}_0^1(\mu) \in \mathbb{R}^2 \mid & (\bar{z}_0^1(\mu) - \bar{z}_0^2(\mu))^T n^2(\mu) < 0, \\ & (\bar{z}_0^1(\mu) - \bar{z}_0^3(\mu))^T n^3(\mu) < 0, \\ & (\bar{z}_0^1(\mu) - \bar{z}_0^{2,3}(\mu))^T n^{2,3}(\mu) < 0 \} \end{aligned} \quad (50)$$

and

$$\begin{aligned} \mathcal{R}_{\text{Out}}(\mu) = \{ \bar{z}_0^1(\mu) \in \mathbb{R}^2 \mid & (\bar{z}_0^1(\mu) - \bar{z}_0^2(\mu))^T n^2(\mu) > 0, \\ & (\bar{z}_0^1(\mu) - \bar{z}_0^3(\mu))^T n^3(\mu) > 0 \} \end{aligned} \quad (51)$$

for $\bar{z}_0^m(\mu)$, $1 \leq m \leq 3$, given by (41). Here $n^2(\mu)$ and $n^3(\mu)$ are the outwards-oriented normals to the ellipse at $\bar{z}_0^2(\mu)$ and $\bar{z}_0^3(\mu)$, respectively, $\bar{z}_0^{2,3}(\mu) = \frac{1}{2}(\bar{z}_0^2(\mu) + \bar{z}_0^3(\mu))$, and $n^{2,3}(\mu)$ is the “outward” normal to the line segment $\bar{z}_0^2(\mu)\bar{z}_0^3(\mu)$ at $\bar{z}_0^{2,3}(\mu)$.

It is simple to derive, based on elementary trigonometric identities, explicit conditions on ω such that the angle conditions (50), (51) shall be satisfied. Several happy circumstances conspire to ensure that our condition is μ -independent: hence if the angle condition is honored for $\mu_{\text{ref}} \in \mathcal{D}$ corresponding to the reference domain, then the angle condition is satisfied for all $\mu \in \mathcal{D}$. We provide here

the condition for the particular choice $t^1 = \frac{1}{2}(t^2 + t^3)$. For this choice, recalling that $0 < t^3 - t^2 < \pi$, (50) reduces to

$$\omega < \cos\left(\frac{t^3 - t^2}{2}\right)$$

for the Inwards case, and (51) reduces to

$$\omega > \frac{1}{\cos(\frac{t^3 - t^2}{2})}$$

for the Outwards case. (It is clear that, in the multidomain context, $t^3 - t^2$ must be not just less than π but in fact well away from π .)

It is evident that we must exercise some care in the construction of Elliptical Triangles to ensure controlled elliptical arcs, continuous (Inwards/Outwards) mappings, and well-defined internal angles. We do note (as some consolation) that Elliptical Triangles are “consistent” under refinement: if we split either a straight edge or the elliptical arc of an Elliptical Triangle described by (41) for ω satisfying the internal angle conditions (50), (51), we obtain two daughter Elliptical Triangles each described by (41) for (different) ω satisfying the internal angle conditions (50), (51). (In fact, the daughter Elliptical Triangles share the same affine mapping as the parent.) This edge-split consistency property provides, in the multidomain context, a simple mechanism by which to integrate Elliptical Triangles into domain decompositions which respect boundaries and internal interfaces. We discuss this further below.

$\Omega_0(\mu)$: “Curvy” Triangles

The extension from “Elliptical” Triangles to (general) “Curvy” Triangles is, at least formally, very straightforward: we simply replace $\cos t$, $\sin t$ in (41) (and in all subsequent occurrences) with a general parametrization $g_1(t)$, $g_2(t)$. We restrict our attention to curvy arcs which, within a particular “Curvy” Triangle, are either strictly convex (“Inwards”) or strictly concave (“Outwards”) for all $\mu \in \mathcal{D}$; this condition is easily articulated in terms of the sign of the derivative of the normal, or equivalently the curvature. It can then be shown that our internal angle conditions (50), (51) are directly applicable and (as in the Elliptical case) *parameter independent*. We can further demonstrate that, for a proper choice of “center,” these angle conditions reduce to a simple set of algebraic equations—in terms of $g_1(t)$, $g_2(t)$, $g_1'(t)$, and $g_2'(t)$ evaluated at t^1 , t^2 , and t^3 —that indeed admit a feasible solution. However, even in the convex/concave case, we can not express the Curvy Triangle solution in any simple closed form analogous to the Elliptical Triangle result.

5.1.3 Piecewise-Affine Mappings: Multiple Subdomains

A single affine mapping can treat only a very limited family of parametrized domains $\Omega_0(\mu)$. However, piecewise affine

mappings—in our case (typically) based on a domain decomposition in Standard, Elliptical, and Curvy Triangles—can address a much larger class of geometric variations. We restrict attention here to two space dimensions (in the examples of Sect. 6 we consider a specific three-dimensional problem).

We shall consider “Elliptical-Edge” domains: domains and associated regions for which the boundary and internal interfaces can be represented by either straight edges or the elliptical arcs described by (39); we shall naturally choose for our building blocks (Standard and) Elliptical Triangles. (An analogous family of “Curvy-Edge” domains built from Curvy Triangles can also be considered. The Curvy-Edge procedure is very similar to the Elliptic-Edge procedure, but on somewhat less firm theoretical ground. We consider an example in the next section.)

There are three steps to the multidomain mapping process. First (our emphasis in this section), we shall generate an “RB triangulation” (17) of the *reference* domain Ω , (16), and associated reference regions, (18), that is compatible with the mapping continuity condition (20); second (as already developed in the previous section), we will construct the parameter-dependent affine mappings (21) for each subdomain following the recipe (27)–(29); third (as described in the next section), we will translate these parametric mappings into PDE coefficients, and then “optimize”—coalesce similar mapping-induced PDE coefficients into single terms—to arrive at an economical affine expansion, (2). In fact, the first and second steps are implicitly coupled—most notably (but not exclusively) for Elliptical Triangles in which the point selection must be constrained, (41) and (50), (51), to ensure consistent, continuous, and invertible affine mappings.

Our focus here is thus on the RB triangulation of Ω . We wish to generate a domain decomposition (17) of our reference geometry Ω , (16), and associated reference regions, (18), that is compatible with the mapping global continuity condition (20) (and in particular satisfies the Elliptical Triangle consistency/continuity (41) and internal angle, (50), (51), conditions). The User input is a set of (parametrized) “control” points and edges that completely specify the domain boundary and region internal interfaces. There are two stages to the algorithm (which shares some features with classical FE triangulation):

Stage 1. In this stage we focus on the Elliptical Triangles required by *non-straight* edges on the boundary or internal interfaces. We introduce an (Inwards or Outwards) Elliptical Triangle for each elliptical-arc boundary edge and two (an Inwards and an Outwards) Elliptical Triangles for each elliptical-arc internal interface edge; we perform “splits” of the elliptical-arc boundary/interface edges as necessary to ensure reference Elliptical Triangles that satisfy both the internal angle conditions (50), (51) and the region integrity

conditions (18). The new points created by introduction of the Elliptical Triangles are denoted interior control points. *Stage 2.* In this stage we “fill” in the remainder of the domain with Standard Triangles: (a) We perform a Delaunay triangulation on the boundary/internal interface control points (initially provided by the User) *and* the interior control points (initially introduced in Stage 1 of the algorithm). (b) We next search for a (any) “disrespectful” edge: here a “disrespectful” edge is an edge that belongs to the domain boundary or an internal interface *or* an Elliptical Triangle but does not (yet) belong to the Delaunay triangulation. (c) We then split the “disrespectful” edge; note this step will create additional boundary/internal interface or interior control points (and Elliptical Triangles). (d) We repeat (a)–(c) until no disrespectful edges remain. (Note that since Stage 1 already enforces all curved-edge considerations, Stage 2 need only consider “logical” edges associated with the control point connectivity graph.)

Note the mapping continuity condition (20) is satisfied for both (i) Elliptical Triangles—the constrained choice of points, (41), ensures Inwards/Outwards compatibility, and (ii) Standard Triangles—two affine mappings that agree at the endpoints of a segment also agree at all points on the segment.

Unfortunately, our algorithm does *not* guarantee—even for very simple parametric domains $\Omega_o(\mu)$ with straight-edge boundaries decomposed in Standard Triangles—that the Jacobians of the associated affine mappings will remain (strictly) positive for all $\mu \in \mathcal{D}$, and in particular for μ far from μ_{ref} . Equivalently, our algorithm does *not* guarantee that a “valid” domain decomposition of Ω , (16)–(18), will induce a “valid” domain decomposition of $\Omega_o(\mu)$, (13)–(15), for all $\mu \in \mathcal{D}$.⁴ Furthermore, even if the Jacobian does not vanish, “small” Jacobians corresponding to excessive distortion will lead to FE approximations—which we recall shall be built on a very fine “ $\mathcal{N}$ ” FE subtriangulation of the coarse “ K_{dom} ” RB triangulation—that are at best inefficient and at worst very ill-conditioned.

We further caution that even well-behaved/non-singular mappings can be quite “inefficient” as regards ultimate performance of the reduced basis approximation. Inefficient RB triangulations are characterized by many parametrically “dissimilar” triangles that, in turn, generate many distinct affine mappings (21): we obtain a large value for Q in (2) and ultimately (as we shall see) poor Offline and Online RB

⁴The Elliptical Triangles of Stage 1 satisfy the parameter-independent angle conditions (50), (51): hence the Elliptical Triangles of Stage 1 remain non-singular. However, the Standard Triangles of Stage 2 are not in general of the form (41) and will not satisfy the conditions (50), (51): hence the Standard Triangles of Stage 2 may degenerate to a line—zero Jacobian—and then proceed to “penetrate” other (Standard or Elliptical) Triangles.

performance.⁵ In contrast, efficient RB triangulations are characterized by (not necessarily relatively few triangles, but rather) relatively few parametrically “dissimilar” triangles that, in turn, generate relatively few distinct affine mappings (21)—in particular, relatively few distinct $J^{\text{aff},\cdot}$ and $D^{\text{aff},\cdot}$; we obtain a smaller value for Q in (2) and ultimately better Offline and Online RB performance.

Fortunately, proper selection of the initial control point/edge data perhaps supplemented by the introduction of “artificial” regions—regions motivated by mapping considerations rather than physical/mathematical (e.g., discontinuous property/coefficient) considerations—can ensure both well-behaved/non-singular and efficient transformations; in most cases, good (and bad) choices are rather self-evident. We shall consider several examples in the next section, but we must first understand the connection between the domain decomposition and associated mapping coefficients developed here and the final affine representation of the PDE bilinear form, (2).

5.2 Bilinear Form

As already indicated, we shall consider here only the scalar case; the vector case (linear elasticity) admits an analogous treatment [61].

5.2.1 Formulation on “Original” Domain

Our problem is initially posed on the “original” domain $\Omega_o(\boldsymbol{\mu})$, which we assume realizes the Affine Geometry Precondition as described in the previous section. We shall assume for simplicity that $X_o^e(\boldsymbol{\mu}) = H_0^1(\Omega_o(\boldsymbol{\mu}))$, which corresponds to homogeneous Dirichlet boundary conditions over the entire boundary $\partial\Omega_o(\boldsymbol{\mu})$; we subsequently discuss natural (Neumann and Robin) conditions.

Given $\boldsymbol{\mu} \in \mathcal{D}$, we evaluate

$$s_o^e(\boldsymbol{\mu}) = f_o(u_o^e(\boldsymbol{\mu})),$$

where $u_o^e(\boldsymbol{\mu}) \in X_o^e(\boldsymbol{\mu})$ satisfies

$$a_o(u_o^e(\boldsymbol{\mu}), v; \boldsymbol{\mu}) = f_o(v), \quad \forall v \in X_o^e(\boldsymbol{\mu}).$$

We now place conditions on a_o and f_o such that, in conjunction with the Affine Geometry Precondition, we are ensured an affine expansion of the bilinear form.

In particular, we require that $a_o(\cdot, \cdot; \boldsymbol{\mu}): H^1(\Omega_o(\boldsymbol{\mu})) \times H^1(\Omega_o(\boldsymbol{\mu})) \rightarrow \mathbb{R}$ can be expressed as

⁵As discussed in Sects. 7.1 and 9.4, the RB Offline computational complexity scales as Q , and the RB Online computational complexity scales as Q^2 .

$$a_o(w, v; \boldsymbol{\mu})$$

$$= \sum_{\ell=1}^{L_{\text{reg}}} \int_{\bigcup_{k \in \mathbb{K}_\ell} \Omega_o^k(\boldsymbol{\mu})} \begin{bmatrix} \frac{\partial w}{\partial x_{o1}} & \frac{\partial w}{\partial x_{o2}} & w \end{bmatrix} \mathcal{K}_{o,\ell ij}(\boldsymbol{\mu}) \begin{bmatrix} \frac{\partial v}{\partial x_{o1}} \\ \frac{\partial v}{\partial x_{o2}} \\ v \end{bmatrix}, \quad (52)$$

where $\mathbf{x}_o = (x_{o1}, x_{o2})$ denotes a point in $\Omega_o(\boldsymbol{\mu})$. Here, for $1 \leq \ell \leq L_{\text{reg}}$, $\mathcal{K}_{o,\ell}: \mathcal{D} \rightarrow \mathbb{R}^{3 \times 3}$ is a given symmetric positive semidefinite matrix (we must ensure coercivity of our bilinear form): the upper 2×2 principal submatrix of $\mathcal{K}_{o,\ell}$ is the usual tensor conductivity/diffusivity; the $(3, 3)$ element of $\mathcal{K}_{o,\ell}$ represents the identity operator (“mass matrix”); and the $(3, 1)$, $(3, 2)$ (and $(1, 3)$, $(2, 3)$) elements of $\mathcal{K}_{o,\ell}$ —which we can choose here as zero thanks to our current restriction to symmetric operators—permit first derivative (or “convection”) terms. (In Sect. 11 we consider non-compliant and in particular non-symmetric bilinear forms, in which case the convection contributions are non-zero.)

Similarly, we require that $f_o: H^1(\Omega_o(\boldsymbol{\mu})) \rightarrow \mathbb{R}$ can be expressed as

$$f_o(v) = \sum_{\ell=1}^{L_{\text{reg}}} \int_{\bigcup_{k \in \mathbb{K}_\ell} \Omega_o^k(\boldsymbol{\mu})} \mathcal{F}_{o,\ell}(\boldsymbol{\mu}) v,$$

where, for $1 \leq \ell \leq L_{\text{reg}}$, $\mathcal{F}_{o,\ell}: \mathcal{D} \rightarrow \mathbb{R}$. (As we discuss below, somewhat greater generality is in fact permitted.)

5.2.2 Formulation on Reference Domain

We now apply standard techniques to transform the problem statement over the original domain to an equivalent problem statement over the reference domain: Given $\boldsymbol{\mu} \in \mathcal{D}$, we find

$$s^e(\boldsymbol{\mu}) = f(u^e(\boldsymbol{\mu})),$$

where $u^e(\boldsymbol{\mu}) \in X^e \equiv H_0^1(\Omega)$ satisfies

$$a(u^e(\boldsymbol{\mu}), v; \boldsymbol{\mu}) = f(v), \quad \forall v \in X^e.$$

We may then identify $s^e(\boldsymbol{\mu}) = s_o^e(\boldsymbol{\mu})$ and $u^e(\boldsymbol{\mu}) = u_o^e(\boldsymbol{\mu}) \circ \mathcal{T}^{\text{aff}}(\cdot; \boldsymbol{\mu})$.

The transformed bilinear form, a , can be expressed as

$$a(w, v; \boldsymbol{\mu}) = \sum_{k=1}^{K_{\text{dom}}} \int_{\Omega^k} \begin{bmatrix} \frac{\partial w}{\partial x_1} & \frac{\partial w}{\partial x_2} & w \end{bmatrix} \mathcal{K}_{ij}^k(\boldsymbol{\mu}) \begin{bmatrix} \frac{\partial v}{\partial x_1} \\ \frac{\partial v}{\partial x_2} \\ v \end{bmatrix}, \quad (53)$$

where $\mathbf{x} = (x_1, x_2)$ denotes a point in Ω . Here the $\mathcal{K}^k: \mathcal{D} \rightarrow \mathbb{R}^{3 \times 3}$, $1 \leq k \leq K_{\text{dom}}$, are symmetric positive semidefinite matrices given by

$$\mathcal{K}^k(\boldsymbol{\mu}) = J^{\text{aff},k}(\boldsymbol{\mu}) \mathcal{G}^k(\boldsymbol{\mu}) \mathcal{K}_{o,\ell}(\boldsymbol{\mu}) (\mathcal{G}^k(\boldsymbol{\mu}))^T, \quad \forall k \in \mathbb{K}_\ell, \quad (54)$$

for $1 \leq \ell \leq L_{\text{reg}}$; the $\mathcal{G}^k: \mathcal{D} \rightarrow \mathbb{R}^{3 \times 3}$, $1 \leq k \leq K_{\text{dom}}$, are given by

$$\mathcal{G}^k(\boldsymbol{\mu}) = \begin{pmatrix} D^{\text{aff},k}(\boldsymbol{\mu}) & 0 \\ 0 & 0 & 1 \end{pmatrix}; \quad (55)$$

$J^{\text{aff},k}(\boldsymbol{\mu})$ and $D^{\text{aff},k}(\boldsymbol{\mu})$, $1 \leq k \leq K_{\text{dom}}$, are given by (22) and (23), respectively; and T denotes transpose.

Similarly, the transformed linear form can be expressed as

$$f(v) = \sum_{k=1}^{K_{\text{dom}}} \int_{\Omega^k} \mathcal{F}^k(\boldsymbol{\mu}) v.$$

Here $\mathcal{F}^k: \mathcal{D} \rightarrow \mathbb{R}$, $1 \leq k \leq K_{\text{dom}}$, is given by

$$\mathcal{F}^k = J^{\text{aff},k}(\boldsymbol{\mu}) \mathcal{F}_{o,\ell}(\boldsymbol{\mu}), \quad \forall k \in \mathbb{K}_\ell, \quad 1 \leq \ell \leq L_{\text{reg}}.$$

(We recall that, for $1 \leq \ell \leq L_{\text{reg}}$, $\mathbb{K}_\ell$ is the set of subdomains associated with region ℓ : $\mathbb{K}_\ell \cap \mathbb{K}_{\ell'} = \emptyset$, $1 \leq \ell < \ell' \leq L_{\text{reg}}$, and $\bigcup_{\ell=1}^{L_{\text{reg}}} \mathbb{K}_\ell = \{1, \dots, K_{\text{dom}}\}$.)

We note that, in general, the $\mathcal{K}^k(\boldsymbol{\mu})$ and $\mathcal{F}^k(\boldsymbol{\mu})$, $1 \leq k \leq K_{\text{dom}}$, will be different for each subdomain Ω^k . The differences can arise either due to “property” variation or to geometry variation—or both. We thus require, as already indicated earlier, that the FE approximation be built upon a *subtriangulation* of the RB triangulation: discontinuities in PDE coefficients are thereby restricted to element edges to ensure (more) rapid convergence; and identification/extraction of the terms in the affine expansion (2) is more readily effected—as we now discuss.

5.2.3 Affine Form

We focus here on a , though f admits a similar treatment. We simply expand the form (53) by considering in turn each subdomain Ω^k and each entry of the diffusivity tensor $\mathcal{K}_{ij}^k$, $1 \leq i, j \leq 3$, $1 \leq k \leq K_{\text{dom}}$. Thus,

$$\begin{aligned} a(w, v; \boldsymbol{\mu}) &= \mathcal{K}_{11}^1(\boldsymbol{\mu}) \int_{\Omega^1} \frac{\partial w}{\partial x_1} \frac{\partial v}{\partial x_1} \\ &\quad + \mathcal{K}_{12}^1(\boldsymbol{\mu}) \int_{\Omega^1} \frac{\partial w}{\partial x_1} \frac{\partial v}{\partial x_2} + \dots \\ &\quad + \mathcal{K}_{33}^{K_{\text{dom}}}(\boldsymbol{\mu}) \int_{\Omega^{K_{\text{dom}}}} w v. \end{aligned} \quad (56)$$

We can then identify each component in the affine expansion: for each term in (56), the pre-factor represents $\Theta^q(\boldsymbol{\mu})$, while the integral represents a^q .

Taking into account the symmetry of the bilinear form, such that (effectively) only the (1, 1), (1, 2) (= (2, 1)), (2, 2), and (3, 3) entries of $\mathcal{K}_{o,\ell}(\boldsymbol{\mu})$ —and hence $\mathcal{K}^k(\boldsymbol{\mu})$ —must be accommodated, there are (at most) $Q = 4K$ terms

in the affine expansion. The $\Theta^q(\boldsymbol{\mu})$ are given by (for the obvious numbering scheme) $\Theta^1(\boldsymbol{\mu}) = \mathcal{K}_{11}^1(\boldsymbol{\mu})$, $\Theta^2(\boldsymbol{\mu}) = \mathcal{K}_{12}^1(\boldsymbol{\mu})$, $\dots$, $\Theta^5(\boldsymbol{\mu}) = \mathcal{K}_{11}^2(\boldsymbol{\mu})$, $\dots$, $\Theta^Q(\boldsymbol{\mu}) = \mathcal{K}_{33}^{K_{\text{dom}}}(\boldsymbol{\mu})$; the $a^q(w, v)$ are given by

$$\begin{aligned} a^1(w, v) &= \int_{\Omega^1} \frac{\partial w}{\partial x_1} \frac{\partial v}{\partial x_1}, \\ a^2(w, v) &= \int_{\Omega^1} \frac{\partial w}{\partial x_1} \frac{\partial v}{\partial x_2}, \\ &\vdots \\ a^5(w, v) &= \int_{\Omega^2} \frac{\partial w}{\partial x_1} \frac{\partial v}{\partial x_1}, \\ &\vdots \\ a^Q(w, v) &= \int_{\Omega^{K_{\text{dom}}}} w v. \end{aligned}$$

This identification constitutes a constructive proof that the Affine Geometry Precondition and the property/coefficient variation permitted by (52) do indeed yield a bilinear form which can be expressed in the requisite affine form, (2).

In fact, many of the terms in the development (56) may indeed vanish: often, many entries in $\mathcal{K}_{o,\ell}$, $1 \leq \ell \leq L_{\text{reg}}$, or in $\mathcal{G}^k$, $1 \leq k \leq K_{\text{dom}}$, will be zero. For example, for a “pure Laplacian” with isotropic diffusivity and “pure dilation” geometry transformations, we immediately reduce the number of non-zero terms in (56) to $Q = 2K_{\text{dom}}$. In practice, even in more complicated/general situations, not only will there often be many zero entries, but there will also be many “duplicate”—linearly dependent—entries: in our development (56), if (say) $\mathcal{K}_{11}^2(\boldsymbol{\mu}) = \text{Const} \mathcal{K}_{11}^1(\boldsymbol{\mu})$, we may eliminate $\Theta^5(\boldsymbol{\mu})$ and redefine

$$a^1(w, v) = \int_{\Omega^1} \frac{\partial w}{\partial x_1} \frac{\partial v}{\partial x_1} + \text{Const} \int_{\Omega^2} \frac{\partial w}{\partial x_1} \frac{\partial v}{\partial x_1},$$

thereby reducing Q . Symbolic manipulation techniques can identify and eliminate all (zero and) redundant terms in (56) to arrive at a minimal- Q affine expansion. However, “good” choice of the User-provided initial control points/edges for our RB triangulation—a choice that honors symmetries and isolates geometric variation—remains important: we describe an example below. (We re-emphasize that fewer subdomains is not important; many parametrically similar subdomains will generate “like terms” that will be automatically coalesced by symbolic economization procedures.)

We close with a discussion of generality. In fact, the conditions we provide are sufficient but not necessary. For example, we can permit affine polynomial dependence on $\mathbf{x}_o$ in both $\mathcal{K}_{o,\ell}(\mathbf{x}_o; \boldsymbol{\mu})$ and $\mathcal{F}_{o,\ell}(\mathbf{x}_o; \boldsymbol{\mu})$ and still ensure an affine development, (2); furthermore, in the absence of geometric variation (in a particular region ℓ), $\mathcal{K}_{o,\ell}(\mathbf{x}_o; \boldsymbol{\mu})$ and

Fig. 5 (a) Original domain $\Omega_o(\mu)$ with single region $\mathcal{R}_{o,1}(\mu)$ and initial control points/edges indicated, and (b) $K_{\text{dom}} = 2$ RB triangulation $\overline{\Omega} = \bigcup_{k=1}^{K_{\text{dom}}} \overline{\Omega}^k$ (the numbers refer to subdomains)

Fig. 6 (a) Original domain $\Omega_o(\mu)$ with $L_{\text{reg}} = 2$ regions $\mathcal{R}_{o,1}(\mu)$ (left triangle), and $\mathcal{R}_{o,2}(\mu)$ (right triangle) and initial control points/edges indicated, and (b) $K_{\text{dom}} = 4$ RB triangulation $\overline{\Omega} = \bigcup_{k=1}^{K_{\text{dom}}} \overline{\Omega}^k$

$\mathcal{F}_{o,\ell}(\mathbf{x}_o; \mu)$ can take on *any* “separable” form in $\mathbf{x}$, μ . However, the affine expansion (2) is by no means completely general: for more complicated data parametric dependencies, non-affine techniques [18, 53, 138] must be invoked.

Another *memento mori* is provided by inhomogeneous natural boundary conditions. Homogeneous Neumann conditions obviously pose no problem, as neither a nor f is affected in this case: we may thus consider homogeneous Neumann conditions on any straight, circular, elliptical, or generally curvy edge of $\partial\Omega$. However, for inhomogeneous Neumann conditions (which modify f) or Robin conditions (which modify a), the situation is less satisfactory: effecting the usual mapping techniques (in this case involving the “edge” Jacobian), we find that the transformed equations admit an affine form (2) *only* in the case of straight or circular edges. Again, to treat the more general case, non-affine techniques [18] must be invoked.

5.3 Computational Results

We consider here several examples of parametrized geometries that illustrate the RB triangulation, mapping, and affine decomposition (and economization) procedures described in the previous sections. We consider either homogeneous Dirichlet boundary conditions, homogeneous Neumann conditions, or inhomogeneous Neumann conditions on straight edges. We restrict our attention to the affine representation associated with the bilinear form a , (2), though

in all cases considered f also admits an affine representation.

As our first example, we consider the original domain $\Omega_o(\mu)$ shown in Fig. 5a: an obelisk with vertices $(0,0)$, $(1,1)$, $(0,\mu_1)$, $(-1,1)$. The single ($P = 1$) parameter μ_1 corresponds to the height of the upper part of the obelisk; the parameter domain is given by $\mathcal{D} = [0.5, 3]$. We choose for our reference parameter $\mu_{\text{ref}} = \mu_{1\text{ref}} = 2.5$.

In the first instance we choose a single region, hence $L_{\text{reg}} = 1$, corresponding to the entire domain shown in Fig. 5a; the User-provided control points $((0,0)$, $(1,1)$, $(0,\mu_1)$, $(-1,1)$) and control edges (sequential point pairs) are also indicated in Fig. 5a. In this case we obtain the RB triangulation (we need only perform one “sweep” of Stage 2 of the algorithm of Sect. 5.1.3) shown in Fig. 5b; clearly, at $\mu_1 = 1$ the mapping will become singular.

In the second instance we now consider the two regions, $L_{\text{reg}} = 2$, shown in Fig. 6a; the User-provided control points (in fact, still $(0,0)$, $(1,1)$, $(0,\mu_1)$, $(-1,1)$) and control edges—which now include a vertical bisector—are also indicated in Fig. 6a. We now obtain the RB triangulation—in this case we require an edge split, and hence two sweeps of Stage 2—shown in Fig. 6b; the resulting mapping is well defined, even if somewhat “distorted,” for all $\mu \in \mathcal{D}$. (We re-iterate that the regions primarily serve to identify different/discontinuous PDE coefficients; however, the regions can also serve to “stimulate” better mappings.)

Fig. 7 (a) Original domain $\Omega_o(\mu)$ with single region $\mathcal{R}_{o,1}(\mu)$ and initial control points/edges indicated, and (b) $K_{\text{dom}} = 8$ RB triangulation $\overline{\Omega} = \bigcup_{k=1}^{K_{\text{dom}}} \overline{\Omega}^k$ (the numbers refer to subdomains)

We now consider a second, only slightly more involved example, to illustrate how the choice of control points can affect not only well-posedness but also efficiency of the “economized” affine representation, (2). (The dependence of the algorithm on the User-provided control points is in fact a liability, not a feature: we would prefer complete automation. Our goal here is only to demonstrate that *without too much effort* we can guide the triangulation in the good direction.) We shall consider both (i) the Laplacian with isotropic diffusivity corresponding to $K_{o,\ell 11} = K_{o,\ell 22} = 1$ and all other entries of $K_{o,\ell}$ zero for $1 \leq \ell \leq L_{\text{reg}}$, and (ii) linear elasticity for an isotropic (either plane stress or plane strain) material [61]. We consider the original domain $\Omega_o(\mu) =]-2, 2[\times]-2, 2[\setminus [-\mu_1, \mu_1] \times [-\mu_2, \mu_2]$: a square with a variable rectangular hole; we choose a single region, hence $L_{\text{reg}} = 1$, corresponding to the entire domain. The two ($P = 2$) parameters correspond to the dimensions of the rectangular hole; the parameter domain is given by $\mathcal{D} = [0.5, 1.5] \times [0.5, 1.5]$. We choose $\mu_{\text{ref}} = (1.0, 1.0)$.

In the first instance, we choose the User-provided control points/edges as shown in Fig. 7a, which yields the $K_{\text{dom}} = 8$ RB triangulation of Ω shown in Fig. 7b. Upon economization, there are $Q = 10$ terms and $Q = 19$ terms in our affine expansion (2) for the Laplacian and elasticity problems, respectively. There is some symmetry in the RB triangulation, and the economization does reduce the number of terms in the affine expansion (say, for the Laplacian) from the maximum possible of 24 to 10. However, the RB triangulation includes many (unnecessary) shear terms, which leads to the rather large Q for this relatively simple problem.

In the second instance, we choose the User-provided control points/edges as shown in Fig. 8a, which yields the $K_{\text{dom}} = 16$ RB triangulation shown in Fig. 8b. Upon economization, there are $Q = 6$ terms and $Q = 7$ terms in our affine expansion (2) for the Laplacian and elasticity problems, respectively. In this case, the RB triangulation both retains the available symmetries *and* avoids (unnecessary) shear contributions. The control points of Fig. 8a do exhibit one drawback: the effect of the hole “propagates” to the

Fig. 8 (a) Original domain $\Omega_o(\mu)$ with single region $\mathcal{R}_{o,1}(\mu)$ and initial control points/edges indicated, and (b) $K_{\text{dom}} = 16$ RB triangulation $\overline{\Omega} = \bigcup_{k=1}^{K_{\text{dom}}} \overline{\Omega}^k$

Fig. 9 (a) Original domain $\Omega_o(\mu)$ with $L_{\text{reg}} = 2$ regions $\mathcal{R}_{o,1}(\mu)$ and $\mathcal{R}_{o,2}(\mu)$ and initial control points/edges indicated, and (b) $K_{\text{dom}} = 34$ RB triangulation $\overline{\Omega} = \bigcup_{k=1}^{K_{\text{dom}}} \overline{\Omega}^k$

outer boundary, and hence would not be efficient if $\Omega_o(\mu)$ in fact represents a *region* within a larger more complex domain; in the latter case, the control points of Fig. 7a, which localize the geometric variation, are preferred.

We now turn to an Elliptical-Edge domain. We consider a Laplacian with a “sink/reaction” term: $K_{o,\ell ij} = \delta_{ij}$, $1 \leq i, j \leq 3$, $1 \leq \ell \leq L_{\text{reg}}$; here δ_{ij} is the Kronecker delta. We consider the original domain $\Omega_o(\mu)$ with two regions: $\mathcal{R}_{o,1}(\mu)$ is the ellipse defined by $x_{o1}^2/\mu_1^2 + x_{o2}^2/\mu_2^2 < 1$; $\mathcal{R}_{o,2}(\mu) =]-2, 2[\times]-2, 2[\setminus \overline{\mathcal{R}_{o,1}(\mu)}$. In actual practice, we would introduce different (parameter-dependent) material properties/discontinuous coefficients in the PDE in the two regions: for simplicity here, we consider just the two ($P = 2$) geometric parameters which define the ellipse major/minor axes; the parameter domain is given by $\mathcal{D} = [0.5, 1.5] \times [0.5, 1.5]$. We choose $\mu_{\text{ref}} = (1.0, 1.0)$.

We select the User-provided control points/edges shown in Fig. 9a; we obtain the $K_{\text{dom}} = 34$ RB triangulation of Ω shown in Fig. 9b. (Note the many control points/edges in Fig. 9b relative to Fig. 9a: we perform many edge splits in both Stage 1 and Stage 2.) Note that all the subdomains inside the ellipse are Inwards Elliptical Triangles, and the subdomains $k = 16, 18, 20, 21, 24, 25, 29, 31$ are all Outwards Elliptical Triangles.

Fig. 10 (a) Original domain $\Omega_o(\mu)$ with single region $\mathcal{R}_{o,1}(\mu)$ and initial control points/edges indicated, and (b) $K_{\text{dom}} = 8$ RB triangulation $\overline{\Omega} = \bigcup_{k=1}^{K_{\text{dom}}} \overline{\Omega}^k$

In this case, we obtain $Q = 15$, which is quite modest; furthermore, the geometric feature is “localized” in the sense that the control points on the outer square are parameter-independent. We observe that many subdomains does not necessarily imply large Q ; in this case, all the Inwards and Outwards Elliptical Triangles share the same (anisotropic) dilation mapping, and hence generate similar mapping coefficients that can be coalesced in the final affine representation.

As our last example, we consider a more general Curvy-Edge domain with a single region. We consider an isotropic Laplacian corresponding to $\mathcal{K}_{o,111} = \mathcal{K}_{o,122} = 1$ (all other entries of $\mathcal{K}_{o,1}$ are set to zero). We consider the original domain $\Omega_o(\mu)$ given by $\{(x_{o1}, x_{o2}) \mid 0 < x_{o1} < 1, -1 < x_{o2} < \mu_1 \cos \pi x_{o1}\}$. The single ($P = 1$) parameter represents the amplitude of the cosinusoidal top boundary; the parameter domain is given by $\mathcal{D} = [\frac{1}{6}, \frac{1}{2}]$. We choose $\mu_{\text{ref}} = 1/3$.

We select the User-provided control points and edges shown in Fig. 10a. Note for the Curvy Edge associated with control points $(0, \mu_1)$ and $(1/2, 0)$ we choose center $(0, 0)$ and hence our parameterization is given by

$$\begin{pmatrix} x_{o1} \\ x_{o2} \end{pmatrix} = \begin{bmatrix} 1 & 0 \\ 0 & \mu_1 \end{bmatrix} \begin{bmatrix} t \\ \cos \pi t \end{bmatrix} \quad (57)$$

for $0 \leq t \leq 1/2$; for the Curvy Edge associated with control points $(0, 1/2)$ and $(1, -\mu_1)$ we choose center $(1, 0)$ and hence our parameterization is given by

$$\begin{pmatrix} x_{o1} \\ x_{o2} \end{pmatrix} = \begin{bmatrix} 1 \\ 0 \end{bmatrix} + \begin{bmatrix} 1 & 0 \\ 0 & \mu_1 \end{bmatrix} \begin{bmatrix} t - 1 \\ \cos \pi t \end{bmatrix} \quad (58)$$

for $1/2 \leq t \leq 1$. These choices yield the $K_{\text{dom}} = 8$ RB triangulation of Ω shown in Fig. 10b. We note that subdomains $k = 5, 6$ correspond to Outwards and Inwards Curvy Triangles, respectively.

In this case we obtain $Q = 9$, again rather modest. This particular problem is more obviously treated by a “stretch” mapping (which is not affine); however, for our purposes we clearly prefer the piecewise-affine mapping presented here. We emphasize that arbitrary Curvy Edge domains will not,

Fig. 11 Thermal Block problem for $B_1 = B_2 = 3$

in general, permit affine treatment: the Curvy Edges must admit a (parameter-independent) concave/convex decomposition and appropriate centers in order for our approach to be directly applicable.

6 Working Examples

6.1 Scalar Problems

We consider two “working” examples. The first shall serve as a vehicle for the convergence theory; the second is intended to illustrate an application. We note that in all cases we provide the formulation for the “exact” problem (superscript e); the FE approximation is then derived from the exact statement (and RB triangulation) following the procedures described earlier. Note also that all problems are presented in non-dimensional form.

6.1.1 Thermal Block

We consider heat conduction in a square domain. The square comprises $B_1 \times B_2$ blocks: each block is a different region with different thermal conductivity; the geometry is depicted in Fig. 11. Inhomogeneous Neumann (non-zero flux) boundary conditions are imposed on (so as to match notation in Fig. 11) Γ_{base} ; homogeneous Dirichlet (temperature) conditions are imposed on Γ_{top} ; and homogeneous (zero flux) Neumann conditions are imposed on the two sides. The output of interest is the average temperature over (so as to match notation in Fig. 11) Γ_{base} [6, 7].

The parameters are then the conductivities in the first $B_1 B_2 - 1$ blocks (with the blocks numbered as shown in Fig. 11); note in our non-dimensionalization the conductivity of the last block, which serves for normalization, is unity. Hence $P = B_1 B_2 - 1$ and $\mu = (\mu_1, \dots, \mu_P)$, where μ_p is the conductivity of block p ; the parameter domain is then given by $\mathcal{D} = [\mu^{\min}, \mu^{\max}]^P$ with $\mu^{\min} = 1/\sqrt{\mu_r}$, $\mu^{\max} = \sqrt{\mu_r}$ for $\mu_r = 100$ (hence $\mu^{\max}/\mu^{\min} = 100$).

We immediately recover our abstract statement of (1): we identify $X^c \equiv \{v \in H^1(\Omega) \mid v|_{\Gamma_{\text{top}}} = 0\}$, which imposes the Dirichlet conditions;

$$a(w, v; \boldsymbol{\mu}) = \sum_{i=1}^P \mu_i \int_{\mathcal{R}_i} \nabla w \cdot \nabla v + \int_{\mathcal{R}_{P+1}} \nabla w \cdot \nabla v, \quad (59)$$

which represents the Laplacian and homogeneous Neumann conditions (as well as internal flux continuity conditions); and

$$f(v) \equiv f^{\text{Neu}}(v) \equiv \int_{\Gamma_{\text{bot}}} v, \quad (60)$$

which imposes the inhomogeneous Neumann conditions. Here

$$\overline{\Omega} = \bigcup_{i=1}^{P+1} \overline{\mathcal{R}}_i,$$

where the $\mathcal{R}_i$, $i = 1, \dots, P+1$, correspond to the regions associated with the respective blocks/conductivities, as indicated in Fig. 11.

The problem is readily demonstrated to be coercive, symmetric, and compliant (the Neumann boundary condition functional indeed corresponds to the average temperature functional). The problem is also affine in parameter, (59): in this case no geometric transformations are required, though the regions still serve to accommodate the discontinuous PDE coefficients; we directly observe that $Q = P+1$ with $\Theta^q(\boldsymbol{\mu}) = \mu_q$, $1 \leq q \leq P$, $\Theta^{P+1} = 1$, and

$$a^q(w, v) = \int_{\mathcal{R}_q} \nabla w \cdot \nabla v, \quad 1 \leq q \leq P+1.$$

(Note for this problem f is indeed independent of $\boldsymbol{\mu}$.)

This problem shall serve to illustrate the convergence rate of the RB discretization, both for $P = 1$ in which we can compare with available *a priori* theory, and for $P = 8$ corresponding to “many parameters.”

6.1.2 Inviscid Flow: Added Mass

We consider inviscid incompressible flow induced by small motions of a square body in a “basin.” In this case, there is only a single region: the physical (flow) domain $\Omega_o(\boldsymbol{\mu})$ —the basin/the body—is depicted in Fig. 12a. The governing equation is Laplace’s equation for the pressure. Inhomogeneous Neumann boundary conditions on the pressure (corresponding to inhomogeneous conditions on the normal velocity) are imposed on $\Gamma_{o,1}^\pm$; homogeneous Dirichlet (“free surface”) conditions on the pressure are imposed on $\Gamma_{o,s}$; and homogeneous Neumann conditions on the pressure (zero normal flow) are imposed on all other boundaries. The output of interest is the classical added mass [94].

Fig. 12 (a) Original domain $\Omega_o(\boldsymbol{\mu})$ with single region $\mathcal{R}_{o,1}(\boldsymbol{\mu})$ and initial control points/edges indicated, and (b) $K_{\text{dom}} = 22$ RB triangulation $\overline{\Omega} = \bigcup_{k=1}^{K_{\text{dom}}} \overline{\Omega}^k$

The three parameters—in this case, all the parameters are *geometric*—are indicated in Fig. 12a. Hence $P = 3$ and $\boldsymbol{\mu} = (\mu_1, \mu_2, \mu_3)$: μ_1 is the lateral extent of the basin, μ_2 is the minor axis of the elliptical arc describing the bottom of the basin, and μ_3 is the horizontal location of the (center of the) square block; the parameter domain is given by $\mathcal{D} = [1.5, 3] \times [0.5, 1.5] \times [-0.35, -0.35]$. We choose $\boldsymbol{\mu}_{\text{ref}} = (2, 1, 0)$ which in turn defines our reference domain $\Omega = \Omega_o(\boldsymbol{\mu}_{\text{ref}})$.

We again realize the abstraction of (1), however in this case—and in fact, in all cases involving geometric variation—it shall prove more convenient to define the constituents with respect to the “original” (parameter-dependent) domain. In particular, we identify $X_o^c(\boldsymbol{\mu}) \equiv \{v \in H^1(\Omega_o(\boldsymbol{\mu})) \mid v|_{\Gamma_{o,s}} = 0\}$, which imposes the Dirichlet conditions;

$$a_o(w, v; \boldsymbol{\mu}) = \int_{\Omega_o(\boldsymbol{\mu})} \nabla w \cdot \nabla v,$$

which represents the Laplacian and homogeneous Neumann/zero normal flow conditions; and

$$f_o(v) \equiv \int_{\Gamma_{o,1}^+} v - \int_{\Gamma_{o,1}^-} v,$$

which imposes a x_{o1} -velocity of unity on both $\Gamma_{o,1}^+$ and $\Gamma_{o,1}^-$. The problem is clearly coercive, symmetric, and compliant (the normal velocity functional also corresponds to the pressure force—and hence added mass—functional).

We may then apply our methods of Sect. 5.1: for the control points/edges indicated in Fig. 12a, the algorithm of Sect. 5.1.3 yields the $K_{\text{dom}} = 22$ RB triangulation of Ω depicted in Fig. 12b; we then construct the associated affine mappings according to the recipe of Sect. 5.1.2; we next effect the re-formulation on the reference domain, as described in Sect. 5.2.2; finally, we extract and coalesce/economize the terms in the affine expansion, (2), following the process defined in Sect. 5.2.3. In this particular case, we obtain an affine expansion (2) with $Q = 34$ terms. (For this problem, f does not depend on μ , as the edges $\Gamma_{o,1}^\pm$ do not depend on the parameters.) Note this proliferation of terms from $P = 3$ to $Q = 34$ —typical of more geometrically complex problems—adversely affects Online performance (which depends on Q), and in practice limits the scope of our approach; without economization, the situation would be much worse.

We note that the current added-mass example also illustrates the many query context: in practice, the output might appear as a position-dependent (μ_3 -dependent) coefficient in the acceleration term of the dynamical equation for the position of the body; numerical integration of these equations of motion could thus require repeated evaluation of the added-mass output. (Applications can be conceived at both the large ocean engineering scale or the small bio-engineering scale.) Reduced basis methods are very efficient for treatment of this class of problems.

6.2 Vector Problems: Linear Elasticity

We consider two “working” examples. The first (in $d = 3$ space dimensions) shall serve as a vehicle for quantifying computational performance; the second (in $d = 2$ dimensions) is intended to illustrate an application. We note that in all cases we provide the formulation for the “exact” problem (superscript c); the FE approximation is then derived from the exact statement (and RB triangulation) following the procedures described earlier—suitably extended to the vector case. Note also that all problems are presented in non-dimensional form.

6.2.1 Elastic Block 3D

We consider linear elasticity in a cubic domain $]0, 3[^3$ with a centered parallelepiped of different material—an inhomogeneous inclusion—and of variable dimension. Each material

Fig. 13 Domain (and two regions) for Elastic Block 3D

is associated with a different region; we depict the geometry $\Omega_o(\mu)$ and two regions in Fig. 13. The governing equation is the equation of linear elasticity (isotropic material) for the displacement. Inhomogeneous Neumann boundary conditions on the displacement (corresponding to uniform normal stress) are imposed on $\Gamma_{o,T}$; homogeneous Dirichlet conditions on the displacement (“clamped”) are imposed on $\Gamma_{o,D}$; and homogeneous Neumann conditions on the displacement (zero normal and tangential stress) are imposed on all other boundaries. The output of interest is the integrated normal displacement over $\Gamma_{o,T}$ [150].

There are four parameters. The first three parameters are geometric, while the fourth parameter relates to material properties. Hence $P = 4$ and $\mu = (\mu_1, \mu_2, \mu_3, \mu_4)$: μ_1, μ_2 , and μ_3 are the extent of the parallelepiped inclusion (with center $(3/2, 3/2, 3/2)$) in the x_{o1}, x_{o2} , and x_{o3} coordinates, while μ_4 is the Young’s modulus of the inclusion relative to the Young’s modulus of the “continuous phase” (the two materials are assumed to have the same Poisson ratio $\nu^{\text{Po}} = 0.3$); the parameter domain is given by $\mathcal{D} = [0.5, 2]^3 \times [0.1, 10]$. We choose $\mu_{\text{ref}} = (1, 1, 1, \cdot)$.

We again realize the abstraction of (1). As before, it is most convenient to define the constituents with respect to the “original” (parameter-dependent) domain. In particular, we identify $X_o^c(\mu) \equiv \{v \in (H^1(\Omega_o(\mu)))^3 | v|_{\Gamma_{o,D}} = 0\}$, which imposes the clamped conditions;

$$a_o(w, v; \mu) = \int_{\mathcal{R}_{o,1}(\mu)} \frac{\partial w_i}{\partial x_{oj}} C_{1ijmn}(\mu) \frac{\partial v_m}{\partial x_{on}} + \int_{\mathcal{R}_{o,2}(\mu)} \frac{\partial w_i}{\partial x_{oj}} C_{2ijmn}(\mu) \frac{\partial v_m}{\partial x_{on}},$$

where for $\ell = 1, 2$,

$$C_{\ell ijmn}(\mu) = \lambda_\ell^1(\mu) \delta_{ij} \delta_{mn} + \lambda_\ell^2(\mu) (\delta_{im} \delta_{jn} + \delta_{in} \delta_{jm}),$$

with δ_{ij} the Kronecker delta symbol and

$$\lambda_1^1(\mu) = \frac{\mu_4 \nu^{\text{Po}}}{(1 + \nu^{\text{Po}})(1 - 2\nu^{\text{Po}})}, \quad \lambda_1^2(\mu) = \frac{\mu_4}{2(\nu^{\text{Po}} + 1)},$$

$$\lambda_2^1(\mu) = \frac{\nu^{\text{Po}}}{(1 + \nu^{\text{Po}})(1 - 2\nu^{\text{Po}})}, \quad \lambda_2^2(\mu) = \frac{1}{2(\nu^{\text{Po}} + 1)},$$

the non-dimensional Lamé constants for an isotropic elastic material; finally,

$$f_o(v; \boldsymbol{\mu}) \equiv \int_{\Gamma_{o,T}} v_1,$$

which imposes the normal stress of unity on $\Gamma_{o,T}$. The problem is clearly coercive, symmetric, and compliant (the normal stress functional also corresponds to the integrated displacement functional).

We now apply (the vector version of) our methods of Sect. 5, but in this case “by hand”: we choose 27 cubic subdomains as our uniform RB “triangulation” of Ω —the inclusion region $\mathcal{R}_{o,1}(\boldsymbol{\mu})$ (respectively, the continuous phase region $\mathcal{R}_{o,2}(\boldsymbol{\mu})$) comprises the “middle” subdomain $\ell = 14$ (respectively, the remaining subdomains $\ell = 1, \dots, 13, 15, \dots, 27$); we then construct the associated affine mappings by inspection—anisotropic dilations; we next effect the re-formulation on the reference domain, as described in Sect. 5.2.2 (suitably extended to the vector case); finally, we extract and coalesce/economize the affine expansion, (2), following the process defined in Sect. 5.2.3. In this particular example, we obtain an affine expansion for a , (2), with $Q = 48$ terms. (For this problem, f does depend on $\boldsymbol{\mu}$: $f(v; \boldsymbol{\mu})$ admits a 4-term affine expansion.) As is typically the case—in particular for rectilinear problems—the $\Theta^q(\boldsymbol{\mu})$ are low-order rational polynomials in $\boldsymbol{\mu}$.

This example shall serve to (quantitatively) assess computational performance of the RB method relative to the FE method. This problem is also quite similar to the cell problems that must be solved in the many-query non-periodic homogenization context; see [26] for application of the RB approach to non-periodic homogenization theory.

6.2.2 Elastic Crack: Stress Intensity Factor

We consider linear elasticity for a crack emanating from a hole in “Mode I” tension; we consider only one quarter of the domain thanks to symmetry of geometry and loading. We depict the geometry (and single region) $\Omega_o(\boldsymbol{\mu})$ in Fig. 14; note the crack corresponds to boundary segment $\Gamma_{o,1}$. The governing equation is the equation of linear elasticity (isotropic material) for the displacement. Inhomogeneous Neumann boundary conditions on the displacement (corresponding to uniform normal stress) are imposed on $\Gamma_{o,4}$; mixed homogeneous Dirichlet/Neumann conditions on the displacement (“symmetry”) are imposed on $\Gamma_{o,2}$ and $\Gamma_{o,5}$; homogeneous Neumann conditions on the displacement (zero normal and tangential stress) are imposed on all other boundaries (including the crack and hole). The ultimate output of interest is the Stress Intensity Factor for the crack, which we shall derive from an intermediate (compliant) energy output by application of the virtual crack extension approach [109].

Fig. 14 Elastic Crack: (a) Original domain $\Omega_o(\boldsymbol{\mu})$ (single region) with initial control points/edges indicated, and (b) Boundary segments on $\partial\Omega_o(\boldsymbol{\mu})$

There are two parameters; both parameters are geometric in origin, as shown in Fig. 14. (In our scaling, the Young’s modulus of the material is unity; the Poisson ratio is fixed at $\nu^{\text{Po}} = 0.3$.) Hence $P = 2$ and $\boldsymbol{\mu} = (\mu_1, \mu_2)$: μ_1 is the length of the crack (as measured from the hole), while μ_2 is the radius of the hole; the parameter domain is given by $\mathcal{D} = [0.1, 0.3] \times [0.1, 0.5]$. We now choose $\boldsymbol{\mu}_{\text{ref}} = [0.2, 0.3]$ which in turn defines our reference domain $\Omega = \Omega_o(\boldsymbol{\mu}_{\text{ref}})$. (Note that as μ_1 varies for μ_2 fixed the length of the crack changes but the radius of the hole is invariant; this shall be important in evaluating the Stress Intensity Factor.)

We again realize the abstraction of (1). As before, it is most convenient to define the constituents with respect to the “original” (parameter-dependent) domain. In particular, we identify $X_o^e(\boldsymbol{\mu}) \equiv \{v \in (H^1(\Omega_o(\boldsymbol{\mu})))^2 \mid v_1|_{\Gamma_{o,5}} = 0, v_2|_{\Gamma_{o,2}} = 0\}$, which imposes the Dirichlet component of the symmetry conditions;

$$a_o(w, v; \boldsymbol{\mu}) = \int_{\Omega_o(\boldsymbol{\mu})} \frac{\partial w_i}{\partial x_{oj}} C_{ijkl} \frac{\partial v_k}{\partial x_{ol}},$$

where

$$C_{ijkl} = \lambda^1 \delta_{ij} \delta_{kl} + \lambda^2 (\delta_{ik} \delta_{jl} + \delta_{il} \delta_{jk}),$$

with δ_{ij} the Kronecker delta symbol and

$$\lambda^1 = \frac{\nu^{\text{Po}}}{(1 - \nu^{\text{Po}})^2}, \quad \lambda^2 = \frac{1}{2(\nu^{\text{Po}} + 1)}$$

the non-dimensional Lamé constants for an isotropic linear elastic material in plane stress ($\nu^{\text{Po}} = 0.3$); finally,

$$f_o(v; \boldsymbol{\mu}) \equiv \int_{\Gamma_{o,4}} v_1,$$

which imposes the normal stress of unity on $\Gamma_{o,4}$. The problem is clearly coercive, symmetric, and—by construction—compliant.

Fig. 15 Elastic Crack:
 $K_{\text{dom}} = 6$ RB triangulation
 $\overline{\Omega} = \bigcup_{k=1}^{K_{\text{dom}}} \overline{\Omega}^k$

We may then apply (the vector version of) our methods of Sect. 5: for the control points/edges indicated in Fig. 14a (recall there is a single region in this example), the algorithm of Sect. 5.1.3 yields the $K_{\text{dom}} = 6$ RB triangulation of Ω depicted in Fig. 15; we then construct the associated affine mappings according to the recipe of Sect. 5.1.2; we next effect the re-formulation on the reference domain, as described in Sect. 5.2.2 (suitably extended to the vector case); finally, we extract and coalesce/economize the affine expansion, (2), following the process defined in Sect. 5.2.3. In this particular case, we obtain an affine expansion for a , (2), with $Q = 25$ terms. (For this problem, f does not depend on μ .) We again note the modest but non insignificant amplification from $P = 2$ (geometric) parameters to $Q = 25$ terms in the affine expansion.

As already indicated, the (transformed) output $s^e(\mu) = f(u^e(\mu))$ is, in this example, only an intermediate result “on the way” to the Stress Intensity Factor (SIF). In particular, the virtual crack extension method [109]—based on the energy formulation of the Stress Intensity Factor—yields the Energy Release Rate (ERR^e), G^e , as

$$G^e(\mu) = - \left(\frac{\partial s^e(\mu)}{\partial \mu_1} \right)$$

in terms of which (in our non-dimensionalization) the Stress Intensity Factor can then be expressed as $\text{SIF}^e(\mu) = \sqrt{G^e(\mu)}$. We can thus construct, for suitably small $\delta\mu_1$, a finite-difference approximation to the ERR^e as

$$\widehat{G}^e(\mu) = - \left(\frac{s^e(\mu + \delta\mu_1) - s^e(\mu)}{\delta\mu_1} \right)$$

and subsequently to the SIF^e as $\widehat{\text{SIF}}^e(\mu) = \sqrt{\widehat{G}^e(\mu)}$. We shall later develop first a FE approximation and subsequently a corresponding RB approximation—and associated RB error bound—for $\widehat{G}^e(\mu)$.

The ERR and SIF often serve in the real-time or many-query contexts. In the real-time context, we might require “in the field” evaluation of fracture for given (varying) environmental conditions [20]. In the many-query context, we

might require evaluation of crack growth from Paris’s law in which the SIF—itsself a function a crack length—is the crucial term.

7 The Reduced Basis Method

7.1 Reduced Basis Approximation

7.1.1 RB Spaces

We assume that we are given a FE approximation space of dimension $\mathcal{N}$, $X^{\mathcal{N}}$. (We shall subsequently confirm that all our results are computationally and mathematically stable as $\mathcal{N} \rightarrow \infty$; however, in order to define a particular reduced basis space, we may consider a fixed $\mathcal{N}$.) We then introduce, given a positive integer N_{max} , an associated sequence of (what shall ultimately be reduced basis) approximation spaces: for $N = 1, \dots, N_{\text{max}}$, $X_N^{\mathcal{N}}$ is a N -dimensional subspace of $X^{\mathcal{N}}$; we further suppose that

$$X_1^{\mathcal{N}} \subset X_2^{\mathcal{N}} \subset \dots \subset X_{N_{\text{max}}}^{\mathcal{N}} \subset X^{\mathcal{N}}. \quad (61)$$

As we shall see, the nested or *hierarchical* condition (61) is important in ensuring (memory) efficiency of the resulting reduced basis approximation. (In one instance—which we will clearly mark—we will for purely theoretical purposes consider a non-hierarchical sequence of spaces.)

We recall from Sect. 2 that there are several classical RB proposals—Taylor, Lagrange, and Hermite spaces—as well as several more recent contenders—such as POD spaces. All of these spaces “focus” in one fashion or another on the low-dimensional, smooth parametric manifold, $\mathcal{M}^{\mathcal{N}} = \{u(\mu) \mid \mu \in \mathcal{D}\}$, already identified in Opportunity I of Sect. 2: indeed, the central role of this parametric manifold is the defining (albeit somewhat imprecisely defined here) characteristic of “proper” reduced basis spaces. Much of what we present—in particular, all the material of this section related to optimality, discrete equations, conditioning, and Offline-Online procedures, and all the material of Sect. 9 related to *a posteriori* error estimation—shall be relevant to any of these reduced basis spaces/approximations.

However, some of what we shall present, in particular related to sampling strategies in Sect. 7.2, is restricted to the particular reduced basis space which shall be our primary focus: the Lagrange reduced basis spaces [119], which we shall denote by $(X_N^{\mathcal{N}} \Rightarrow) W_N^{\mathcal{N}}$. In order to define a (hierarchical) sequence of Lagrange spaces $W_N^{\mathcal{N}}$, $1 \leq N \leq N_{\text{max}}$, we first introduce a “master set” of parameter points $\mu^n \in \mathcal{D}$, $1 \leq n \leq N_{\text{max}}$. We then define, for given $N \in \{1, \dots, N_{\text{max}}\}$, the Lagrange parameter samples

$$S_N = \{\mu^1, \dots, \mu^N\}, \quad (62)$$

Fig. 16 The “snapshots” $u^{\mathcal{N}}(\mu^n)$, $1 \leq n \leq N$, on the parametric manifold $\mathcal{M}^{\mathcal{N}}$

and associated Lagrange RB spaces

$$W_N^{\mathcal{N}} = \text{span}\{u^{\mathcal{N}}(\mu^n), 1 \leq n \leq N\}. \quad (63)$$

We observe that, by construction, these Lagrange spaces $X_N^{\mathcal{N}} = W_N^{\mathcal{N}}$ satisfy (61): the samples (62) are nested— $S_1 = \{\mu^1\} \subset S_2 = \{\mu^1, \mu^2\} \subset \dots \subset S_{N_{\max}}$; the Lagrange RB spaces (63) are hierarchical— $W_1^{\mathcal{N}} = \text{span}\{u^{\mathcal{N}}(\mu^1)\} \subset W_2^{\mathcal{N}} = \text{span}\{u^{\mathcal{N}}(\mu^1), u^{\mathcal{N}}(\mu^2)\} \subset \dots \subset W_{N_{\max}}^{\mathcal{N}}$.

The $u^{\mathcal{N}}(\mu^n)$, $1 \leq n \leq N_{\max}$, are often referred to as “snapshots” of the parametric manifold $\mathcal{M}^{\mathcal{N}}$. For reasons that will become clear subsequently, we shall denote these snapshots more precisely as “retained snapshots.” We depict the retained snapshots graphically in Fig. 16. It is clear that, if indeed the manifold is low-dimensional and smooth (a point we return to later), then we would expect to well approximate any member of the manifold—any solution $u^{\mathcal{N}}(\mu)$ for some $\mu \in \mathcal{D}$ —in terms of relatively few retained snapshots. However, we must first ensure that we can choose a good combination of the available retained snapshots (Sect. 7.1.2), that we can represent the retained snapshots in a stable RB basis (Sect. 7.1.2), that we can efficiently obtain the associated RB basis coefficients (Sect. 7.1.3), and finally that we can choose our retained snapshots—in essence, the parameter sample $S_{N_{\max}}$ —optimally (Sect. 7.2.2). (Note only the last item is specific to Lagrange RB spaces.)

7.1.2 Galerkin Projection

For our particular class of equations, Galerkin projection is arguably the best approach. Given $\mu \in \mathcal{D}$, evaluate (recalling our compliance assumption)

$$s_N^{\mathcal{N}}(\mu) = f(u_N^{\mathcal{N}}(\mu)),$$

where $u_N^{\mathcal{N}}(\mu) \in X_N^{\mathcal{N}} \subset X^{\mathcal{N}}$ (or more precisely, $u_{X_N^{\mathcal{N}}}^{\mathcal{N}}(\mu) \in X_N^{\mathcal{N}}$) satisfies

$$a(u_N^{\mathcal{N}}(\mu), v; \mu) = f(v), \quad \forall v \in X_N^{\mathcal{N}}. \quad (64)$$

We emphasize that our ultimate interest is the output prediction: the field variable serves as an intermediary. (We discuss below a simpler but less rigorous and typically less efficient alternative: direct approximation/ interpolation of the input-output relation $\mu \rightarrow s(\mu)$.)

We immediately obtain the classical optimality result in the energy norm (5):

$$\|u^{\mathcal{N}}(\mu) - u_N^{\mathcal{N}}(\mu)\|_{\mu} \leq \inf_{w \in X_N^{\mathcal{N}}} \|u^{\mathcal{N}}(\mu) - w\|_{\mu}; \quad (65)$$

in the energy norm, the Galerkin procedure automatically selects the *best* combination of snapshots. (Similar but sub-optimal results obtain in other equivalent norms.) It is also readily derived that

$$s^{\mathcal{N}}(\mu) - s_N^{\mathcal{N}}(\mu) \equiv \|u^{\mathcal{N}}(\mu) - u_N^{\mathcal{N}}(\mu)\|_{\mu}^2; \quad (66)$$

the output—our quantity of interest—converges as the “square” of the energy error. (Although this latter result depends critically on the compliance assumption, extension *via* adjoint approximations to the non-compliant case is possible; we discuss this further in Sect. 11.)

We now consider the discrete equations associated with the Galerkin approximation (64). We must first choose an appropriate basis for our space: incorrect choice of the RB basis can lead to very poorly conditioned systems; this is immediately apparent in the Lagrange case—if $W_N^{\mathcal{N}}$ provides rapid convergence then, by construction, the snapshots of (63) will be increasingly co-linear as N increases. Towards this end, we apply the Gram-Schmidt process [88, 149] in the $(\cdot, \cdot)_X$ inner product to our snapshots $u^{\mathcal{N}}(\mu^n)$, $1 \leq n \leq N_{\max}$, to obtain mutually orthonormal functions $\zeta_n^{\mathcal{N}}$, $1 \leq n \leq N_{\max}$: $(\zeta_n^{\mathcal{N}}, \zeta_m^{\mathcal{N}})_X = \delta_{nm}$, $1 \leq n, m \leq N_{\max}$, where δ_{nm} is the Kronecker delta symbol. We then choose the sets $\{\zeta_n^{\mathcal{N}}\}_{n=1, \dots, N}$ as our bases for $W_N^{\mathcal{N}}$, $1 \leq N \leq N_{\max}$.

We now insert

$$u_N^{\mathcal{N}}(\mu) = \sum_{m=1}^N u_{Nm}^{\mathcal{N}}(\mu) \zeta_m^{\mathcal{N}}, \quad (67)$$

and $v = \zeta_n^{\mathcal{N}}$, $1 \leq n \leq N$, into (64) to obtain the RB “stiffness” equations

$$\sum_{m=1}^N a(\zeta_m^{\mathcal{N}}, \zeta_n^{\mathcal{N}}; \mu) u_{Nm}^{\mathcal{N}}(\mu) = f(\zeta_n^{\mathcal{N}}), \quad 1 \leq n \leq N, \quad (68)$$

for the RB coefficients $u_{Nm}^{\mathcal{N}}(\mu)$, $1 \leq m \leq N$; we can subsequently evaluate the RB output prediction as

$$s_N^{\mathcal{N}}(\mu) = \sum_{m=1}^N u_{Nm}^{\mathcal{N}}(\mu) f(\zeta_m^{\mathcal{N}}). \quad (69)$$

It can be readily proven [112] that the condition number of the matrix $a(\zeta_m^{\mathcal{N}}, \zeta_n^{\mathcal{N}}; \mu)$, $1 \leq n, m \leq N$, is bounded by $\gamma^e(\mu)/\alpha^e(\mu)$ independent of N and $\mathcal{N}$.

7.1.3 Offline-Online Procedure

The system (68) is nominally of small size: a set of N linear algebraic equations in N unknowns. However, the formation of the stiffness matrix, and indeed the load vector, involves entities $\zeta_n^{\mathcal{N}}$, $1 \leq n \leq N$, associated with our $\mathcal{N}$ -dimensional FE approximation space. If we must invoke FE fields in order to form the RB stiffness matrix *for each new value of μ* the marginal cost per input-output evaluation $\mu \rightarrow s_N(\mu)$ will remain unacceptably large.

Fortunately, we can appeal to affine parameter dependence to construct very efficient Offline-Online procedures, as we now discuss. In particular, we note that our system (68) can be expressed, thanks to (2), as

$$\sum_{m=1}^N \left(\sum_{q=1}^Q \Theta^q(\mu) a^q(\zeta_m^{\mathcal{N}}, \zeta_n^{\mathcal{N}}) \right) u_{Nm}^{\mathcal{N}}(\mu) = f(\zeta_n^{\mathcal{N}}), \quad 1 \leq n \leq N. \quad (70)$$

We observe that the $\zeta_n^{\mathcal{N}}$ are now isolated in terms that are independent of μ and hence that can be *pre-computed* in an Offline-Online procedure.

In the Offline stage, we first compute the $u^{\mathcal{N}}(\mu^n)$, $1 \leq n \leq N_{\max}$, and subsequently the $\zeta_n^{\mathcal{N}}$, $1 \leq n \leq N_{\max}$; we then form and store the

$$f(\zeta_n^{\mathcal{N}}), \quad 1 \leq n \leq N_{\max}, \quad (71)$$

and

$$a^q(\zeta_m^{\mathcal{N}}, \zeta_n^{\mathcal{N}}), \quad 1 \leq n, m \leq N_{\max}, \quad 1 \leq q \leq Q. \quad (72)$$

The Offline operation count depends on $N_{\max}$, Q , and $\mathcal{N}$.

In the Online (or “deployed”) stage, we retrieve (72) to form

$$\sum_{q=1}^Q \Theta^q(\mu) a^q(\zeta_m^{\mathcal{N}}, \zeta_n^{\mathcal{N}}), \quad 1 \leq n, m \leq N; \quad (73)$$

we solve the resulting $N \times N$ stiffness system (70) to obtain the $u_{Nm}^{\mathcal{N}}(\mu)$, $1 \leq m \leq N$; and finally we access (71) to evaluate the output (69). The Online operation count is $O(QN^2)$ to perform the sum (73), $O(N^3)$ to invert (70)—note that the RB stiffness matrix is full, and finally $O(N)$ to effect the inner product (69). The Online storage (the data archived in the Offline stage) is—thanks to our hierarchical condition (61)—only $O(QN_{\max}^2) + O(N_{\max})$: for any given N , we may extract the necessary RB $N \times N$ matrices (respectively, N -vectors) as principal submatrices (respectively, principal subvectors) of the corresponding $N_{\max} \times N_{\max}$ (respectively, $N_{\max}$) quantities.

The Online cost (operation count and storage)—and hence marginal cost and also asymptotic average cost—to

evaluate $\mu \rightarrow s_N^{\mathcal{N}}(\mu)$ is thus independent of $\mathcal{N}$. The implications are two-fold: first, *if* N is indeed small, we will achieve very fast response in the real-time and many-query contexts; second, we may choose $\mathcal{N}$ very conservatively—to effectively eliminate the error between the exact and FE predictions—without adversely affecting the Online (marginal) cost. We now turn to a more detailed discussion of sampling and (in Sect. 8) convergence in order to understand how, and to a certain extent why, we can achieve high accuracy for N independent of $\mathcal{N}$ and indeed $N \ll \mathcal{N}$.

7.2 “Sampling” Strategies

We first indicate a few preliminaries. We then turn to two examples of sampling strategies.

We shall denote by $\mathcal{E}$ a finite sample of points in $\mathcal{D}$. These “test” samples $\mathcal{E}$ shall serve as surrogates for $\mathcal{D}$ in the calculation and presentation of errors (and, in Sect. 9, error bounds and effectivities) over the parameter domain. Typically these samples are chosen by Monte Carlo methods with respect to a uniform or log-uniform density. For brevity we will often report neither the specific distribution, nor the precise number, of points in $\mathcal{E}$: we always ensure (and empirically confirm) that $\mathcal{E}$ is sufficiently large that the reported results are insensitive to further refinement of the parameter sample.

Given a function $y: \mathcal{D} \rightarrow \mathbb{R}$, we define

$$\|y\|_{L^\infty(\mathcal{E})} \equiv \max_{\mu \in \mathcal{E}} |y(\mu)|,$$

and

$$\|y\|_{L^p(\mathcal{E})} \equiv \left(|\mathcal{E}|^{-1} \sum_{\mu \in \mathcal{E}} |y(\mu)|^p \right)^{1/p}.$$

Given a function $z: \mathcal{D} \rightarrow X^{\mathcal{N}}$ (or X^e), we then define

$$\|z\|_{L^\infty(\mathcal{E}; X)} \equiv \max_{\mu \in \mathcal{E}} \|z(\mu)\|_X,$$

and

$$\|z\|_{L^p(\mathcal{E}; X)} \equiv \left(|\mathcal{E}|^{-1} \sum_{\mu \in \mathcal{E}} \|z(\mu)\|_X^p \right)^{1/p}.$$

Here $|\mathcal{E}|$ denotes the cardinality of (the finite number of elements in) the test sample $\mathcal{E}$.

We denote the particular samples which shall serve to select our RB space—or “train” our RB approximation—by $\mathcal{E}_{\text{train}}$. The cardinality of $\mathcal{E}_{\text{train}}$ will be denoted $|\mathcal{E}_{\text{train}}| = n_{\text{train}}$. We note that although the “test” samples $\mathcal{E}$ serve primarily to understand and assess the quality of the RB approximation and *a posteriori* error estimators, the “train” samples $\mathcal{E}_{\text{train}}$ serve to *generate* the RB approximation. The choice of n_{train} and $\mathcal{E}_{\text{train}}$ thus have important Offline and Online computational implications.

7.2.1 POD RB Spaces: $X_N^{\mathcal{N}\text{POD}}$

These POD RB spaces can be defined very simply [37, 58, 72, 73, 78]: Given $\mathcal{E}_{\text{train}}$, the $X_N^{\mathcal{N}\text{POD}}$, $1 \leq N \leq N_{\text{max}}$, are the solution to the optimization problem

$$X_N^{\mathcal{N}\text{POD}} = \arg \inf_{X_N^{\mathcal{N}} \subset \text{span}\{u^{\mathcal{N}}(\mu) | \mu \in \mathcal{E}_{\text{train}}\}} \|u^{\mathcal{N}} - \Pi_{X_N^{\mathcal{N}}} u^{\mathcal{N}}\|_{L^2(\mathcal{E}_{\text{train}}; X)}, \quad (74)$$

where $\Pi_{X_N^{\mathcal{N}}}: X^{\mathcal{N}} \rightarrow X_N^{\mathcal{N}}$ refers to the projection in the X inner product. We shall denote the functions $u^{\mathcal{N}}(\mu)$, $\mu \in \mathcal{E}_{\text{train}}$, as “candidate snapshots.” Although certainly not immediately apparent, it can be readily demonstrated that the POD spaces are indeed hierarchical, (61).

It is a well-known but nevertheless remarkable fact [58] that the optimization problem (74)—ostensibly of combinatorial complexity in n_{train} —can be reduced to the solution of an $n_{\text{train}} \times n_{\text{train}}$ correlation-matrix eigenproblem—of algebraic complexity in n_{train} .⁶ However, the POD approach is nevertheless very expensive: first, and most importantly, we must compute all n_{train} (FE) candidate snapshots—and perform $n_{\text{train}}^2/2$ candidate snapshot-candidate snapshot X inner products—in order to form the correlation matrix; second, we must solve for the first N_{max} eigenvalues/eigenproblems of the very large— $n_{\text{train}} \times n_{\text{train}}$ —correlation matrix. In practice, this can severely limit n_{train} , which in turn (in particular for higher parameter dimensions P) can result in poor spaces—slow convergence, and large errors over $\mathcal{D} \setminus \mathcal{E}_{\text{train}}$.

7.2.2 Greedy Lagrange RB Spaces: $W_N^{\mathcal{N}\text{Greedy}}$ and $W_N^{\mathcal{N}\text{Greedy, en}}$

We now develop a sample strategy particular to RB Lagrange spaces. The method can be viewed as a “heuristic” (more precisely, sub-optimal) solution to the $L^\infty(\mathcal{E}_{\text{train}}; X)$ optimization problem analogous to the $L^2(\mathcal{E}_{\text{train}}; X)$ POD optimization problem (74).

We are given $\mathcal{E}_{\text{train}}$ and N_{max} , as well as $S_1 = \{\mu^1\}$, $W_1^{\mathcal{N}\text{Greedy}} = \text{span}\{u^{\mathcal{N}}(\mu^1)\}$. (In actual practice we may set N_{max} either directly, or indirectly through a prescribed error tolerance.) Then, for $N = 2, \dots, N_{\text{max}}$, we find

$$\mu^N = \arg \max_{\mu \in \mathcal{E}_{\text{train}}} \Delta_{N-1}(\mu),$$

⁶The eigenfunctions associated with the N_{max} largest eigenvalues—with minor postprocessing—directly provide the requisite orthonormal basis functions [112]. Note, however, that each basis function in general will combine all n_{train} candidate snapshots $u^{\mathcal{N}}(\mu)$, $\mu \in \mathcal{E}_{\text{train}}$, and thus the $X_N^{\mathcal{N}\text{POD}}$ will not correspond to Lagrange RB spaces.

set $S_N = S_{N-1} \cup \mu^N$, and update $W_N^{\mathcal{N}\text{Greedy}} = W_{N-1}^{\mathcal{N}\text{Greedy}} + \text{span}\{u^{\mathcal{N}}(\mu^N)\}$. As we shall describe in detail in Sect. 9, $\Delta_N(\mu)$ is a sharp, (asymptotically) *inexpensive a posteriori* error bound for $\|u^{\mathcal{N}}(\mu) - u_{W_N^{\mathcal{N}\text{Greedy}}}^{\mathcal{N}}(\mu)\|_X$.

Roughly, at iteration N the greedy algorithm appends to the *retained* snapshots that particular candidate snapshot—over all candidate snapshots $u^{\mathcal{N}}(\mu)$, $\mu \in \mathcal{E}_{\text{train}}$ —which is (predicted by the *a posteriori* error bound to be the) least well approximated by (the RB prediction associated to) $W_{N-1}^{\mathcal{N}\text{Greedy}}$. Typically, this greedy approach will be much less expensive than the POD approach: in the greedy approach, we need compute only the N —typically very few—FE *retained* snapshots; in the POD approach, we must compute all n_{train} —typically/desirably very many—FE *candidate* snapshots.⁷ Furthermore we shall observe empirically, although at present we have no proof, that the short-horizon “greedy” selection criterion in fact engenders rapidly convergent RB approximation spaces—typically as good as the POD spaces. (The temporal evolution case is quite different: the greedy approach [52, 55] can encounter difficulties best treated by incorporating elements of the POD selection process [60].)

We can develop an analogous greedy procedure in the energy norm; the energy norm is particularly relevant in the “compliant” case, since the error in the energy norm is directly related to the error in the output (see Sect. 7.1.2). As before, we are given $\mathcal{E}_{\text{train}}$ and N_{max} , as well as $S_1 = \{\mu^1\}$, $W_1^{\mathcal{N}\text{Greedy, en}} = \text{span}\{u^{\mathcal{N}}(\mu^1)\}$. Then, for $N = 2, \dots, N_{\text{max}}$, we find

$$\mu^N = \arg \max_{\mu \in \mathcal{E}_{\text{train}}} (\omega_N(\mu))^{-1} \Delta_{N-1}^{\text{en}}(\mu),$$

set $S_N = S_{N-1} \cup \mu^N$, and update $W_N^{\mathcal{N}\text{Greedy, en}} = W_{N-1}^{\mathcal{N}\text{Greedy, en}} + \text{span}\{u^{\mathcal{N}}(\mu^N)\}$. As described in detail in Sect. 9.3, $\Delta_N^{\text{en}}(\mu)$ is a sharp, (asymptotically) *inexpensive a posteriori* error bound for $\|u^{\mathcal{N}}(\mu) - u_{W_N^{\mathcal{N}\text{Greedy, en}}}^{\mathcal{N}}(\mu)\|_{\mu}$. The prefactor $\omega_N(\mu)$ is typically chosen either as unity or (as in all cases in this paper) $\|u_{W_N^{\mathcal{N}\text{Greedy, en}}}^{\mathcal{N}}(\mu)\|_{\mu}$; for the latter choice, the greedy selects on relative energy error (the square of which yields the relative output error).

We now turn to theoretical and computational evidence that this greedy selection process generates spaces which are, if not optimal, at least very good.

⁷Clearly the accuracy and cost of the *a posteriori* error estimator $\Delta_N(\mu)$ are crucial to the success of the greedy algorithm: we provide theoretical results for the effectivity and asymptotic average operation count in Sect. 9. Detailed operation counts for the aggregate greedy procedure, including all *a posteriori* estimation contributions, are provided in [112].

Fig. 17 Thermal Block problem: $B_1 = 2$, $B_2 = 1$

8 Convergence of RB Approximations

8.1 Single Parameter Case: $P = 1$

8.1.1 Model Problem

We shall consider the Thermal Block problem of Sect. 6.1.1 for the case in which $B_1 = 2$, $B_2 = 1$, as shown in Fig. 17. The governing equations are then given by (59), (60) for two blocks/regions $\mathcal{R}_1$ and $\mathcal{R}_2$, the single parameter $\mu = \mu = \mu_1$ representing the conductivity of region $\mathcal{R}_1$ (the conductivity of region $\mathcal{R}_2$ is unity), and the parameter domain $\mathcal{D} = [\mu^{\min}, \mu^{\max}] \equiv [1/\sqrt{\mu_r}, \sqrt{\mu_r}]$ for $\mu_r = 100$; the associated affine expansion (2) now comprises only $Q = 2$ terms. The FE discretization is then described by (3): for most of our results, we consider a FE approximation with $\mathcal{N} = 1024$ degrees of freedom; as we describe in Sect. 8.2.1, our results are largely insensitive to $\mathcal{N}$ for sufficiently large $\mathcal{N}$ (and any fixed N).

Our RB approximation is then given by (64). We shall consider three choices for our space: the hierarchical POD spaces $X_N^{\mathcal{N}} = X_N^{\mathcal{N}\text{POD}}$, $1 \leq N \leq N_{\max}$, generated by the algorithm described in Sect. 7.2.1; the hierarchical Lagrange spaces $X_N^{\mathcal{N}} = W_N^{\mathcal{N}\text{Greedy,en}}$, $1 \leq N \leq N_{\max}$, generated by the (energy version) of the greedy algorithm described in Sect. 7.2.2; and finally, purely for theoretical purposes, the non-hierarchical Lagrange spaces $W_N^{\mathcal{N}\text{nh,ln}}$, $1 \leq N \leq N_{\max}$, given by

$$W_N^{\mathcal{N}\text{nh,ln}} = \text{span}\{u^{\mathcal{N}}(\mu_N^n), 1 \leq n \leq N\}, \quad (75)$$

for the (non-nested, or only occasionally nested) parameter points given by

$$\mu_N^n = \mu^{\min} \exp \left\{ \frac{n-1}{N-1} \ln \left(\frac{\mu^{\max}}{\mu^{\min}} \right) \right\}, \quad 1 \leq n \leq N, \quad 2 \leq N \leq N_{\max}. \quad (76)$$

We denote the corresponding RB approximations by $u_{X_N^{\mathcal{N}\text{POD}}}^{\mathcal{N}}$ or simply $u_N^{\mathcal{N}\text{POD}}$, $u_{W_N^{\mathcal{N}\text{Greedy,en}}}^{\mathcal{N}}$ or simply $u_N^{\mathcal{N}\text{Greedy,en}}$, and $u_{W_N^{\mathcal{N}\text{nh,ln}}}^{\mathcal{N}}$ or simply $u_N^{\mathcal{N}\text{nh,ln}}$, respectively.

The *a priori* theory described below suggests that the spaces (75)—which we shall denote “equi-ln” spaces—contain certainly optimality properties, though we shall observe that our more automatic sample selection procedures do just as well (and perhaps even better for larger N). We note the analysis presented here in fact is relevant to a large class of single parameter coercive problems [85, 87, 112].

8.1.2 Equi-ln Spaces: A Priori Theory

We present from [85, 87, 112] an *a priori* theory for RB approximations associated with the specific non-hierarchical equi-ln spaces (75). (In fact, the parameters need only be quasi-uniform in $\ln \mu$ [85, 87].) In particular, for the model problem of Sect. 8.1.1, given general data f (of which f^{Neu} of (60) is a particular example), we obtain [112]

Proposition 1 For any $N \geq N_{\text{crit}}$, and $\forall \mu \in \mathcal{D}$,

$$\frac{\|u^{\mathcal{N}}(\mu) - u_N^{\mathcal{N}\text{nh,ln}}(\mu)\|_{\mu}}{\|u^{\mathcal{N}}(\mu)\|_{\mu}} \leq \exp \left\{ - \frac{N-1}{N_{\text{crit}}-1} \right\}, \quad (77)$$

where

$$N_{\text{crit}} = 1 + [2e \ln \mu_r]_+.$$

(Here $[\cdot]_+$ returns the smallest integer greater than or equal to its real argument.)

Note we can directly derive from Proposition 1 and (66) a bound on the relative (compliant) output error.

The proof is a “parameter” version of the standard (finite element) variational arguments. In particular, we first invoke (65); we then take as our surrogate for the best fit a high-order polynomial interpolant in the parameter μ (in fact, in the mapped parameter $\ln \mu$) of $u^{\mathcal{N}}(\mu)$; we next apply the standard Lagrange interpolant remainder formula; finally, we appeal to an eigenfunction expansion to bound the parametric (sensitivity) derivatives and optimize the order of the polynomial interpolant. The result is not particularly sharp: independent tests [112] indicate that the interpolant in parameter is in fact reasonably accurate (though of course sub-optimal in the energy norm relative to the Galerkin projection); however, our estimate for the parametric derivatives and hence interpolant error is rather crude—thus yielding a rather pessimistic bound.

Nevertheless, we can draw several important conclusions from Proposition 1. First, the notion of RB approximation as a “parameter domain” analogue to FE approximation in the “physical domain” is not only qualitatively but also quantitatively relevant. Second, and related, RB convergence relies on smoothness in parameter but not on spatial regularity—recall that Proposition 1 is valid for any $f \in (X^e)'$. Third, the RB convergence rate—more precisely, our bound (77)

for the RB convergence rate—does not depend on $\mathcal{N}$. (The actual convergence rate does in fact depend on the underlying FE approximation space, however for any fixed N this dependence vanishes as $\mathcal{N}$ increases.) Fourth, the RB convergence rate depends quite weakly on the extent of the parameter domain: the exponent in the convergence rate decreases only *logarithmically* with μ_r ; the RB approximation is indeed global. Fifth, and finally, the RB approximation can converge very quickly.

8.1.3 Greedy Lagrange Spaces: $X^{\mathcal{N}} = W_N^{\mathcal{N}\text{Greedy,en}}$

In actual practice, we invoke the hierarchical, automatically generated, spaces $W_N^{\mathcal{N}\text{Greedy,en}}$ —RB approximations associated with Lagrange spaces generated by the (energy version of the) greedy algorithm described in Sect. 7.2.2—not the non-hierarchical, “hand-crafted” equi-ln spaces $W_N^{\mathcal{N}\text{nh,ln}}$.

We present in Fig. 18 the relative energy error in the $L^\infty(\mathcal{E})$ norm as a function of N for the non-hierarchical equi-ln spaces associated with our *a priori* theory of Sect. 8.1.2,

$$\max_{\mu \in \mathcal{E}} \left(\frac{\|u^{\mathcal{N}}(\mu) - u_{W_N^{\mathcal{N}\text{nh,ln}}}^{\mathcal{N}}(\mu)\|_\mu}{\|u^{\mathcal{N}}(\mu)\|_\mu} \right), \quad (78)$$

and for our greedy spaces,

$$\max_{\mu \in \mathcal{E}} \left(\frac{\|u^{\mathcal{N}}(\mu) - u_{W_N^{\mathcal{N}\text{Greedy,en}}}^{\mathcal{N}}(\mu)\|_\mu}{\|u^{\mathcal{N}}(\mu)\|_\mu} \right). \quad (79)$$

(Here $\mathcal{E}$ is a suitably large log-random sample over $\mathcal{D}$.) We observe, first, that both spaces provide extremely fast convergence: the relative energy error is $1\text{E}-6$ (and hence from (66) the relative output error is $1\text{E}-12$) already for $N = 6$. Second, we observe that the greedy spaces are initially very slightly worse than the equi-ln spaces, but for “larger” N the greedy and equi-ln results are indistinguishable.

The nearly identical convergence results could signal either that there are many possibly very different spaces all of which provide very rapid convergence, or that the equi-ln and greedy spaces are in fact quite similar. In Fig. 19 we present (in log-lin format) (i) the sample points $\mu_{N=6}^n$ for the equi-log space, (76), and (ii) the sample points $\mu^{n'}$, $1 \leq n' \leq 6$, for the hierarchical space generated by the energy version of the greedy algorithm. (The greedy points are plotted in ascending order for easier comparison.) We observe that the greedy algorithm selects points quite close to the equi-ln distribution: this result can serve (somewhat circularly) as evidence of the optimality of either the equi-ln distribution or the greedy distribution. In fact, it appears that

Fig. 18 Thermal Block problem for $B_1 = 2$, $B_2 = 1$: $L^\infty(\mathcal{E})$ relative energy error as a function of N for the spaces $W_N^{\mathcal{N}\text{nh,ln}}$ (*) and $W_N^{\mathcal{N}\text{Greedy,en}}$ (o) ((78) and (79), respectively)

Fig. 19 Thermal Block problem for $B_1 = 2$, $B_2 = 1$: Sample points $\mu_{N=6}^n$ for $W_N^{\mathcal{N}\text{nh,ln}}$ (*) and $\mu^{n'}$ for $W_{N=6}^{\mathcal{N}\text{Greedy,en}}$ (o)

the greedy distribution is tending to a Chebyshev in $\ln \mu$ distribution, with clustering near the endpoints of the parameter domain: the Chebyshev-ln distribution perhaps performs better than the equi-ln distribution.

8.1.4 POD Spaces: $X^{\mathcal{N}} = X_N^{\mathcal{N}\text{POD}}$

We recall that the POD *truly* minimizes the *projection error* in $L^2(\mathcal{E}_{\text{train}}; X)$, whereas the greedy algorithm *heuristically* minimizes the *RB error bound* in $L^\infty(\mathcal{E}_{\text{train}}; X)$. (Note in this subsection, in order to more directly compare with the POD result, we consider the X -norm rather than energy version of the greedy algorithm of Sect. 7.2.2.) We might thus expect that, if the greedy heuristic is performing well, then the POD RB error should be smaller (respectively, larger) than the greedy RB error in the $L^2(\mathcal{E}_{\text{train}}; X)$ norm (respectively $L^\infty(\mathcal{E}_{\text{train}}; X)$ norm); equivalently, we might conclude that if the POD RB error is in fact smaller (respectively, larger) than the greedy RB error in the $L^2(\mathcal{E}_{\text{train}}; X)$ norm (respectively $L^\infty(\mathcal{E}_{\text{train}}; X)$ norm), then the greedy heuristic is indeed performing well.

Fig. 20 Thermal Block problem for $B_1 = 2$, $B_2 = 1$: (a) $\|u^\mathcal{N} - u_N^{\mathcal{N}^{\text{POD}}}\|_{L^2(\mathcal{E}; X)}$ (x) and $\|u^\mathcal{N} - u_N^{\mathcal{N}^{\text{Greedy}}}\|_{L^2(\mathcal{E}; X)}$ (◇) as a function of N ; (b) $\|u^\mathcal{N} - u_N^{\mathcal{N}^{\text{POD}}}\|_{L^\infty(\mathcal{E}; X)}$ (o) and $\|u^\mathcal{N} - u_N^{\mathcal{N}^{\text{Greedy}}}\|_{L^\infty(\mathcal{E}; X)}$ (*) as a function of N

We present in Fig. 20a the $L^2(\mathcal{E}; X)$ errors for the RB approximations associated with the POD and greedy Lagrange spaces, $\|u^\mathcal{N} - u_N^{\mathcal{N}^{\text{POD}}}\|_{L^2(\mathcal{E}; X)}$ and $\|u^\mathcal{N} - u_N^{\mathcal{N}^{\text{Greedy}}}\|_{L^2(\mathcal{E}; X)}$, respectively, as a function of N ; we present in Fig. 20b the $L^\infty(\mathcal{E}; X)$ errors for the RB approximations associated with the POD and greedy Lagrange spaces, $\|u^\mathcal{N} - u_N^{\mathcal{N}^{\text{POD}}}\|_{L^\infty(\mathcal{E}; X)}$ and $\|u^\mathcal{N} - u_N^{\mathcal{N}^{\text{Greedy}}}\|_{L^\infty(\mathcal{E}; X)}$, respectively, as a function of N . In both cases, $\mathcal{E}$ is a suitably large test sample. We observe, as anticipated, that the POD spaces perform a bit better than the greedy spaces in the $L^2(\mathcal{E}; X)$ norm, while the greedy spaces perform very slightly better than the POD spaces in the (somewhat stronger) $L^\infty(\mathcal{E}; X)$ norm. Presuming that this performance is “generic,” and taking into account the much lower cost of the greedy algorithm in particular for larger n_{train} , we conclude that the greedy approach—despite the very short-horizon heuristic—is a computationally attractive alternative to the POD. (Further improvement to the greedy algorithm is possible: [32, 33] replaces the simple enumeration maximization of the error bound described in Sect. 7.2.2 with a much more efficient gradient-based search procedure.)

Fig. 21 Thermal Block problem for $B_1 = 2$, $B_2 = 1$: Relative output error as a function of N for the output interpolant (*) and RB output prediction (◇) ((80) and (81), respectively)

8.1.5 Output Interpolation

It might appear that the reduced basis approach is full of sound and fury (and complicated algorithms), signifying little computational advantage. Indeed, there is a much simpler alternative that avoids the field-variable “intermediary”: direct interpolation (in parameter) of the output $s^\mathcal{N}(\mu)$.

We compare in Fig. 21 the N -point Gauss-Lobatto Chebyshev polynomial interpolant in $\ln \mu$ [112, 125] of $s^\mathcal{N}(\mu)$, denoted $s_{N,\text{int}}^\mathcal{N}(\mu)$, to the RB output associated with the greedy (energy) spaces, $s_N^{\mathcal{N}^{\text{Greedy},\text{en}}}(\mu)$: we present

$$\max_{\mu \in \mathcal{E}} \frac{|s^\mathcal{N}(\mu) - s_{N,\text{int}}^\mathcal{N}(\mu)|}{s^\mathcal{N}(\mu)} \quad (80)$$

and

$$\max_{\mu \in \mathcal{E}} \frac{|s^\mathcal{N}(\mu) - s_N^{\mathcal{N}^{\text{Greedy},\text{en}}}(\mu)|}{s^\mathcal{N}(\mu)} \quad (81)$$

for a suitably fine test sample $\mathcal{E}$. We observe that the Chebyshev output interpolant is in fact more accurate than the RB output, and certainly more efficient—evaluation $\mu \rightarrow s_{N,\text{int}}^\mathcal{N}(\mu)$ requires $O(N^2)$ operations whereas (Online) evaluation $\mu \rightarrow s_N^{\mathcal{N}^{\text{Greedy},\text{en}}}(\mu)$ requires $O(N^3)$ operations.

However, in higher parameter dimensions ($P > 1$), it will not be possible to develop an efficient output interpolant. First, the obvious tensor product sample point distributions are very inefficient, and optimal scattered data alternatives are difficult to find. Even for $P = 1$, we must be astute or at least awake in first effecting the $\ln \mu$ transform before appealing to the Chebyshev distribution; in contrast, the greedy algorithm *automatically* uncovers the good choice. Second, the generation of good interpolation procedures for general data in $\mathbb{R}^P$ is not a simple task; in contrast, the Galerkin projection automatically chooses the best combination of retained snapshots, at least in the energy norm. Third, and finally, we note that it is rather difficult to obtain rigorous

Fig. 22 Schematic of the RB approximation process

and sharp *a posteriori* bounds for the output interpolant error $|s^N(\mu) - s_{N,\text{int}}^N(\mu)|$; in contrast, as we shall describe in Sect. 9, we can readily develop rigorous and sharp bounds for the RB output error, $|s^N(\mu) - s_N^N(\mu)|$.

8.2 Convergence: $P > 1$

As already highlighted in the previous section, the key to RB convergence in higher parameter dimensions is the role of the PDE and field variable in determining appropriate sample points and combinations of (retained) snapshots. We illustrate the process schematically in Fig. 22: the RB field approximation, *via* the PDE residual, yields the error bound; the error bound, in turn, facilitates the greedy selection of good sample points; the Galerkin projection then provides the optimal combination of retained snapshots; finally, the RB output approximation—application of the output functional—inherits the good properties of the RB field variable, (66). As we shall observe shortly, the Greedy sample points are quite non-intuitive and very far from the obvious (and inefficient) tensor-product recipes. (In general, however, we do observe clustering near the boundaries of $\mathcal{D}$, as we might expect from classical approximation theory.)

The computational success of the (implicit) complicated process described by Fig. 22 is in fact also responsible for the failure, at present, to provide any general *a priori* convergence theory: we can not construct a best-fit surrogate since *a priori* we can neither predict an efficient sample nor construct an effective parametric interpolant.

We can anticipate that for a good set of points (and from Galerkin a good combination of retained snapshots), we should obtain rapid convergence: as already identified in Opportunity I of Sect. 2, $u^N(\mu) \in X^N$ —the field we wish to approximate by the RB method—perforce resides on the parametrically induced low-dimensional, *smooth* manifold $\mathcal{M}^N = \{u^N(\mu) \mid \mu \in \mathcal{D}\}$ ⁸; the essential role of parametric smoothness—already exploited in Sect. 8.1 for the

⁸As regards smoothness, we note that for $\Theta^q \in C^\infty(\mathcal{D})$, $1 \leq q \leq Q$, it can be shown under our coercivity, continuity, and affine hypotheses of Sect. 4 that $\|D^\sigma u^N(\mu)\|_X$ is bounded by a constant $C^{|\sigma|}$ (independent of N) for all $\mu \in \mathcal{D}$; here $D^\sigma u^N(\mu)$ refers to the σ multi-index derivative of u^N with respect to μ .

Fig. 23 Thermal Block problem for $B_1 = B_2 = 3$: (upper bound for the) $L^\infty(\mathcal{E}_{\text{train}})$ relative energy error, (82), as a function of N

single-parameter case—was identified in the first theoretical (typically asymptotic) analyses of Taylor RB spaces [47] and subsequently Lagrange RB spaces [119]. However, it is not obvious that a good set of points (and hence a good Lagrange RB space) must exist, and even less obvious that the greedy algorithm will identify this good set (or even a slightly less good set) of points. At present, we have only empirical evidence for particular examples, as we now describe. Note in all cases we consider RB approximations $u_N^N = u_N^{\text{Greedy, en}}$ associated with the spaces $W_N^N = W_N^{\text{Greedy, en}}$.

8.2.1 Thermal Block: $B_1 = B_2 = 3$

We first consider the Thermal Block problem introduced in Sect. 6.1.1 and (for $B_1 = B_2 = 3$) depicted in Fig. 11; note that now there are $P = 8$ parameters. For problems in one parameter, it is simple to choose “sufficiently rich” test and train samples; in the current situation, with $P = 8$ parameters, it is very difficult to afford—even with the greedy algorithm—a sufficiently rich test/train sample. We choose for $\mathcal{E}_{\text{train}}$ a log-uniform random sample of size $n_{\text{train}} = 5000$; note that, in any event, we always have recourse to our *a posteriori* error bounds for any new $\mu \in \mathcal{D}$ visited Online.

We present in Fig. 23 the error measure

$$\max_{\mu \in \mathcal{E}_{\text{train}}} \left(\frac{\Delta_N^{\text{en}}(\mu)}{\|u_N^N(\mu)\|_\mu} \right) \quad (82)$$

as a function of N ; note that $\Delta_N^{\text{en}}(\mu)$ is an upper bound for $\|u^N(\mu) - u_N^N(\mu)\|_\mu$ and $\|u_N^N(\mu)\|_\mu$ is a lower bound for $\|u^N(\mu)\|_\mu$, and hence $\Delta_N^{\text{en}}(\mu)/\|u_N^N(\mu)\|_\mu$ is in fact an upper bound for the relative error in the energy norm. (We note that, implicitly, $\mathcal{E}_{\text{train}}$ serves both as train and, in Fig. 23, test

Fig. 24 Thermal Block problem for $B_1 = B_2 = 3$: (upper bound for the) $L^\infty(\mathcal{E}_{\text{train}})$ relative energy error, (82), as a function of N ; $\mathcal{N} = 137$ (dotted), $\mathcal{N} = 453$ (dashed), and $\mathcal{N} = 661, 1737, 2545, 6808$ (all quite similar)

sample; since $n_{\text{train}} \gg N$, this statistical crime will have little sensible effect on our assessment.)

We observe in Fig. 23 that, despite the rather large parameter dimension, and extensive parameter domain, the RB approximation still converges very rapidly with N . We achieve an accuracy of $1\text{E}-2$ in the relative energy error (and hence an accuracy of $1\text{E}-4$ in the relative output error) with only $N \approx 40$ degrees of freedom. Clearly, if we compare Fig. 23 for $P = 8$ to Fig. 18 for $P = 1$, there is an effect, even a significant effect, due to the increased parameter dimension. However, the decrease in convergence rate with the increase in P , at least for this particular—admittedly rather simple—problem is rather modest. This example is proof of concept that the reduced basis method can in fact treat problems with “many” parameters.

This problem can also serve to verify the weak effect of $X^{\mathcal{N}}$. The results of Fig. 23 are obtained for a rather modest FE discretization corresponding to $\mathcal{N} = 661$ degrees of freedom. We now repeat in Fig. 24 the calculations of Fig. 23 but for several different $\mathcal{N}$ from $\mathcal{N} = 137$ to $\mathcal{N} = 6808$. We note that for this particular problem the dimension of $\text{span}\{\mathcal{M}^{\mathcal{N}}\}$ is in fact $\text{Const} \times \sqrt{\mathcal{N}}$ —only the boundaries of the blocks “count”—with $\text{Const} \approx 7$. We observe, as expected, that if N approaches the dimension of $\text{span}\{\mathcal{M}^{\mathcal{N}}\}$, then the RB error will “artificially” plummet to zero. However, for any fixed N , the RB convergence rate is largely insensitive to $\mathcal{N}$ as $\mathcal{N} \rightarrow \infty$. It follows that the reduced basis can replicate an *arbitrarily rich* finite element approximation to *any desired accuracy* for N independent of $\mathcal{N}$.

8.2.2 Inviscid Flow: Added Mass

We next consider the inviscid flow added mass example of Sect. 6.1.2. For this problem, with only $P = 3$ parameters, we can now visualize the greedy-predicted sample. We show in Fig. 25 the sample $S_{N_{\text{max}}}$ obtained by application of the (energy version of the) greedy algorithm of Sect. 7.2.2 for

Fig. 25 Inviscid flow example: greedy (energy version) sample $S_{N_{\text{max}}}$; note the value of μ_2 ($0.5 \leq \mu_2 \leq 1.5$) is proportional to the radius of the circle

Fig. 26 Inviscid flow example: $\max_{\mu \in \mathcal{E}_{\text{train}}} (\Delta_N^{\text{en}}(\mu) / \|u_N^{\mathcal{N}}(\mu)\|_{\mu})$ as a function of N for the Lagrange RB approximations associated with the sample of Fig. 25; here $\mathcal{E}_{\text{train}}$ is a log-uniform random sample of size $n_{\text{train}} = 3000$

$\mathcal{E}_{\text{train}}$ a log-uniform random sample of size $n_{\text{train}} = 3000$. Clearly, the point distribution is very far from tensor-product in form: there is some clustering near the boundaries of the parameter domain, however the interior of the domain is very sparsely populated. We also note that the sample $S_{N_{\text{max}}}$ reflects the particular problem of interest, as would be expected from the “adaptive” greedy procedure: the densest clustering of points is near $\mu_1 = 1.5$ and for $\mu_3 = \pm 0.35$, corresponding to regions of $\mathcal{D}$ in which the parametric sensitivity is largest.

We plot in Fig. 26 $\max_{\mu \in \mathcal{E}_{\text{train}}} (\Delta_N^{\text{en}}(\mu) / \|u_N^{\mathcal{N}}(\mu)\|_{\mu})$ for the Lagrange RB approximations associated with the sample of Fig. 25. We again observe very rapid, exponential convergence.

8.2.3 Elastic Crack

We next consider the two dimensional elasticity crack problem described in Sect. 6.2.2. We recall that this problem is characterized by $P = 2$ parameters, in which the first parameter μ_1 corresponds to the length of the crack. As described in Sect. 6.2.2, the (now FE) Energy Release Rate (ERR) $G^{\mathcal{N}}(\mu)$ can be calculated from the compliant output

$s^{\mathcal{N}}(\boldsymbol{\mu})$ as $G^{\mathcal{N}}(\boldsymbol{\mu}) = -(\partial s^{\mathcal{N}}(\boldsymbol{\mu})/\partial \mu_1)$; we consider a finite-difference approximation $\widehat{G}^{\mathcal{N}}(\boldsymbol{\mu})$ to $G^{\mathcal{N}}(\boldsymbol{\mu})$ given by

$$\widehat{G}^{\mathcal{N}}(\boldsymbol{\mu}) = \frac{s^{\mathcal{N}}(\boldsymbol{\mu}) - s^{\mathcal{N}}(\boldsymbol{\mu} + \delta\mu_1)}{\delta\mu_1} \quad (83)$$

for some given (small) $\delta\mu_1$. For our purposes here, we shall take $\widehat{G}^{\mathcal{N}}(\boldsymbol{\mu})$ as “equivalent” to $G^{\mathcal{N}}(\boldsymbol{\mu})$.⁹

We next define our reduced basis ERR approximation as

$$\widehat{G}_N^{\mathcal{N}}(\boldsymbol{\mu}) = \frac{s_N^{\mathcal{N}}(\boldsymbol{\mu}) - s_N^{\mathcal{N}}(\boldsymbol{\mu} + \delta\mu_1)}{\delta\mu_1}, \quad (84)$$

where $s_N^{\mathcal{N}}(\boldsymbol{\mu})$ is our usual RB compliant output. We then define

$$\Delta_{\widehat{G}}^{\mathcal{N}}(\boldsymbol{\mu}) = \frac{\Delta_N^s(\boldsymbol{\mu}) + \Delta_N^s(\boldsymbol{\mu} + \delta\mu_1)}{\delta\mu_1}, \quad (85)$$

where $\Delta_N^s(\boldsymbol{\mu})$ is the RB compliant output error bound defined in Sect. 9. It readily follows that $\Delta_{\widehat{G}}^{\mathcal{N}}(\boldsymbol{\mu})$ is a rigorous bound for the RB ERR prediction: $|\widehat{G}^{\mathcal{N}}(\boldsymbol{\mu}) - \widehat{G}_N^{\mathcal{N}}(\boldsymbol{\mu})| \leq \Delta_{\widehat{G}}^{\mathcal{N}}(\boldsymbol{\mu})$. We note that the choice of our virtual crack displacement value $\delta\mu_1$ is not arbitrary. It is clear that $\delta\mu_1$ must to be chosen as small as possible to provide an accurate finite-difference approximation—to ensure that $\widehat{G}_N^{\mathcal{N}}(\boldsymbol{\mu})$ is indeed sufficiently close to $G_N^{\mathcal{N}}(\boldsymbol{\mu})$; however, if the value $\delta\mu_1$ is too small, our ERR error bound (85) will suffer due to amplification of the RB contribution.

We first construct a reduced basis approximation to the compliant output $s(\boldsymbol{\mu})$. We choose for $\mathcal{E}_{\text{train}}$ a log-uniform random sample of size $n_{\text{train}} = 3000$. We again observe rapid convergence: we achieve a relative (compliant) output error of $1\text{E}-5$ for $N \approx 70$; note that in this example our parameter range represents a relatively large geometric variation. We next choose $\delta\mu_1 = 0.01$ for our ERR prediction $\widehat{G}^{\mathcal{N}}(\boldsymbol{\mu})$ of (83), and hence RB ERR prediction $\widehat{G}_N^{\mathcal{N}}(\boldsymbol{\mu})$ of (84).¹⁰

We plot in Fig. 27 $\widehat{G}_N^{\mathcal{N}}(\boldsymbol{\mu})$ and the error bar interval $[\widehat{G}_N^{\mathcal{N}}(\boldsymbol{\mu}) - \Delta_{\widehat{G}}^{\mathcal{N}}(\boldsymbol{\mu}), \widehat{G}_N^{\mathcal{N}}(\boldsymbol{\mu}) + \Delta_{\widehat{G}}^{\mathcal{N}}(\boldsymbol{\mu})]$ for $\mu_1 \in [0.1, 0.45]$, $\mu_2 = 0.25$, and $N = 15$. It is clear that our error bound $\Delta_{\widehat{G}}^{\mathcal{N}}(\boldsymbol{\mu})$ is too large, mostly due to the relatively small value of $\delta\mu_1$ compared to the errors in the associated RB compliant output. We next plot, in Fig. 28, $\widehat{G}_N^{\mathcal{N}}(\boldsymbol{\mu})$ and the error bar interval $[\widehat{G}_N^{\mathcal{N}}(\boldsymbol{\mu}) - \Delta_{\widehat{G}}^{\mathcal{N}}(\boldsymbol{\mu}), \widehat{G}_N^{\mathcal{N}}(\boldsymbol{\mu}) + \Delta_{\widehat{G}}^{\mathcal{N}}(\boldsymbol{\mu})]$ for

Fig. 27 Elastic Crack problem: $\widehat{G}_N^{\mathcal{N}}(\boldsymbol{\mu})$ and the error bar interval $[\widehat{G}_N^{\mathcal{N}}(\boldsymbol{\mu}) - \Delta_{\widehat{G}}^{\mathcal{N}}(\boldsymbol{\mu}), \widehat{G}_N^{\mathcal{N}}(\boldsymbol{\mu}) + \Delta_{\widehat{G}}^{\mathcal{N}}(\boldsymbol{\mu})]$ for $\mu_1 \in [0.1, 0.45]$, $\mu_2 = 0.25$, and $N = 15$

Fig. 28 Elastic Crack problem: $\widehat{G}_N^{\mathcal{N}}(\boldsymbol{\mu})$ and the error bar interval $[\widehat{G}_N^{\mathcal{N}}(\boldsymbol{\mu}) - \Delta_{\widehat{G}}^{\mathcal{N}}(\boldsymbol{\mu}), \widehat{G}_N^{\mathcal{N}}(\boldsymbol{\mu}) + \Delta_{\widehat{G}}^{\mathcal{N}}(\boldsymbol{\mu})]$ for $\mu_1 \in [0.1, 0.45]$, $\mu_2 = 0.25$, and $N = 30$

$\mu_1 \in [0.1, 0.45]$, $\mu_2 = 0.25$, and $N = 30$; it is now observed that the estimated error is significantly improved—thanks to a better reduced basis approximation that compensates for the small value of $\delta\mu_1$. This example demonstrates how, in practice, the error bounds serve in the Online stage to confirm and, if necessary improve, the RB accuracy. We emphasize that $\widehat{G}^{\mathcal{N}}(\boldsymbol{\mu}) \in [\widehat{G}_N^{\mathcal{N}}(\boldsymbol{\mu}) - \Delta_{\widehat{G}}^{\mathcal{N}}(\boldsymbol{\mu}), \widehat{G}_N^{\mathcal{N}}(\boldsymbol{\mu}) + \Delta_{\widehat{G}}^{\mathcal{N}}(\boldsymbol{\mu})]$ for all $\boldsymbol{\mu} \in \mathcal{D}$ and all $N \in \{1, \dots, N_{\text{max}}\}$: the high-fidelity FE ERR prediction *must* lie within the error bounds provided.

⁹We note that a completely rigorous formulation of the exact and subsequently FE and RB crack problems—a construction that eliminates the nuisance, error, and uncertainty associated with $\delta\mu_1$ —is described in detail in [61]. This formulation is, however, *non-coercive* (and also rather complicated) and hence beyond of the scope of the current paper.

¹⁰We compare our results with reference data [91] available for several $\boldsymbol{\mu}$: our ERR results are in very good agreement with the reference results; for this particular $\delta\mu_1$ value, the maximum relative error compared to the reference results is only 2%. This confirms that our choice of $\delta\mu_1$ is adequately small.

8.2.4 Elastic Block

We now consider the three-dimensional elasticity problem described in Sect. 6.2.1; we recall that this problem is characterized by $P = 4$ parameters, three of which are geometric in nature. We choose for $\mathcal{E}_{\text{train}}$ a log-uniform random sample of size $n_{\text{train}} = 4000$. We focus in this (three-dimensional) example on computational performance.

Fig. 29 Elastic Block problem: (upper bound for the) $L^\infty(\mathcal{E}_{\text{train}})$ relative output error, (86), as a function of N

We first present in Fig. 29 the error measure

$$\max_{\mu \in \mathcal{E}_{\text{train}}} \frac{\Delta_N^s(\mu)}{s_N^{\mathcal{N}}(\mu)}; \quad (86)$$

note that $\Delta_N^s(\mu)$ is an upper bound for $|s^{\mathcal{N}}(\mu) - s_N^{\mathcal{N}}(\mu)|$ and $s_N^{\mathcal{N}}(\mu)$ is a lower bound for $s^{\mathcal{N}}(\mu)$, and therefore $\Delta_N^s(\mu)/s_N^{\mathcal{N}}(\mu)$ is in fact an upper bound for the relative error in the output. We observe that we again achieve rapid convergence: to obtain a relative output error of $1\text{E}-4$, we require only $N \approx 30$ points.

The spatial dimensionality plays little role in RB convergence: it follows that the relative efficiency of the RB approach—relative to direct FE evaluation—increases with increasing spatial dimension. We consider three different FE approximations: $X^{\mathcal{N}_{\text{vf}}}$ (very fine) corresponding to $\mathcal{N}_{\text{vf}} = 106,754$; $X^{\mathcal{N}_{\text{f}}}$ (fine) corresponding to $\mathcal{N}_{\text{f}} = 26,952$; and $X^{\mathcal{N}_{\text{c}}}$ (coarse) corresponding to $\mathcal{N}_{\text{c}} = 6,315$. We shall take the “very fine” approximation as “exact”: we then conclude (admittedly from a very coarse parameter test sample—since evaluation of $u^{\mathcal{N}_{\text{vf}}}(\mu)$ is very expensive) that the relative output error on the “fine” mesh is roughly 0.01 and on the “coarse” mesh roughly 0.09.

We now invoke the (energy version of the) greedy algorithm of Sect. 7.2.2 to construct two RB approximations spaces: $W_N^{\mathcal{N}_{\text{f}} \text{ Greedy, en}}, 1 \leq N \leq N_{\text{max f}}$, such that

$$\max_{\mu \in \mathcal{E}_{\text{train}}} \frac{\Delta_{N_{\text{max f}}}^s(\mu)}{s_{N_{\text{max f}}}^{\mathcal{N}_{\text{f}} \text{ Greedy, en}}(\mu)} = \frac{1}{4} \times 0.09$$

and $W_N^{\mathcal{N}_{\text{c}} \text{ Greedy, en}}, 1 \leq N \leq N_{\text{max c}}$, such that

$$\max_{\mu \in \mathcal{E}_{\text{train}}} \frac{\Delta_{N_{\text{max c}}}^s(\mu)}{s_{N_{\text{max c}}}^{\mathcal{N}_{\text{c}} \text{ Greedy, en}}(\mu)} = \frac{1}{4} \times 0.1,$$

where in both cases $\mathcal{E}_{\text{train}}$ is our random uniform train sample of size $n_{\text{train}} = 4000$; we obtain $N_{\text{max f}} = 22$ and

$N_{\text{max c}} = 15$. In essence, the RB output approximation associated with $W_{N_{\text{max f}}}^{\mathcal{N}_{\text{f}} \text{ Greedy, en}}$ replicates the “fine” FE approximation (relative to the exact result), and the RB output approximation associated with $W_{N_{\text{max c}}}^{\mathcal{N}_{\text{c}} \text{ Greedy, en}}$ replicates the “coarse” FE approximation (relative to the exact result). We thus can evaluate the efficiency of RB prediction relative to FE prediction *at fixed accuracy*.

We first consider the coarse case. We find that Online RB evaluation $\mu \rightarrow s_{N_{\text{max c}}}^{\mathcal{N}_{\text{c}}}(\mu), \Delta_{N_{\text{max c}}}^s(\mu)$ is roughly 50 times faster than direct FE evaluation $\mu \rightarrow s^{\mathcal{N}_{\text{c}}}(\mu)$: a considerable speed-up. Proceeding now to the fine case, the Online RB evaluation $\mu \rightarrow s_{N_{\text{max f}}}^{\mathcal{N}_{\text{f}}}(\mu), \Delta_{N_{\text{max f}}}^s(\mu)$ is now roughly 500 times faster than direct FE evaluation $\mu \rightarrow s^{\mathcal{N}_{\text{f}}}(\mu)$: as our (linear element) FE approximation converges only algebraically, and the RB approximation converges exponentially, the RB advantage increases as the error tolerance decreases. (In the fine case, the RB output error bound calculation $\Delta_{N_{\text{max f}}}^s(\mu)$ consumes roughly 70% of the Online time; we describe the Online error bound procedure in detail in Sect. 9.4.) Finally, we observe that for the fine case, the Offline effort is roughly 54 times more expensive than a single FE evaluation $\mu \rightarrow s^{\mathcal{N}_{\text{f}}}(\mu)$; equivalently, the many-query “break-even” point at which the RB approximation is first computationally “interesting” is a rather modest 54 evaluations.¹¹

Finally, we note that our comparison is skewed a bit towards the reduced basis—and hence the savings of 500 are probably too optimistic. First, we consider here only linear finite elements; at least quadratic finite elements should be considered. Second, we apply (minimum fill-in) Cholesky factorization to solve the FE systems [125]; certainly, at least for the finer FE meshes, preconditioned conjugate gradients or other efficient iterative solvers must be pursued. (However, the Cholesky factorization does have some benefits as regards the FE pseudo-solutions related to *a posteriori* error estimation: one factorization serves many right-hand sides, as we shall see in the next section.)

9 A Posteriori Error Estimation

9.1 Role

Effective *a posteriori* error bounds for the quantity of interest—our output—are crucial both for the efficiency and the reliability of RB approximations. As regards *efficiency*

¹¹Note the Offline operation count depends on $\mathcal{N}$ (the FE solutions and FE “pseudo”-solutions related to *a posteriori* error estimation) but also n_{train} (the greedy error bound maximization); ideally, the computational time associated with these two components should be in balance.

(related to the concept of “adaptivity” within the FE context), error bounds play a role in both the Offline and Online stages. In the greedy algorithms of Sect. 7.2.2, the application of error bounds (as surrogates for the actual error) permits significantly larger training samples $\mathcal{E}_{\text{train}} \subset \mathcal{D}$ at greatly reduced Offline computational cost. These more extensive training samples in turn engender RB approximations which provide high accuracy at greatly reduced Online computational cost. The error bounds also serve directly in the Online stage—to find the smallest RB dimension N that achieves the requisite accuracy—to further optimize Online performance. In short, *a posteriori* error estimation permits us to (inexpensively) control the error which in turn permits us to minimize the computational effort.

As regards *reliability*, it is clear that our Offline sampling procedures can not be exhaustive: for larger parameter dimensions P there will be large “parts” of the parameter set $\mathcal{D}$ that remain unexplored—the output error uncharacterized; we must admit that we will only encounter most parameter values in $\mathcal{D}$ Online. Our *a posteriori* estimation procedures ensure that we can rigorously and efficiently bound the output error in the Online (deployed/application) stage. We can thus be sure that constraints are satisfied, feasibility (and safety/failure) conditions are verified, and prognoses are valid: *real-time or design decisions are endowed with the full assurances of the high-fidelity FE solution*. In short, *a posteriori* error bounds permit us to confidently—with certainty—exploit the rapid predictive power of the RB approximation.

We should emphasize that *a posteriori* output error bounds are particularly important for RB approximations. First, RB approximations are *ad hoc*: each problem is different as regards discretization. Second, RB approximations are typically pre-asymptotic: we will choose N quite small—before any “tail” in the convergence rate. Third, the RB basis functions can not be directly related to any spatial or temporal scales: physical intuition is of little value. And fourth and finally, the RB approach is typically applied in the real-time context: there is no time for Offline verification; errors are immediately manifested and often in deleterious ways. There is, thus, even greater need for *a posteriori* error estimation in the RB context than in the much more studied FE context [2, 3, 11–13, 23].

Our motivations for error estimation in turn place requirements on our error bounds. First, the error bounds must be *rigorous*—valid for all N and for all parameter values in the parameter domain $\mathcal{D}$: non-rigorous error “indicators” may suffice for adaptivity, but not for reliability. Second, the bounds must be reasonably *sharp*: an overly conservative error bound can yield inefficient approximations (N too large) or suboptimal engineering results (unnecessary safety margins); design should be dictated by the output and not the output error. And third, the bounds must be very *efficient*: the

Online operation count and storage to compute the RB error bounds—the marginal or asymptotic average cost—must be independent of $\mathcal{N}$ (and hopefully commensurate with the cost associated with the RB output prediction). We do re-emphasize here that our RB error bounds are defined relative to the underlying “truth” FE approximation; however, we also recall that the RB Online cost is independent of $\mathcal{N}$, and hence the truth approximation can and should be chosen conservatively.

9.2 Preliminaries

The central equation in *a posteriori* theory is the error residual relationship. In particular, it follows from the problem statements for $u^{\mathcal{N}}(\boldsymbol{\mu})$, (3), and $u_N^{\mathcal{N}}(\boldsymbol{\mu})$, (64), that the error ($e^{\mathcal{N}}(\boldsymbol{\mu}) \equiv e(\boldsymbol{\mu}) \equiv u^{\mathcal{N}}(\boldsymbol{\mu}) - u_N^{\mathcal{N}}(\boldsymbol{\mu}) \in X^{\mathcal{N}}$ satisfies

$$a(e(\boldsymbol{\mu}), v; \boldsymbol{\mu}) = r(v; \boldsymbol{\mu}), \quad \forall v \in X^{\mathcal{N}}. \quad (87)$$

Here $r(v; \boldsymbol{\mu}) \in (X^{\mathcal{N}})'$ (the dual space to $X^{\mathcal{N}}$) is the residual,

$$r(v; \boldsymbol{\mu}) \equiv f(v; \boldsymbol{\mu}) - a(u_N^{\mathcal{N}}(\boldsymbol{\mu}), v; \boldsymbol{\mu}), \quad \forall v \in X^{\mathcal{N}}. \quad (88)$$

(Indeed, (87) directly follows from the definition (88), $f(v; \boldsymbol{\mu}) = a(u^{\mathcal{N}}(\boldsymbol{\mu}), v; \boldsymbol{\mu})$, $\forall v \in X^{\mathcal{N}}$, bilinearity of a , and the definition of $e(\boldsymbol{\mu})$.)

It shall prove convenient to introduce the Riesz representation of $r(v; \boldsymbol{\mu})$: $\hat{e}(\boldsymbol{\mu}) \in X^{\mathcal{N}}$ [112] satisfies

$$(\hat{e}(\boldsymbol{\mu}), v)_X = r(v; \boldsymbol{\mu}), \quad \forall v \in X^{\mathcal{N}}. \quad (89)$$

We can thus also write the error residual equation (87) as

$$a(e(\boldsymbol{\mu}), v; \boldsymbol{\mu}) = (\hat{e}(\boldsymbol{\mu}), v)_X, \quad \forall v \in X^{\mathcal{N}}. \quad (90)$$

It also follows that

$$\|r(\cdot; \boldsymbol{\mu})\|_{(X^{\mathcal{N}})'} \equiv \sup_{v \in X^{\mathcal{N}}} \frac{r(v; \boldsymbol{\mu})}{\|v\|_X} = \|\hat{e}(\boldsymbol{\mu})\|_X; \quad (91)$$

the evaluation of the dual norm of the residual through the Riesz representation is central to the Offline-Online procedures developed in Sect. 9.4 below.

We recall the definition of the exact and FE coercivity constants, (7) and (8), respectively. We shall require a lower bound to the coercivity constant $\alpha^{\mathcal{N}}(\boldsymbol{\mu})$, $\alpha_{\text{LB}}^{\mathcal{N}}: \mathcal{D} \rightarrow \mathbb{R}$, such that (i) $0 < \alpha_{\text{LB}}^{\mathcal{N}}(\boldsymbol{\mu}) \leq \alpha^{\mathcal{N}}(\boldsymbol{\mu})$, $\forall \boldsymbol{\mu} \in \mathcal{D}$, and (ii) the Online computational time to evaluate $\boldsymbol{\mu} \rightarrow \alpha_{\text{LB}}^{\mathcal{N}}(\boldsymbol{\mu})$ is independent of $\mathcal{N}$. In Sect. 10 we provide a methodology [64] to construct the requisite lower bound.

9.3 Error Bounds

We define error estimators for the energy norm and output as

$$\Delta_N^{\text{en}}(\boldsymbol{\mu}) \equiv \|\hat{e}(\boldsymbol{\mu})\|_X / (\alpha_{\text{LB}}^{\mathcal{N}}(\boldsymbol{\mu}))^{1/2},$$

and

$$\Delta_N^s(\mu) \equiv \|\hat{e}(\mu)\|_X^2 / \alpha_{\text{LB}}^{\mathcal{N}}(\mu),$$

respectively. We next introduce the effectivities associated with these error estimators as

$$\eta_N^{\text{en}}(\mu) \equiv \Delta_N^{\text{en}}(\mu) / \|u^{\mathcal{N}}(\mu) - u_N^{\mathcal{N}}(\mu)\|_{\mu},$$

and

$$\eta_N^s(\mu) \equiv \Delta_N^s(\mu) / (s^{\mathcal{N}}(\mu) - s_N^{\mathcal{N}}(\mu)),$$

respectively.

Clearly, the effectivities are a measure of the quality of the proposed estimator: for rigor, we shall insist upon effectivities ≥ 1 ; for sharpness, we desire effectivities as close to unity as possible. We can prove

Proposition 2 *For any $N = 1, \dots, N_{\max}$, the effectivities satisfy*

$$1 \leq \eta_N^{\text{en}}(\mu) \leq \sqrt{\frac{\gamma^e(\mu)}{\alpha_{\text{LB}}^{\mathcal{N}}(\mu)}}, \quad \forall \mu \in \mathcal{D}, \quad (92)$$

$$1 \leq \eta_N^s(\mu) \leq \frac{\gamma^e(\mu)}{\alpha_{\text{LB}}^{\mathcal{N}}(\mu)}, \quad \forall \mu \in \mathcal{D}. \quad (93)$$

Proof It follows directly from (90) for $v = e(\mu)$ and the Cauchy-Schwarz inequality that

$$\|e(\mu)\|_{\mu}^2 \leq \|\hat{e}(\mu)\|_X \|e(\mu)\|_X. \quad (94)$$

But $(\alpha_{\text{LB}}^{\mathcal{N}}(\mu))^{\frac{1}{2}} \|e(\mu)\|_X \leq a^{\frac{1}{2}}(e(\mu), e(\mu); \mu) \equiv \|e(\mu)\|_{\mu}$, and hence from (94) we obtain $\|e(\mu)\|_{\mu} \leq \Delta_N^{\text{en}}(\mu)$ or $\eta_N^{\text{en}}(\mu) \geq 1$. We now again consider (90)—but now for $v = \hat{e}(\mu)$ —and the Cauchy-Schwarz inequality to obtain

$$\|\hat{e}(\mu)\|_X^2 \leq \|\hat{e}(\mu)\|_{\mu} \|e(\mu)\|_{\mu}. \quad (95)$$

But from continuity $\|\hat{e}(\mu)\|_{\mu} \leq (\gamma^e(\mu))^{\frac{1}{2}} \|\hat{e}(\mu)\|_X$, and hence from (95) $\Delta_N^{\text{en}}(\mu) \equiv (\alpha_{\text{LB}}^{\mathcal{N}}(\mu))^{-\frac{1}{2}} \|\hat{e}(\mu)\|_X \leq (\alpha_{\text{LB}}^{\mathcal{N}}(\mu))^{-\frac{1}{2}} (\gamma^e(\mu))^{\frac{1}{2}} \|e(\mu)\|_{\mu}$, or $\eta_N^{\text{en}}(\mu) \leq \sqrt{\frac{\gamma^e(\mu)}{\alpha_{\text{LB}}^{\mathcal{N}}(\mu)}}$.

Next, we know from (66) that $s^{\mathcal{N}}(\mu) - s_N^{\mathcal{N}}(\mu) = \|e(\mu)\|_{\mu}^2$, and hence since $\Delta_N^s(\mu) = (\Delta_N^{\text{en}}(\mu))^2$

$$\eta_N^s(\mu) \equiv \frac{\Delta_N^s(\mu)}{s^{\mathcal{N}}(\mu) - s_N^{\mathcal{N}}(\mu)} = \frac{(\Delta_N^{\text{en}}(\mu))^2}{\|e(\mu)\|_{\mu}^2} = (\eta_N^{\text{en}}(\mu))^2; \quad (96)$$

(93) directly follows from (92) and (96). $\square$

Similar results can be obtained for $\Delta_N(\mu)$, the *a posteriori* error bound in the X norm.

It is important to observe that our effectivity upper bounds, (92) and (93), are *independent* of N , and hence stable with respect to *RB refinement*. Furthermore, it is sometimes possible (see Sect. 10) to provide a rigorous lower bound for $\alpha_{\text{LB}}^{\mathcal{N}}(\mu)$ that depends only on μ : in this case we obtain an upper bound for the effectivity which is (not only independent of N but also) independent of $\mathcal{N}$, and hence stable with respect to *FE refinement*; the latter reflects our proper choice of $(H^1(\Omega)$ -equivalent) inner product/norm. More generally, our construction for $\alpha_{\text{LB}}^{\mathcal{N}}(\mu)$ (see Sect. 10) is designed to ensure that $\alpha^{\mathcal{N}}(\mu)/\alpha_{\text{LB}}^{\mathcal{N}}(\mu)$ will be bounded by a constant—typically 4 for our choice of tolerances—for most μ in $\mathcal{D}$; since (say, for the output effectivity)

$$\eta_N^s(\mu) \leq \frac{\gamma^e(\mu)}{\alpha_{\text{LB}}^{\mathcal{N}}(\mu)} \leq \frac{\alpha^{\mathcal{N}}(\mu)}{\alpha_{\text{LB}}^{\mathcal{N}}(\mu)} \frac{\gamma^e(\mu)}{\alpha^e(\mu)}, \quad \forall \mu \in \mathcal{D}, \quad (97)$$

we again obtain stability with respect to FE refinement.

9.4 Offline-Online: $\|\hat{e}(\mu)\|_X$

The error bounds of the previous section are of no utility without an accompanying Offline-Online computational approach.

9.4.1 Ingredients

The computationally crucial component of all the error bounds of the previous section is $\|\hat{e}(\mu)\|_X$, the dual norm of the residual. (Offline-Online treatment of $\alpha_{\text{LB}}^{\mathcal{N}}(\mu)$ is addressed in Sect. 10.)

To develop an Offline-Online procedure for the dual norm of the residual we first expand the residual (88) according to (67) and (2):

$$\begin{aligned} r(v; \mu) &\equiv f(v) - a(u_N^{\mathcal{N}}(\mu), v; \mu) \\ &= f(v) - a\left(\sum_{n=1}^N u_{Nn}^{\mathcal{N}}(\mu) \zeta_n^{\mathcal{N}}, v; \mu\right) \\ &= f(v) - \sum_{n=1}^N u_{Nn}^{\mathcal{N}}(\mu) a(\zeta_n^{\mathcal{N}}, v; \mu) \\ &= f(v) - \sum_{n=1}^N u_{Nn}^{\mathcal{N}}(\mu) \sum_{q=1}^Q \Theta^q(\mu) a^q(\zeta_n^{\mathcal{N}}, v). \end{aligned} \quad (98)$$

If we insert (98) in (89) and apply linear superposition, we obtain

$$(\hat{e}(\mu), v)_X = f(v) - \sum_{q=1}^Q \sum_{n=1}^N \Theta^q(\mu) u_{Nn}^{\mathcal{N}}(\mu) a^q(\zeta_n^{\mathcal{N}}, v),$$

or

$$\hat{e}(\mu) = \mathcal{C} + \sum_{q=1}^Q \sum_{n=1}^N \Theta^q(\mu) u_{Nn}^{\mathcal{N}}(\mu) \mathcal{L}_n^q,$$

where $(\mathcal{C}, v)_X = f(v)$, $\forall v \in X^{\mathcal{N}}$, and $(\mathcal{L}_n^q, v)_X = -a^q(\zeta_n^{\mathcal{N}}, v)$, $\forall v \in X^{\mathcal{N}}$, $1 \leq n \leq N$, $1 \leq q \leq Q$. We denote the $\mathcal{C}$, $\mathcal{L}_n^q$, $1 \leq n \leq N$, $1 \leq q \leq Q$, as FE “pseudo”-solutions—solutions of “associated” FE Poisson problems.

We thus obtain

$$\begin{aligned} \|\hat{e}(\mu)\|_X^2 &= \left(\mathcal{C} + \sum_{q=1}^Q \sum_{n=1}^N \Theta^q(\mu) u_{Nn}^{\mathcal{N}}(\mu) \mathcal{L}_n^q, \bullet \right)_X \\ &= (\mathcal{C}, \mathcal{C})_X + \sum_{q=1}^Q \sum_{n=1}^N \Theta^q(\mu) u_{Nn}^{\mathcal{N}}(\mu) \\ &\quad \times \left\{ 2(\mathcal{C}, \mathcal{L}_n^q)_X + \sum_{q'=1}^Q \sum_{n'=1}^N \Theta^{q'}(\mu) u_{Nn'}^{\mathcal{N}}(\mu) (\mathcal{L}_n^q, \mathcal{L}_{n'}^{q'})_X \right\}, \end{aligned} \quad (99)$$

from which we can directly calculate the requisite dual norm of the residual through (91).

9.4.2 Computational Procedure

The Offline-Online decomposition is now clear. In the Offline stage we form the parameter-independent quantities. In particular, we compute the FE “pseudo”-solutions $\mathcal{C}$, $\mathcal{L}_n^q$, $1 \leq n \leq N_{\max}$, $1 \leq q \leq Q$, and form/store $(\mathcal{C}, \mathcal{C})_X$, $(\mathcal{C}, \mathcal{L}_n^q)_X$, $(\mathcal{L}_n^q, \mathcal{L}_{n'}^{q'})_X$, $1 \leq n, n' \leq N_{\max}$, $1 \leq q, q' \leq Q$. (Note that, in the direct context, a single factorization suffices to obtain all $1 + QN_{\max}$ FE pseudo-solutions.) The Offline operation count depends on $N_{\max}$, Q , and $\mathcal{N}$.

In the Online stage, given any “new” value of μ —and $\Theta^q(\mu)$, $1 \leq q \leq Q$, $u_{Nn}^{\mathcal{N}}(\mu)$, $1 \leq n \leq N$ —we simply retrieve the stored quantities $(\mathcal{C}, \mathcal{C})_X$, $(\mathcal{C}, \mathcal{L}_n^q)_X$, $(\mathcal{L}_n^q, \mathcal{L}_{n'}^{q'})_X$, $1 \leq n, n' \leq N$, $1 \leq q, q' \leq Q$, and then evaluate the sum (99). The Online operation count, and hence also the marginal cost and asymptotic average cost, is $O(Q^2 N^2)$ —and independent of $\mathcal{N}$.¹² (Note again the advantage of the hierarchical spaces: the necessary quantities for any $N \in \{1, \dots, N_{\max}\}$ can be simply extracted from the corresponding quantities for $N = N_{\max}$.)

¹²It thus follows that the *a posteriori* error estimation contribution to the cost of the greedy algorithm of Sect. 7.2.2 is $O(QN_{\max}\mathcal{N}) + O(Q^2 N_{\max}^2 \mathcal{N}) + O(n_{\text{train}} Q^2 N_{\max}^3)$: we may thus choose $\mathcal{N}$ and n_{train} independently (and large).

9.5 Numerical Results

In fact, in Sect. 8 we have already presented several numerical examples illustrating the relevance, (Offline and Online) application, and computational efficiency of the *a posteriori* error bounds. In the current section (and again in Sect. 11.3 for the non-compliant case) we restrict attention to a brief study of *sharpness* as measured by the effectivity.

We consider here the Thermal Block problem of Sect. 6.1.1 with $B_1 = 3$, $B_2 = 3$, and hence $P = 8$. Here $1/\mu_i^{\min} = \mu_i^{\max} = \sqrt{\mu_r}$ ($= 10$), $1 \leq i \leq P$, and hence $\mu_i^{\max}/\mu_i^{\min} = \mu_r$ ($= 100$), $1 \leq i \leq P$; we choose $\bar{\mu}_i = 1$, $1 \leq i \leq P$, and $\tau = 0$ in our inner product, (6). For the truth discretization, we take $\mathcal{N} = 661$: we confirm, per the theory, that the effectivities are insensitive to $\mathcal{N}$ for sufficiently large $\mathcal{N}$. (For the construction of $\alpha_{\text{LB}}^{\mathcal{N}}(\mu)$ for this particular problem, see Sect. 10.2.)

We then define, for a given (fine) test sample $\mathcal{E}$,

$$\begin{aligned} \Delta_{N,\max}^s &= \max_{\mu \in \mathcal{E}} \Delta_N^s(\mu), \\ \eta_{N,\max}^s &= \max_{\mu \in \mathcal{E}} \eta_N^s(\mu), \end{aligned}$$

and

$$\eta_{N,\text{ave}}^s = \frac{1}{|\mathcal{E}|} \sum_{\mu \in \mathcal{E}} \eta_N^s(\mu).$$

We shall shortly provide numerical results for these effectivity metrics. However, we first briefly derive a theoretical result for the particular problem of interest.

To wit, we note from (59) that

$$\begin{aligned} |a(w, v; \mu)| &= \left| \sum_{i=1}^P \mu_i \int_{\mathcal{R}_i} \nabla w \cdot \nabla v + \int_{\mathcal{R}_{P+1}} \nabla w \cdot \nabla v \right| \\ &\leq \max(1, \mu_1, \dots, \mu_P) \sum_{i=1}^{P+1} \int_{\mathcal{R}_i} |\nabla w| |\nabla v|. \end{aligned}$$

However, for our choice of norm,

$$\begin{aligned} \sum_{i=1}^{P+1} \int_{\mathcal{R}_i} |\nabla w| |\nabla v| &\leq \sum_{i=1}^{P+1} \left(\int_{\mathcal{R}_i} |\nabla w|^2 \right)^{\frac{1}{2}} \left(\int_{\mathcal{R}_i} |\nabla v|^2 \right)^{\frac{1}{2}} \\ &\leq \|w\|_X \|v\|_X \end{aligned}$$

by repeated repeated application of the Cauchy-Schwarz inequality. We conclude that $\gamma^e(\mu) \leq \max(1, \mu_1, \dots, \mu_P)$. We shall further demonstrate in Sect. 10.2 that $\alpha_{\text{LB}}^{\mathcal{N}}(\mu) \geq \min(1, \mu_1, \dots, \mu_P)$. Hence,

$$\eta_{N,\max}^s \leq \max_{\mu \in \mathcal{D}} \frac{\gamma^e(\mu)}{\alpha_{\text{LB}}^{\mathcal{N}}(\mu)}$$

Table 1 Thermal Block problem for $B_1 = B_2 = 3$: Output error bound and effectivity metrics as a function of N

N	$\Delta_{N,\max}^s$	$\eta_{N,\max}^s$	$\eta_{N,\text{ave}}^s$
10	2.2036E+00	31.2850	6.7067
20	2.0020E-01	37.3024	7.5587
30	1.5100E-02	62.2537	12.1138
40	1.2000E-03	73.1151	14.4598
50	1.0000E-04	57.5113	10.2566

$$\leq \max_{\mu \in \mathcal{D}} \frac{\max(1, \mu_1, \dots, \mu_P)}{\min(1, \mu_1, \dots, \mu_P)} \leq \mu_r, \quad (100)$$

since we can readily find a $\mu \in \mathcal{D}$ such that $\max(1, \mu_1, \dots, \mu_P) = \sqrt{\mu_r}$ and $\min(1, \mu_1, \dots, \mu_P) = 1/\sqrt{\mu_r}$.

We now turn to the numerical results. We present in Table 1 $\Delta_{N,\max}^s$, $\eta_{N,\max}^s$, and $\eta_{N,\text{ave}}^s$ as a function of N . The effectivities are of course greater than unity: our error estimates are rigorous upper bounds. We further observe that $\eta_{N,\max}^s \leq \mu_r = 100$: unfortunately, the theoretical estimate (100) is reasonably accurate, and hence $\eta_{N,\max}^s$ can be quite large; however, on “average,” the effectivities are quite moderate even for this rather extensive parameter variation. Finally, we note that the effect of error over-estimation on (decreased) efficiency—due to choice of an unnecessarily large N —will be modest given the rapid convergence of the RB approximation [112].

10 Lower Bounds for the Coercivity Constant

As introduced in Sect. 9, our *a posteriori* error analysis of reduced basis approximations to (affinely) parametrized partial differential equations requires a lower bound for the coercivity constant.

In essence, the discrete coercivity constant (8) is a generalized minimum eigenvalue [112]. There are many classical techniques for the estimation of minimum eigenvalues or minimum singular values. One class of methods is based on Gershgorin’s theorem and variants [110]. Within our particular context, these approaches are not optimal: generalized eigenvalue and singular value problems are difficult to treat; the operation count will scale with $\mathcal{N}$; and finally, the Gershgorin-like bounds are often not useful for elliptic PDEs. A second class of methods is based on eigenfunction/eigenvalue (e.g., Rayleigh Ritz) approximation and subsequent residual evaluation [65, 71]. Unfortunately, the lower bounds are not truly rigorous: we obtain lower bounds not for the smallest eigenvalue, but rather for the eigenvalue closest to the proposed approximate eigenvalue.

In this section we shall describe the Successive Constraint Method (SCM), an approach to the construction of

lower bounds for coercivity (and, in the non-coercive case, inf-sup stability) constants [64]. The method—based on an Offline-Online strategy relevant in the many-query and real-time context—reduces the Online (real-time/deployed) calculation to a small Linear Program for which the operation count is *independent* of $\mathcal{N}$. The SCM is more efficient and general than earlier proposals [97, 121, 142]; also the SCM is much more easily implemented [64].

10.1 Successive Constraint Method (SCM)

10.1.1 Preliminaries

We start by recalling the definition of our FE coercivity constant of Sect. 4,

$$\alpha^{\mathcal{N}}(\mu) = \inf_{w \in X^{\mathcal{N}}} \frac{a(w, w; \mu)}{\|w\|_X^2}, \quad \forall \mu \in \mathcal{D}. \quad (101)$$

We assume that a is coercive, $\alpha^{\mathcal{N}}(\mu) > 0$, $\forall \mu \in \mathcal{D}$, and continuous; we also recall that a is “affine” in the parameter μ , (2), and symmetric. All of these properties shall be exploited in our development here. However, we note that the SCM can be readily extended to non-symmetric coercive operators (see Sect. 11) as well as general non-coercive operators [61, 64, 141].

We first introduce an objective function $\mathcal{J}^{\text{obj}} : \mathcal{D} \times \mathbb{R}^Q \rightarrow \mathbb{R}$ given by

$$\mathcal{J}^{\text{obj}}(\mu; y) = \sum_{q=1}^Q \Theta^q(\mu) y_q, \quad (102)$$

where $y = (y_1, \dots, y_Q)$. We may then express our coercivity constant as

$$\alpha^{\mathcal{N}}(\mu) = \inf_{y \in \mathcal{Y}} \mathcal{J}^{\text{obj}}(\mu; y), \quad (103)$$

where the set $\mathcal{Y} \subset \mathbb{R}^Q$ is defined by

$$\mathcal{Y} = \left\{ y \in \mathbb{R}^Q \mid \exists w_y \in X^{\mathcal{N}} \text{ s.t. } y_q = \frac{a^q(w_y, w_y)}{\|w_y\|_X^2}, \ 1 \leq q \leq Q \right\}. \quad (104)$$

The equivalence between (101) and (103), (104) is readily confirmed: (2) is the crucial ingredient.

We next introduce the “continuity constraint” box

$$\mathcal{B} = \prod_{q=1}^Q \left[\inf_{w \in X^{\mathcal{N}}} \frac{a^q(w, w)}{\|w\|_X^2}, \sup_{w \in X^{\mathcal{N}}} \frac{a^q(w, w)}{\|w\|_X^2} \right]; \quad (105)$$

from our continuity hypothesis, $\mathcal{B}$ is bounded. We then define our “coercivity constraint” sample,

$$\mathcal{C}_J = \{\mu_{\text{SCM}}^1 \in \mathcal{D}, \dots, \mu_{\text{SCM}}^J \in \mathcal{D}\}. \quad (106)$$

We denote by $\mathcal{C}_J^{M,\mu}$ the set of M (≥ 1) points in $\mathcal{C}_J$ closest (in the usual Euclidian norm) to a given $\mu \in \mathcal{D}$. (Note that if $M > J$, then we set $\mathcal{C}_J^{M,\mu} = \mathcal{C}_J$.)

10.1.2 Lower Bound

Now for given $\mathcal{C}_J$, $M \in \mathbb{N} \equiv \{1, 2, \dots\}$, and any $\mu \in \mathcal{D}$, we define the “lower bound” set $\mathcal{Y}_{\text{LB}}(\mu; \mathcal{C}_J, M) \subset \mathbb{R}^Q$ as

$$\mathcal{Y}_{\text{LB}}(\mu; \mathcal{C}_J, M) \equiv \left\{ y \in \mathbb{R}^Q \mid y \in \mathcal{B}; \sum_{q=1}^Q \Theta^q(\mu') y_q \geq \alpha^{\mathcal{N}}(\mu'), \forall \mu' \in \mathcal{C}_J^{M,\mu} \right\}. \quad (107)$$

We then demonstrate

Lemma 1 Given $\mathcal{C}_J \subset \mathcal{D}$ and $M \in \mathbb{N}$,

$$\mathcal{Y} \subset \mathcal{Y}_{\text{LB}}(\mu; \mathcal{C}_J, M), \quad \forall \mu \in \mathcal{D}. \quad (108)$$

Proof For any $y \in \mathcal{Y}$, $\exists w_y \in X^{\mathcal{N}}$ such that $y_q = \frac{a^q(w_y, w_y)}{\|w_y\|_X^2}$, $1 \leq q \leq Q$. Then, since

$$\inf_{w \in X^{\mathcal{N}}} \frac{a^q(w, w)}{\|w\|_X^2} \leq \underbrace{\frac{a^q(w_y, w_y)}{\|w_y\|_X^2}}_{y_q} \leq \sup_{w \in X^{\mathcal{N}}} \frac{a^q(w, w)}{\|w\|_X^2}, \quad (109)$$

and also

$$\begin{aligned} \sum_{q=1}^Q \Theta^q(\mu) \frac{a^q(w_y, w_y)}{\|w_y\|_X^2} &= \frac{a(w_y, w_y; \mu)}{\|w_y\|_X^2} \\ &\geq \alpha^{\mathcal{N}}(\mu), \quad \forall \mu \in \mathcal{D}, \end{aligned} \quad (110)$$

it follows that every member y of $\mathcal{Y}$ is also a member of $\mathcal{Y}_{\text{LB}}(\mu; \mathcal{C}_J, M)$. This concludes the proof. $\square$

We then define our lower bound

$$\alpha_{\text{LB}}^{\mathcal{N}}(\mu; \mathcal{C}_J, M) = \min_{y \in \mathcal{Y}_{\text{LB}}(\mu; \mathcal{C}_J, M)} \mathcal{J}^{\text{obj}}(\mu; y). \quad (111)$$

We may then obtain

Proposition 3 For given $\mathcal{C}_J \subset \mathcal{D}$, $M \in \mathbb{N}$,

$$\alpha_{\text{LB}}^{\mathcal{N}}(\mu) \leq \alpha^{\mathcal{N}}(\mu), \quad \forall \mu \in \mathcal{D}. \quad (112)$$

Proof It readily follows that, given $\mathcal{C}_J \subset \mathcal{D}$, $M \in \mathbb{N}$,

$$\begin{aligned} \alpha_{\text{LB}}^{\mathcal{N}}(\mu) &= \min_{y \in \mathcal{Y}_{\text{LB}}(\mu; \mathcal{C}_J, M)} \mathcal{J}^{\text{obj}}(\mu; y) \\ &\leq \min_{y \in \mathcal{Y}} \mathcal{J}^{\text{obj}}(\mu; y) \\ &= \alpha^{\mathcal{N}}(\mu), \quad \forall \mu \in \mathcal{D}; \end{aligned} \quad (113)$$

here we invoke the definition of $\alpha_{\text{LB}}^{\mathcal{N}}(\mu)$, (111), Lemma 1, and then (103). $\square$

We note that our lower bound (111) is in fact a linear optimization problem (or Linear Program (LP)); indeed, (111) resembles a discretized linear semi-infinite program [51]. We observe that our LP (111) contains Q design variables and $2Q + M$ (one-sided) inequality constraints. The crucial observation is that the operation count to evaluate $\mu \rightarrow \alpha_{\text{LB}}^{\mathcal{N}}(\mu)$, given $\mathcal{B}$ and the set $\{\alpha^{\mathcal{N}}(\mu') \mid \mu' \in \mathcal{C}_J\}$, is independent of $\mathcal{N}$. We return to this point in Sect. 10.1.5.

We pause here and make an observation for the case in which the coercive bilinear form a is *parametrically* coercive. We say that a is *parametrically coercive* [112] if $\Theta^q(\mu) > 0$, $\forall \mu \in \mathcal{D}$, and $a^q(w, w) \geq 0$, $\forall w \in X^c$, $1 \leq q \leq Q$. In this case, for any $\mu \in \mathcal{D}$, any $\mu' \in \mathcal{C}_J^{M,\mu}$, and any $y \in \mathcal{Y}_{\text{LB}}(\mu; \mathcal{C}_J, M)$,

$$\begin{aligned} \mathcal{J}^{\text{obj}}(\mu; y) &= \sum_{q=1}^Q \Theta^q(\mu) y_q = \sum_{q=1}^Q \frac{\Theta^q(\mu)}{\Theta^q(\mu')} \Theta^q(\mu') y_q \\ &\geq \min_{q \in [1, \dots, Q]} \frac{\Theta^q(\mu)}{\Theta^q(\mu')} \sum_{q=1}^Q \Theta^q(\mu') y_q \\ &\geq \min_{q \in [1, \dots, Q]} \frac{\Theta^q(\mu)}{\Theta^q(\mu')} \alpha^{\mathcal{N}}(\mu') \end{aligned}$$

since $\mathcal{B} \subset \mathbb{R}_+^Q$ in the parametrically coercive case. (Here $\mathbb{R}_+^Q$ is the set $\{y \in \mathbb{R}^Q \mid y_q \geq 0, 1 \leq q \leq Q\}$.) It is observed that, for this particular case, the SCM will provide a positive lower bound even for $J = 1$; indeed for $J \geq 1$,

$$\alpha_{\text{LB}}^{\mathcal{N}}(\mu; \mathcal{C}_J, M) \geq \max_{\mu' \in \mathcal{C}_J^{M,\mu}} \left(\min_{q \in [1, \dots, Q]} \frac{\Theta^q(\mu)}{\Theta^q(\mu')} \alpha^{\mathcal{N}}(\mu') \right). \quad (114)$$

The result (114) will not in general be true for problems that are not parametrically coercive.

10.1.3 Upper Bound

As we shall see, we also require an *upper* bound for the coercivity constant for the (effective) construction of a good “coercivity constraint” sample $\mathcal{C}_J$. For given $\mathcal{C}_J$, $M \in \mathbb{N}$, and any $\mu \in \mathcal{D}$, we introduce our “upper bound” set $\mathcal{Y}_{\text{UB}}(\mu; \mathcal{C}_J, M) \in \mathbb{R}^Q$ as

$$\mathcal{Y}_{\text{UB}}(\mu; \mathcal{C}_J, M) = \{y^*(\mu') \mid \mu' \in \mathcal{C}_J^{M,\mu}\}, \quad (115)$$

where

$$y^*(\mu) = \arg \inf_{y \in \mathcal{Y}} \mathcal{J}^{\text{obj}}(\mu; y)$$

(in the event of non-uniqueness, any minimizer suffices). We can then define our upper bound as

$$\alpha_{\text{UB}}^{\mathcal{N}}(\mu; C_J, M) = \min_{y \in \mathcal{Y}_{\text{UB}}(\mu; C_J, M)} \mathcal{J}^{\text{obj}}(\mu; y). \quad (116)$$

It directly follows from (115) that $\mathcal{Y}_{\text{UB}}(\mu; C_J, M) \subset \mathcal{Y}$ and hence, for given $C_J, M \in \mathbb{N}$, $\alpha_{\text{UB}}^{\mathcal{N}}(\mu; C_J, M) \geq \alpha^{\mathcal{N}}(\mu)$, $\forall \mu \in \mathcal{D}$.

We note that the upper bound (116) is a simple enumeration; the operation count for the Online stage to evaluate $\mu \rightarrow \alpha_{\text{UB}}^{\mathcal{N}}(\mu)$, given the set $\{y^*(\mu') \mid \mu' \in C_J\}$, is *independent* of $\mathcal{N}$.

10.1.4 Greedy Selection of C_J

In some cases, in particular for parametrically coercive problems, we can often specify our lower bound construction “by hand”: we prescribe J, C_J , and M ; the lower bound (111) then directly follows. We shall give an example in Sect. 10.2. However, more generally, we must appeal to a greedy algorithm.

We now present the construction of the set C_J by an Offline greedy algorithm. We shall require a “train” sample $\mathcal{E}_{\text{train,SCM}} = \{\mu_{\text{train,SCM}}^1, \dots, \mu_{\text{train,SCM}}^{n_{\text{train,SCM}}}\} \subset \mathcal{D}$ of $n_{\text{train,SCM}}$ parameter points. We also require a tolerance $\varepsilon_{\text{SCM}} \in (0, 1)$ which shall control the error in the lower bound prediction. We first set $J = 1$ and choose $C_1 = \{\mu_{\text{SCM}}^1\}$ arbitrarily. We then perform

$$\begin{aligned} \text{While } \max_{\mu \in \mathcal{E}_{\text{train,SCM}}} \left[\frac{\alpha_{\text{UB}}^{\mathcal{N}}(\mu; C_J, M) - \alpha_{\text{LB}}^{\mathcal{N}}(\mu; C_J, M)}{\alpha_{\text{UB}}^{\mathcal{N}}(\mu; C_J, M)} \right] &> \varepsilon_{\text{SCM}}: \\ \mu_{\text{SCM}}^{J+1} &= \arg \max_{\mu \in \mathcal{E}_{\text{train,SCM}}} \left[\frac{\alpha_{\text{UB}}^{\mathcal{N}}(\mu; C_J, M) - \alpha_{\text{LB}}^{\mathcal{N}}(\mu; C_J, M)}{\alpha_{\text{UB}}^{\mathcal{N}}(\mu; C_J, M)} \right]; \\ C_{J+1} &= C_J \cup \mu_{\text{SCM}}^{J+1}; \\ J &\leftarrow J + 1; \end{aligned}$$

end.

Set $J_{\text{max}} = J$.

Note that $J_{\text{max}} = J_{\text{max}}(\varepsilon_{\text{SCM}})$ refers to the particular lower bound construction which satisfies the specified error tolerance.

In essence, at each iteration of the greedy procedure, we add to our “coercivity constraint” sample that point in $\mathcal{D}$ for which (roughly) the current lower bound approximation is least accurate; we may anticipate that the “gap” between $\alpha_{\text{LB}}^{\mathcal{N}}(\mu; C_J, M)$ and $\alpha_{\text{UB}}^{\mathcal{N}}(\mu; C_J, M)$ will be reduced at each iteration. In fact, since $\alpha_{\text{UB}}^{\mathcal{N}}(\mu; C_J, M) = \alpha_{\text{LB}}^{\mathcal{N}}(\mu; C_J, M)$, $\forall \mu \in C_J$, it follows from continuity considerations that, for sufficiently large J ($= J_{\text{max}}(\varepsilon_{\text{SCM}})$), our error tolerance ε_{SCM} will be honored. In practice, $J_{\text{max}}(\varepsilon_{\text{SCM}})$ will be reasonably small, as we discuss further in Sect. 10.2.

We note that we choose $\alpha_{\text{UB}}^{\mathcal{N}}(\mu; C_J, M)$, not $\alpha_{\text{LB}}^{\mathcal{N}}(\mu; C_J, M)$, in the denominator of our selection criterion since $\alpha_{\text{LB}}^{\mathcal{N}}(\mu; C_J, M)$ may be negative or zero. Furthermore, our choice of stopping criterion permits us to bound

$$\begin{aligned} &\frac{\alpha^{\mathcal{N}}(\mu)}{\alpha_{\text{LB}}^{\mathcal{N}}(\mu; C_{J_{\text{max}}}, M)} \\ &= \frac{\alpha^{\mathcal{N}}(\mu)}{\alpha_{\text{UB}}^{\mathcal{N}}(\mu; C_{J_{\text{max}}}, M) - (\alpha_{\text{UB}}^{\mathcal{N}}(\mu; C_{J_{\text{max}}}, M) - \alpha_{\text{LB}}^{\mathcal{N}}(\mu; C_{J_{\text{max}}}, M))} \\ &\leq \frac{\alpha^{\mathcal{N}}(\mu)}{\alpha_{\text{UB}}^{\mathcal{N}}(\mu; C_{J_{\text{max}}}, M)} \frac{1}{1 - \varepsilon_{\text{SCM}}} \\ &\leq \frac{1}{1 - \varepsilon_{\text{SCM}}}, \quad \forall \mu \in \mathcal{E}_{\text{train,SCM}}. \end{aligned}$$

We may then replace our output effectivity upper bounds (97)—though rigorously now only for $\mu \in \mathcal{E}_{\text{train,SCM}}$ —with $\gamma^e(\mu)/((1 - \varepsilon_{\text{SCM}})\alpha^e(\mu))$ (independent of N and $\mathcal{N}$). We conclude that even a rather crude lower bound—we often choose $\varepsilon_{\text{SCM}} = 0.75$ —will have relatively little deleterious effect on our error bounds.

10.1.5 Offline-Online Procedure

We conclude our development of the SCM by summarizing the cost of the Offline and Online computations.

In the Offline stage, the notable computations are (i) $2Q$ (respectively, J_{max}) eigenproblems over $X^{\mathcal{N}}$ to form $\mathcal{B}$ (respectively, to form $\{\alpha^{\mathcal{N}}(\mu') \mid \mu' \in C_{J_{\text{max}}}\}$), (ii) $J_{\text{max}}Q$ inner products over $X^{\mathcal{N}}$ to subsequently form $\{y^*(\mu') \mid \mu' \in C_{J_{\text{max}}}\}$, and (iii) $n_{\text{train,SCM}}J_{\text{max}}$ lower bound LP’s of “size” $2Q + M$ (as well as associated upper bound enumerations) to perform the “arg max”. The total computation cost thus roughly scales as $O(\mathcal{N}(2Q + J_{\text{max}})) + O(\mathcal{N}QJ_{\text{max}}) + O(n_{\text{train,SCM}}J_{\text{max}}QM)$, which of course depends on $\mathcal{N}$; however, we note that there is no “cross-term” $O(n_{\text{train,SCM}}\mathcal{N})$, and thus in practice we can choose both $n_{\text{train,SCM}}$ and $\mathcal{N}$ very large. (There is a clear analogy with our RB greedy algorithm of Sect. 7.2.2: by replacing the true error with a surrogate, we can perform a more efficient and/or thorough search.)

The eigenproblems associated with the calculation of the $\alpha^{\mathcal{N}}(\mu')$, $\mu' \in C_J$, can be treated very efficiently by the Lanczos method [44, 149]. In particular, it follows from the Rayleigh quotient (8) that $\alpha^{\mathcal{N}}(\mu')$ is the minimum eigenvalue of a generalized eigenproblem: find $(\chi, \lambda) \in X^{\mathcal{N}} \times \mathbb{R}_+$ such that

$$a(\chi, v; \mu) = \lambda(\chi, v)_X, \quad \forall v \in X^{\mathcal{N}}. \quad (117)$$

For the (parameter-independent) choice

$$\tau = \inf_{w \in X^{\mathcal{N}}} \frac{a(w, w; \bar{\mu})}{(w, w)_{L^2(\Omega)}} \quad (118)$$

in (6), it can be shown that $\lambda_{\min} = \lambda_1$ of (117) is “well-separated” from λ_2 ; the latter, in turn, ensures rapid convergence of the Lanczos procedure. As indicated earlier, the choice of norm (6) will not affect the RB output prediction (except indirectly through the greedy process), but will affect the quality and cost of the *a posteriori* output error bound; we now further understand that the choice of $\bar{\mu}$ affects the effectivity (see Sect. 9) while the choice of τ affects Offline (eigenvalue) efficiency.

In the Online stage, for each evaluation $\mu \rightarrow \alpha_{\text{LB}}^{\mathcal{N}}(\mu; C_J, M)$, (i) we first perform a sort of the $J_{\max}$ points in $C_{J_{\max}}$ to determine the set $C_{J_{\max}}^{M, \mu}$ —the operation count is at most $O(M J_{\max})$, (ii) we must next effect $(M + 1)Q$ evaluations $\mu' \rightarrow \Theta^q(\mu')$, $1 \leq q \leq Q$ —the operation count is $O((M + 1)Q)$, (iii) we then extract the selected M members of the pre-computed set $\{\alpha^{\mathcal{N}}(\mu') \mid \mu' \in C_J\}$ and solve the resulting LP to obtain $\alpha_{\text{LB}}^{\mathcal{N}}(\mu; C_J, M)$ —a variety of standard procedures of both the simplex and interior point variety can be applied to this LP of rather modest size. In closing, and most importantly, we note that the operation count for the Online evaluation of $\alpha_{\text{LB}}^{\mathcal{N}}(\mu; C_J, M)$ does not depend on $\mathcal{N}$ and hence we can retain rapid response in the many-query and real-time contexts.

10.2 Numerical Results

We shall next present some numerical results. We first consider the Thermal Block ($B_1 = B_2 = 3$) example of Sects. 6.1.1, 8.2.1, and 9.5. For this problem we choose ($\bar{\mu} = (1 \ 1 \ \dots \ 1)$) and $\tau = 0$: this choice for τ —corresponding to the numerical results of Sect. 9.5—admits very simple analysis. It is readily observed from (59) that the Thermal Block problem is *parametrically coercive*: we may thus construct a lower bound “by hand”; we can also expect rapid convergence of the automated greedy procedure.

We first construct a lower bound by hand: we choose $J = 1$, $C_J = \{\bar{\mu}\} = \{(1 \ 1 \ \dots \ 1)\}$, $M = 1$; the earlier numerical results reported in Sects. 8.2.1 and 9.5 correspond to this simple lower bound. It follows directly from (114) that

$$\begin{aligned} \alpha_{\text{LB}}^{\mathcal{N}}(\mu; C_J, M) &\geq \min(1, \mu_1, \dots, \mu_P) \alpha^{\mathcal{N}}(\bar{\mu}) \\ &= \min(1, \mu_1, \dots, \mu_P) \end{aligned} \quad (119)$$

since $a(w, v; \bar{\mu}) = (w, v)_X$ for $\tau = 0$ and hence $\alpha^{\mathcal{N}}(\bar{\mu}) = 1$.¹³ We have already taken advantage of this estimate in Sect. 9 to derive a simple effectivity upper bound.

In practice, even for parametrically coercive problems, the greedy procedure will improve upon any “by hand”

proposal—ensuring closer adherence of $\alpha_{\text{LB}}^{\mathcal{N}}(\mu; C_J, M)$ to $\alpha^{\mathcal{N}}(\mu)$ and hence, from Proposition 2, better effectivities. We now apply our Offline SCM Greedy algorithm of Sect. 10.1.4 to the Thermal Block problem: we choose $\mu_{\text{SCM}}^1 = \bar{\mu}$, a random sample $\mathcal{E}_{\text{train, SCM}}$ of size $n_{\text{train, SCM}} = 500$, $\varepsilon_{\text{SCM}} = 0.75$, and $M = 64$. We obtain $J_{\max} = 12$ (which effectively “resets” M to 12 since $C_{J_{\max}}^{\mu, 64} = C_{J_{\max}}^{\mu, 12} = C_{J_{\max}}$); we expect that $\alpha^{\mathcal{N}}(\mu)/\alpha_{\text{LB}}^{\mathcal{N}}(\mu; C_{J_{\max}}, M) \approx 4$ for most $\mu \in \mathcal{D}$. As expected, the greedy converges very quickly for this parametrically coercive problem despite the large number of parameters.¹⁴ (In fact, $n_{\text{train, SCM}} = 500$ is rather small for $P = 8$ parameters; however, for parametrically coercive problems, smaller samples typically suffice.)

None of our remaining examples—inviscid flow of Sect. 6.1.2, center crack of Sect. 6.2.2, or elastic block of Sect. 6.2.1—is parametrically coercive, and hence the greedy SCM construction is indispensable; in all the numerical results presented in Sect. 8, the coercivity lower bound is provided by the greedy SCM result $\alpha_{\text{LB}}^{\mathcal{N}}(\mu; C_{J_{\max}}, M)$. We present numerical evidence here for the inviscid flow added mass example: we choose (as in all our examples except Thermal Block) $\bar{\mu} = \mu_{\text{ref}}$ and τ according to (118); the latter ensures rapid convergence of the requisite Offline eigenproblems.

We consider a random sample $\mathcal{E}_{\text{train, SCM}}$ of size $n_{\text{train, SCM}} = 500$, $\varepsilon_{\text{SCM}} = 0.75$, and $M = 8$. We obtain $J_{\max} = 63$: as expected for this non-parametrically coercive problem, we require at least a few “coercivity constraint” sample points in each parameter direction in order to ensure compliance with our error tolerance. However, the Offline effort is still not excessive: thanks to our choice of τ , the number of Lanczos iterations “per $\alpha^{\mathcal{N}}(\mu')$ ” is quite modest, typically ≈ 10 to achieve a relative accuracy of 10^{-2} . Furthermore, the Online effort is very insensitive to $J_{\max}$, since the size of the LP is dictated solely by Q and M .

We close with several general comments. First, for problems with many parameters that are *not* parametrically coercive, we must choose $\mathcal{E}_{\text{train, SCM}}$ quite large—larger than $n_{\text{train, SCM}} = 500$ —to obtain a lower bound $\alpha_{\text{LB}}^{\mathcal{N}}(\mu)$ that is viable (positive) for all μ in $\mathcal{D}$. Second, as already noted, we do not need high accuracy for $\alpha_{\text{LB}}^{\mathcal{N}}$: an increase in ε_{SCM} from ≈ 0 to 0.75 will only increase our output bound effectivities by a factor of roughly four. Third, we note that M can be chosen to balance Offline and Online effort: M very large will economize Offline performance (reduce $J_{\max}$) but degrade Online performance, while M very small will increase Offline cost (increase $J_{\max}$) but improve Online response; typically, M should be chosen such that $C_J^{M, \mu}$ “covers” variations in all the parametric coordinates.

¹³In fact, the numerical results of Sects. 8.2.1 and 9.5 correspond to a coercivity lower bound $\min(1, \mu_1, \dots, \mu_P)$ which is, from (119), a simple analytical (positive) conservative estimate for the SCM prediction $\alpha_{\text{LB}}^{\mathcal{N}}(\mu; C_J, M)$ [112].

¹⁴However, our choice $\tau = 0$, though theoretically convenient, yields poor Lanczos convergence; in practice (see below), we choose τ according to (118).

Finally, we note that we must anticipate that $J_{\max}$, unlike $N_{\max}$, will increase exponentially with P —at least for problems that are not parametrically coercive. This will certainly limit the number of parameters that may be considered; in the non-coercive case [64] the effect will be even more pronounced. However, the “constants” appear to be quite favorable, at least for coercive problems; for all of our working examples in this paper, $J_{\max} < 100$.

11 Extension to Noncompliant Problems

11.1 General Primal-Dual Approximation Framework

We now consider the more general non-compliant problem: given $\mu \in \mathcal{D}$, find

$$s(\mu) = \ell(u^e(\mu)), \quad (120)$$

where $u^e(\mu) \in X^e$ satisfies

$$a(u^e(\mu), v; \mu) = f(v), \quad \forall v \in X^e. \quad (121)$$

We assume that a is coercive and continuous (and affine, (2)) but not necessarily symmetric. We further assume that both ℓ and f are bounded functionals but we no longer require $\ell = f$.¹⁵ We shall also need the dual problem associated to ℓ : find $\psi^e(\mu) \in X^e$ such that

$$a(v, \psi^e(\mu); \mu) = -\ell(v), \quad \forall v \in X^e;$$

ψ^e is denoted the “adjoint” or “dual” field.

Our FE approximation is then given by

$$s^{\mathcal{N}}(\mu) = \ell(u^{\mathcal{N}}(\mu)),$$

where

$$a(u^{\mathcal{N}}(\mu), v; \mu) = f(v), \quad \forall v \in X^{\mathcal{N}},$$

and

$$a(v, \psi^{\mathcal{N}}(\mu); \mu) = -\ell(v), \quad \forall v \in X^{\mathcal{N}}.$$

For our purposes here a single FE space suffices for both the primal and dual. In actual FE practice—in which the dual approximation serves to improve output accuracy [117] or develop output error bounds [23, 108]—the FE primal and dual spaces may be different.

By way of motivation, we first consider approximation of the primal problem given by (120), (121) under the stated

hypotheses on ℓ , f , and a . We shall require two approximation subspaces $\tilde{X}_1^{\text{pr}} \subset X^{\mathcal{N}}$ and $\tilde{X}_2^{\text{pr}}(\mu) \subset X^{\mathcal{N}}$ (possibly parameter dependent) and a given function $\Phi: \mathcal{D} \rightarrow X^{\mathcal{N}}$. Ultimately $\tilde{X}_1^{\text{pr}}$ and $\tilde{X}_2^{\text{pr}}$ shall be interpreted as RB spaces; however, at present the particular nature of these approximation subspaces is not crucial.

We then introduce a general Petrov-Galerkin approximation: Given $\mu \in \mathcal{D}$, find $\tilde{u}^{\mathcal{N}}(\mu) \in \tilde{X}_1^{\text{pr}}$ such that

$$a(\tilde{u}^{\mathcal{N}}(\mu), v; \mu) = f(v), \quad \forall v \in \tilde{X}_2^{\text{pr}}(\mu), \quad (122)$$

and evaluate the output as

$$\tilde{s}^{\mathcal{N}}(\mu) = \ell(\tilde{u}^{\mathcal{N}}(\mu)) - \tilde{r}^{\text{pr}}(\Phi(\mu); \mu), \quad (123)$$

where

$$\tilde{r}^{\text{pr}}(v; \mu) \equiv f(v) - a(\tilde{u}^{\mathcal{N}}(\mu), v; \mu), \quad \forall v \in X^e, \quad (124)$$

is the primal residual. Ultimately, $\tilde{u}^{\mathcal{N}}(\mu)$ and $\tilde{s}^{\mathcal{N}}(\mu)$ shall be interpreted as our RB approximations to the field $u^{\mathcal{N}}(\mu)$ and output $s^{\mathcal{N}}(\mu)$, respectively.

We note that for coercive problems it shall suffice to consider standard Galerkin RB approximation— $\tilde{X}_2^{\text{pr}}(\mu) = \tilde{X}_1^{\text{pr}}$. (For non-coercive problems, consideration of Petrov-Galerkin RB approximations is much more interesting [86, 134].) However, the different roles of the “trial” and “test” spaces is best illuminated in the full Petrov-Galerkin context; we hence consider the more general framework for the purposes of this conceptual discussion.

We define the discrete coercivity, inf-sup, and continuity constants associated with our spaces as

$$\tilde{\alpha}(\mu) = \inf_{w \in \tilde{X}_1^{\text{pr}}} \frac{a(w, w; \mu)}{\|w\|_X^2}, \quad (125)$$

$$\tilde{\beta}(\mu) = \inf_{w \in \tilde{X}_1^{\text{pr}}} \sup_{v \in \tilde{X}_2^{\text{pr}}(\mu)} \frac{a(w, v; \mu)}{\|w\|_X \|v\|_X}, \quad (126)$$

and

$$\tilde{\gamma}(\mu) = \sup_{w \in \tilde{X}_1^{\text{pr}}} \sup_{v \in \tilde{X}_2^{\text{pr}}(\mu)} \frac{a(w, v; \mu)}{\|w\|_X \|v\|_X}, \quad (127)$$

respectively. Our norm $\|\cdot\|_X$ is defined, as before, by (6) of Sect. 4.

We can then prove (under the assumption that a is coercive and continuous)

Proposition 4 For $\tilde{u}^{\mathcal{N}}(\mu)$ and $\tilde{s}^{\mathcal{N}}(\mu)$ given by (122) and (123), respectively,

$$\begin{aligned} & \|u^{\mathcal{N}}(\mu) - \tilde{u}^{\mathcal{N}}(\mu)\|_X \\ & \leq \left(1 + \frac{\gamma^e(\mu)}{\tilde{\beta}(\mu)}\right) \inf_{\tilde{w} \in \tilde{X}_1^{\text{pr}}} \|u^{\mathcal{N}}(\mu) - \tilde{w}\|_X, \end{aligned} \quad (128)$$

¹⁵Typical output functionals correspond to the “integral” of the field $u(\mu)$ over an area or line (in particular, boundary segment) in $\overline{\Omega}$. However, by appropriate lifting techniques [1, 56], “integrals” of the flux over boundary segments can also be considered.

and

$$\begin{aligned}
& |s^{\mathcal{N}}(\boldsymbol{\mu}) - \tilde{s}^{\mathcal{N}}(\boldsymbol{\mu})| \\
& \leq \gamma^e(\boldsymbol{\mu}) \|u^{\mathcal{N}}(\boldsymbol{\mu}) - \tilde{u}^{\mathcal{N}}(\boldsymbol{\mu})\|_X \\
& \quad \times \inf_{\tilde{v} \in \tilde{X}_2^{\text{pr}}(\boldsymbol{\mu})} \|\psi^{\mathcal{N}}(\boldsymbol{\mu}) - \Phi(\boldsymbol{\mu}) - \tilde{v}\|_X. \tag{129}
\end{aligned}$$

Here $\Phi: \mathcal{D} \rightarrow X^{\mathcal{N}}$ is a (any) given function.

The proof is standard.

We observe that $\tilde{X}_2^{\text{pr}}(\boldsymbol{\mu})$ in fact plays two roles. First, (for given $\tilde{X}_1^{\text{pr}}$ chosen based on approximation considerations) $\tilde{X}_2^{\text{pr}}(\boldsymbol{\mu})$ determines, through $\tilde{\beta}(\boldsymbol{\mu})$, the *stability* of our approximation—both as regards convergence and also algebraic conditioning. (For coercive problems we can ensure stability simply through the Galerkin recipe, $\tilde{X}_2^{\text{pr}}(\boldsymbol{\mu}) = \tilde{X}_1^{\text{pr}}$; for non-coercive problems, alternative non-Galerkin choices for $\tilde{X}_2^{\text{pr}}(\boldsymbol{\mu})$ can be advantageous.) Note that stability is a significant concern in the RB context since the RB spaces do not have general approximation properties. Second, $\tilde{X}_2^{\text{pr}}(\boldsymbol{\mu})$ can improve the *accuracy* of the output prediction: if members of $\tilde{X}_2^{\text{pr}}(\boldsymbol{\mu})$ approximate well the adjoint $\psi^{\mathcal{N}}(\boldsymbol{\mu})$, then the second term in (129) will be small—quite independent of the primal approximation properties of $\tilde{X}_1^{\text{pr}}$ as reflected in (128) and the first term in (129).

This double role of the test space informs the choice of $\tilde{X}_2^{\text{pr}}(\boldsymbol{\mu})$ and the introduction/selection of $\Phi(\boldsymbol{\mu})$. In particular, if we choose $\tilde{X}_2^{\text{pr}}(\boldsymbol{\mu})$ solely based on stability considerations, we forego the approximation advantages associated with the second term in (129). Conversely, if we choose $\tilde{X}_2^{\text{pr}}(\boldsymbol{\mu})$ solely based on adjoint approximation considerations, we shall have no guarantee of stability. It is the additional “degree of freedom” afforded by $\Phi(\boldsymbol{\mu})$ that can resolve the dilemma: we will typically select $\tilde{X}_2^{\text{pr}}(\boldsymbol{\mu})$ to control stability—for coercive problems, by the simple Galerkin recipe $\tilde{X}_2^{\text{pr}}(\boldsymbol{\mu}) = \tilde{X}_1^{\text{pr}}$ —and select $\Phi(\boldsymbol{\mu})$ to control better the second term in (129).

The choice of $\Phi(\boldsymbol{\mu})$ will in most cases take the form of a Petrov-Galerkin (Galerkin for coercive problems) approximation to the *dual* problem: Given $\boldsymbol{\mu} \in \mathcal{D}$, $\tilde{\psi}^{\mathcal{N}}(\boldsymbol{\mu}) \equiv \Phi(\boldsymbol{\mu}) \in \tilde{X}_1^{\text{du}}$ satisfies

$$a(v, \tilde{\psi}^{\mathcal{N}}(\boldsymbol{\mu}); \boldsymbol{\mu}) = -\ell(v), \quad \forall v \in \tilde{X}_2^{\text{du}}(\boldsymbol{\mu}). \tag{130}$$

Here the space $\tilde{X}_1^{\text{du}} \subset X^{\mathcal{N}}$ is chosen to provide good approximation of $\psi^{\mathcal{N}}(\boldsymbol{\mu})$, and the space $\tilde{X}_2^{\text{du}}(\boldsymbol{\mu})$ (possibly parameter-dependent) is chosen to provide good stability of the discrete dual problem. As regards the former, we emphasize that in the *ad hoc* RB context, even more so than in the “generic” FE context, the dual and primal approximation spaces may be fundamentally different—with very little shared approximation properties. As regards the latter, we again observe that, in the coercive case, the Galerkin recipe $\tilde{X}_2^{\text{du}} = \tilde{X}_1^{\text{du}}$ is certainly a viable approach.

11.2 The Reduced Basis Context

At this stage we can now motivate (roughly) the computational motivation for, and benefit of, the adjoint.

11.2.1 Approximation

We choose the primal and dual approximation spaces according to the (greedy) Lagrange RB prescription: $\tilde{X}_1^{\text{pr}} = W_{N^{\text{pr}}}^{\text{pr}}$ —the span of N^{pr} snapshots of $u^{\mathcal{N}}(\boldsymbol{\mu})$, $\boldsymbol{\mu} \in \mathcal{D}$; $\tilde{X}_1^{\text{du}} = W_{N^{\text{du}}}^{\text{du}}$ —the span of N^{du} snapshots of $\psi^{\mathcal{N}}(\boldsymbol{\mu})$, $\boldsymbol{\mu} \in \mathcal{D}$. We choose the corresponding primal and dual test spaces (perforce stably) according to the Galerkin recipe: $\tilde{X}_2^{\text{pr}} = W_{N^{\text{pr}}}^{\text{pr}}$; $\tilde{X}_2^{\text{du}} = W_{N^{\text{du}}}^{\text{du}}$.

The resulting RB primal approximation— $\tilde{u}^{\mathcal{N}}(\boldsymbol{\mu})$ of (122) for our particular RB primal spaces—will be denoted $u_{N^{\text{pr}}}^{\mathcal{N}}(\boldsymbol{\mu})$; the resulting RB dual approximation— $\tilde{\psi}^{\mathcal{N}}$ of (130) for our particular RB dual spaces—will be denoted $\psi_{N^{\text{du}}}^{\mathcal{N}}(\boldsymbol{\mu})$. (We shall adopt the convention that $N^{\text{pr}} = 0$ and $N^{\text{du}} = 0$ corresponds to $u_{N^{\text{pr}}}^{\mathcal{N}} = 0$ and $\psi_{N^{\text{du}}}^{\mathcal{N}} = 0$, respectively.) Finally, the associated RB output approximation— $\tilde{s}^{\mathcal{N}}$ of (123) for $\Phi(\boldsymbol{\mu}) = \psi_{N^{\text{du}}}^{\mathcal{N}}(\boldsymbol{\mu})$ —will take the form

$$s_{N^{\text{pr}}, N^{\text{du}}}^{\mathcal{N}}(\boldsymbol{\mu}) = \ell(u_{N^{\text{pr}}}^{\mathcal{N}}(\boldsymbol{\mu})) - r^{\text{pr}}(\psi_{N^{\text{du}}}^{\mathcal{N}}(\boldsymbol{\mu}); \boldsymbol{\mu}). \tag{131}$$

Recall that r^{pr} is the RB primal residual as defined in (88).

Next, and solely for the purposes of our arguments here, we suppose that

$$\|u^{\mathcal{N}}(\boldsymbol{\mu}) - u_{N^{\text{pr}}}^{\mathcal{N}}(\boldsymbol{\mu})\|_X = \|u^{\mathcal{N}}(\boldsymbol{\mu})\|_X g_{\text{err}}(N^{\text{pr}}), \tag{132}$$

and

$$\|\psi^{\mathcal{N}}(\boldsymbol{\mu}) - \psi_{N^{\text{du}}}^{\mathcal{N}}(\boldsymbol{\mu})\|_X = \|\psi^{\mathcal{N}}(\boldsymbol{\mu})\|_X g_{\text{err}}(N^{\text{du}}). \tag{133}$$

Here $g_{\text{err}}: \mathbb{N}_0 \rightarrow \mathbb{R}$ is a presumed monotonically decreasing “convergence” function with associated inverse g_{err}^{-1} (such that $g_{\text{err}}^{-1}(g_{\text{err}}(N)) = N$); note it follows from our conventions for $N^{\text{pr}} = 0$ and $N^{\text{du}} = 0$ that $g(0) = 1$. (Implicitly in (132), (133) we presume—given our single primal/dual convergence function—that the primal and dual problems are of similar regularity or “difficulty”; we shall return to this point below.) Finally, we shall suppose—as motivated by the result (129)—that

$$\begin{aligned}
& |s^{\mathcal{N}}(\boldsymbol{\mu}) - s_{N^{\text{pr}}, N^{\text{du}}}^{\mathcal{N}}(\boldsymbol{\mu})| \\
& = C_s \|u^{\mathcal{N}}(\boldsymbol{\mu})\|_X \|\psi^{\mathcal{N}}(\boldsymbol{\mu})\|_X g_{\text{err}}(N^{\text{pr}}) g_{\text{err}}(N^{\text{du}}), \tag{134}
\end{aligned}$$

where C_s is a constant that reflects the non-zero angle between the primal and dual errors.

It then follows from our assumptions that, at fixed output error $|s^{\mathcal{N}}(\boldsymbol{\mu}) - s_{N^{\text{pr}}, N^{\text{du}}}^{\mathcal{N}}(\boldsymbol{\mu})| = C_s \|u^{\mathcal{N}}(\boldsymbol{\mu})\|_X \|\psi^{\mathcal{N}}(\boldsymbol{\mu})\|_X \varepsilon$,

Online Cost of $s_N^{\mathcal{N}}(\boldsymbol{\mu})$ WITH ADJOINT
Online Cost of $s_N^{\mathcal{N}}(\boldsymbol{\mu})$ WITHOUT ADJOINT

$$= 2 \left(\frac{g_{\text{err}}^{-1}(\sqrt{\varepsilon})}{g_{\text{err}}^{-1}(\varepsilon)} \right)^3; \quad (135)$$

here “WITH ADJOINT” refers to the optimal—given (132), (133)—choice $N^{\text{pr}} = N^{\text{du}}$, and “WITHOUT ADJOINT” refers to the choice $N^{\text{pr}} \neq 0$, $N^{\text{du}} = 0$ (and hence, according to our convention, $\psi_{N^{\text{pr}}}^{\mathcal{N}} = 0$). The cube in (135) arises from the assumed dominant LU decomposition of the (dense) RB stiffness matrices as described in Sect. 7. Even in the “worst” case, in which $g_{\text{err}}(N)$ is exponential— $g_{\text{err}}(N) = e^{-\omega N}$ for given positive real ω —the ratio in (135) is $\frac{1}{4}$: by “normal” standards, a reduction in effort of 75% is very large. Furthermore, if convergence is only algebraic— $g_{\text{err}}(N) = (N+1)^{-\omega}$ —the ratio (135) is approximately $2\varepsilon^{3/(2\omega)}$, and hence the reduction in effort is even more significant (in particular for smaller ε).

In actual practice the best (most efficient) decomposition between primal and dual degrees of freedom will be strongly problem dependent. In particular, we must in general replace g_{err} by $g_{\text{err}}^{\text{pr}}$ in (132) and by $g_{\text{err}}^{\text{du}}$ in (133); the optimal choice of N^{pr} , N^{du} —the choice that minimizes Online effort—will in general depend on the ratio between $g_{\text{err}}^{\text{pr}}$ and $g_{\text{err}}^{\text{du}}$. For example, if $g_{\text{err}}^{\text{pr}} \ll g_{\text{err}}^{\text{du}}$ then we should choose $N^{\text{pr}} \neq 0$, $N^{\text{du}} = 0$; conversely, if $g_{\text{err}}^{\text{du}} \ll g_{\text{err}}^{\text{pr}}$, then we should choose $N^{\text{pr}} = 0$, $N^{\text{du}} \neq 0$. We further note that these arguments will be influenced by the number of outputs: clearly for *many* outputs the adjoint becomes a good idea that we can no longer afford, largely independent of the relative convergence rates of the primal and dual RB approximations.

11.2.2 A Posteriori Error Estimation

Finally, we briefly discuss *a posteriori* error estimation in order to provide yet another motivation for the adjoint. We first consider the case “without adjoint,” $N^{\text{du}} = 0$, corresponding to $\psi_{N^{\text{du}}}^{\mathcal{N}} = 0$ in (131) and hence $s_{N^{\text{pr}}, N^{\text{du}}}^{\mathcal{N}}(\mu) = \ell(u_{N^{\text{pr}}}^{\mathcal{N}}(\mu))$ and

$$|s^{\mathcal{N}}(\mu) - s_{N^{\text{pr}}, N^{\text{du}}}^{\mathcal{N}}(\mu)| = |\ell(u_{N^{\text{pr}}}^{\mathcal{N}}(\mu) - u^{\mathcal{N}}(\mu))|. \quad (136)$$

Our *a posteriori* error bound in this case shall (must) in general take the form

$$\Delta_N^s(\mu) = \|\ell\|_{(X^{\mathcal{N}})'} \Delta_N(\mu), \quad (137)$$

where $\Delta_N(\mu)$ is the error bound for $\|u^{\mathcal{N}}(\mu) - u_{N^{\text{pr}}}^{\mathcal{N}}(\mu)\|_X$. It is clear that (137) is indeed a rigorous bound for the error in the output.

However, we can demonstrate that the bound (137) can, in some situations, be very poor. In particular, we consider the (compliant) case $\ell = f$: since from (66) $|s^{\mathcal{N}}(\mu) - s_{N^{\text{pr}}, N^{\text{du}}}^{\mathcal{N}}(\mu)| \leq \gamma^e(\mu) \|u^{\mathcal{N}}(\mu) - u_{N^{\text{pr}}}^{\mathcal{N}}(\mu)\|_X^2$, and

from Proposition 3 (extended to the X norm [112]) $\Delta_N(\mu) \geq \|u^{\mathcal{N}}(\mu) - u_{N^{\text{pr}}}^{\mathcal{N}}(\mu)\|_X$, we obtain

$$\frac{\Delta_N^s(\mu)}{|s^{\mathcal{N}}(\mu) - s_{N^{\text{pr}}, N^{\text{du}}}^{\mathcal{N}}(\mu)|} \geq \frac{\|\ell\|_{(X^{\mathcal{N}})'}}{\gamma^e(\mu) \|u^{\mathcal{N}}(\mu) - u_{N^{\text{pr}}}^{\mathcal{N}}(\mu)\|_X}; \quad (138)$$

hence the effectivity of the output error bound (137) *tends to infinity* as $(N \rightarrow \infty)$ and $u_{N^{\text{pr}}}^{\mathcal{N}}(\mu) \rightarrow u^{\mathcal{N}}(\mu)$. We expect similar behavior for any ℓ “close” to f : the failing is that (137) does not reflect the contribution of the test space to the convergence of the output.

The introduction of the RB adjoint approximation will largely cure this problem—and ensure a stable limit $N \rightarrow \infty$. In particular, in the non-compliant case, the output error bound takes the form

$$\Delta_N^s(\mu) \equiv \frac{\|r^{\text{pr}}(\cdot; \mu)\|_{(X^{\mathcal{N}})'}}{(\alpha_{\text{LB}}^{\mathcal{N}}(\mu))^{1/2}} \frac{\|r^{\text{du}}(\cdot; \mu)\|_{(X^{\mathcal{N}})'}}{(\alpha_{\text{LB}}^{\mathcal{N}}(\mu))^{1/2}}, \quad (139)$$

where

$$r^{\text{du}}(v; \mu) \equiv -\ell(v) + a(v, \psi_{N^{\text{du}}}^{\mathcal{N}}(\mu); \mu), \quad \forall v \in X^e, \quad (140)$$

is the dual residual. In effect, (139) is the product of primal and dual errors (with the coercivity constant detributed “symmetrically” for the purposes of the separate primal and dual greedy algorithms).

The effect of the test space, “through” $\Phi = \psi_{N^{\text{du}}}^{\mathcal{N}}$, is now reflected in our error bound (139): we recover, in the compliant case, our earlier (N -independent) effectivity bound of Proposition 2. Unfortunately, even with the adjoint, the effectivities—though bounded—may be quite large, in particular in the very *non*-compliant limit in which the primal and dual errors can be very uncorrelated.

11.3 An Advection-Diffusion Problem

We consider an advection-diffusion example in a rectangular domain $\Omega_0(\mu) =]0, L[\times]0, 1[$ representing a channel. The governing equation for the passive-scalar field (say, temperature) is the advection-diffusion equation with imposed Couette velocity $(x_{02}, 0)$. Neumann (flux) boundary conditions are imposed on the bottom wall $\Gamma_{0,\text{bot}}$; homogeneous Dirichlet conditions are imposed on the top wall $\Gamma_{0,\text{top}}$ and on the “inflow” left boundary $\Gamma_{0,\text{in}}$; and homogeneous (zero flux) Neumann conditions are imposed on the “outflow” right boundary $\Gamma_{0,\text{out}}$. The output of interest is the integral of the temperature over the heated (bottom) surface $\Gamma_{0,\text{bot}}$. This example is a simplified version of a Couette-Graetz problem [140].

We consider two parameters: the length of the channel, L , and the Peclet number, Pe [57]. Hence $P = 2$ and $\mu = (\mu_1, \mu_2)$: μ_1 is the channel length L , and μ_2 is the

Peclet number Pe ; the parameter domain is given by $\mathcal{D} = [1, 10] \times [0.1, 100]$. We now choose $\mu_{\text{ref}} = (1, 1)$, which in turn defines the reference domain $\Omega = \Omega_o(\mu_{\text{ref}})$.

In terms of the original domain, we identify $X_o^e(\mu) \equiv \{v \in H^1(\Omega_o(\mu)) | v|_{\Gamma_{o,\text{top}} \cup \Gamma_{o,\text{in}}} = 0\}$, which imposes the Dirichlet conditions;

$$a_o(w, v; \mu) = \int_{\Omega_o(L)} x_{o2} \frac{\partial w}{\partial x_1} v + \frac{1}{Pe} \int_{\Omega_o(L)} \nabla w \cdot \nabla v, \quad (141)$$

which represents the advection (non-symmetric) and diffusion (symmetric) contributions as well as the homogeneous Neumann conditions; and

$$f_o(v) = \ell_o(v) \equiv \int_{\Gamma_{o,\text{bot}}(L)} v, \quad (142)$$

which reflects the inhomogeneous Neumann conditions and bottom-wall temperature output functional.

In this case we effect the simple dilation-in- x_{o1} affine mapping— $(x_1, x_2) = (x_{o1}/L, x_{o2})$ —“by hand”: we directly obtain the reference domain abstraction (120), (121) for $X^e \equiv \{v \in H^1(\Omega) | v|_{\Gamma_{\text{top}} \cup \Gamma_{\text{in}}} = 0\}$,

$$\begin{aligned} a(w, v, \mu) &\equiv \int_{\Omega} x_2 \frac{\partial w}{\partial x_1} v + \frac{1}{\mu_1 \mu_2} \int_{\Omega} \frac{\partial w}{\partial x_1} \frac{\partial v}{\partial x_1} \\ &\quad + \frac{\mu_1}{\mu_2} \int_{\Omega} \frac{\partial w}{\partial x_2} \frac{\partial v}{\partial x_2} \end{aligned} \quad (143)$$

and

$$f(v; \mu) = \ell(v; \mu) \equiv \mu_1 \int_{\Gamma_{\text{bot}}} v.$$

(We observe that ℓ and f depend affinely on μ ; as in the compliant case, this extension is straightforward.) Note that although $\ell = f$ this problem is *not* compliant since a is *not* symmetric and hence $u^e(\mu) \neq \psi^e(\mu)$: we must invoke the full primal-dual technology described in the previous two sections.

We construct the primal and dual hierarchical RB approximations by (two separate) greedy procedures analogous to the method described in Sect. 7.2.2 for the (primal only) compliant case. We consider for $\mathcal{E}_{\text{train}}$ a random sample of size $n_{\text{train}} = 2000$. We plot in Fig. 30 the maximum of the *a posteriori* output error bound over $\mathcal{E}_{\text{train}}$ as a function of $N = N_{\text{pr}} = N_{\text{du}}$. (In actual practice, in the Offline stage we can form “look-up” tables of cost-optimal $N_{\text{pr}}, N_{\text{du}}$ discretization pairs as a function of maximum output error (bound) over $\mathcal{D}$; in the Online stage we then search this table to select an RB primal-dual discretization pair $N_{\text{pr}}, N_{\text{du}}$ appropriate for the desired accuracy [44].) We observe that the output converges quite rapidly despite the rather large parametric variations in the channel length and in particular the Peclet number.

Fig. 30 Advection-Diffusion problem: $\max_{\mu \in \mathcal{E}_{\text{train}}} \Delta_N^s(\mu)$ as a function of N ; note $|s^N(\mu) - s_N^N(\mu)| \leq \Delta_N^s(\mu)$

Fig. 31 Advection-Diffusion problem: $\alpha_{\text{UB}}^N(\mu)$ (upper curve, solid) and $\alpha_{\text{LB}}^N(\mu)$ (lower curve, dotted) as a function of $\mu' \in \mathcal{E}_{\text{train,SCM}}$ after $J_{\text{max}} = 4$ iterations of the SCM greedy algorithm; here the abscissa represents the index of the point $\mu_{\text{train,SCM}}^k$ in $\mathcal{E}_{\text{train,SCM}}$

We next consider *a posteriori* error estimation for the output. In fact, our procedures of Sects. 9 and 10 for the symmetric case require very little modification. Our output error bound in the noncompliant case is given by (139): we now must perform two dual norm calculations—one for the primal residual, (88), and one for the dual residual, (140); however, the fundamental Offline-Online procedure described in Sect. 9.4 requires no modification. Similarly, our SCM procedure is little changed: to construct a lower bound for the coercivity constant of a *non*-symmetric bilinear form a , we simply apply the SCM procedure of Sect. 10.1 but now to the *symmetric part* of a . For a random sample $\mathcal{E}_{\text{train,SCM}}$ of size $n_{\text{train,SCM}} = 1000$, $\varepsilon_{\text{SCM}} = 0.75$, and $M = 4$ we obtain $J_{\text{max}} = 4$; the symmetric part of a is parametrically coercive—hence the rapid convergence. We plot in Fig. 31 the SCM lower and upper bounds for the coercivity constant.

We present in Table 2 the output error bound (139) and associated effectivity measures (see Sect. 9.5) as a function

Table 2 Advection-Diffusion problem: Output error bound and effectivity metrics as a function of N

N	$\Delta_{N,\max}^s$	$\eta_{N,\max}^s$	$\eta_{N,\text{ave}}^s$
10	1.9E-01	63.1	7.9
15	5.3E-02	46.78	9.3
25	4.0E-03	48.5	5.9
33	1.0E-03	94.3	8.2
40	2.5E-04	81.4	17.8

of $N = N_{\text{pr}} = N_{\text{du}}$ for a random test sample $\mathcal{E}$ of size 2000. We note that in the non-compliant case (without further assumptions on the output convergence) we can no longer obtain an upper bound for the output error bound effectivity, and hence we anticipate that the effectivities might be larger and also perhaps more erratic than in the compliant case; this expectation is (unfortunately) realized in Table 2, though in fact at least the average effectivities remain quite well-behaved. (We note that the higher values of the effectivity obtain for the higher Peclet values—for which the continuity/coercivity ratio is larger and furthermore the primal and dual solutions are less correlated.)

In other studies, the time-dependent advection-diffusion (parabolic) equation is considered [52]. Future work must also address RB approximation of *stabilized* FE approximations (for high Peclet number) [43, 124].

References

1. rbMIT Software (2007) http://augustine.mit.edu/methodology/methodology_rbMIT_System.htm. MIT, Cambridge
2. Ainsworth M, Oden JT (1997) *A posteriori* error estimation in finite element analysis. Comput Methods Appl Mech Eng 142:1–88
3. Ainsworth M, Oden JT (2000) *A posteriori* error estimation in finite element analysis. Wiley-Interscience, New York
4. Almroth BO, Stern P, Brogan FA (1978) Automatic choice of global shape functions in structural analysis. AIAA J 16:525–528
5. Anderson TL (2005) Fracture mechanics: fundamentals and application, 3rd edn. CRC Press, Boca Raton
6. Arpaci VS (1966) Conduction heat transfer. Addison-Wesley, Reading
7. Arpaci VS, Larsen PS (1984) Convection heat transfer. Prentice Hall, Englewood Cliffs
8. Atwell JA, King BB (2001) Proper orthogonal decomposition for reduced basis feedback controllers for parabolic equations. Math Comput Model 33(1–3):1–19
9. Babuška I (1971) Error-bounds for finite element method. Numer Math 16:322–333
10. Babuška I, Osborn J (1991) Eigenvalue problems. In: Handbook of numerical analysis, vol II. Elsevier, Amsterdam, pp 641–787
11. Babuška I, Rheinboldt W (1978) *A posteriori* error estimates for the finite element method. Int J Numer Methods Eng 12:1597–1615
12. Babuška I, Rheinboldt W (1978) Error estimates for adaptive finite element computations. SIAM J Numer Anal 15:736–754
13. Babuška I, Strouboulis T (2001) The finite element method and its reliability. Numerical mathematics and scientific computation. Clarendon, Oxford
14. Bai ZJ (2002) Krylov subspace techniques for reduced-order modeling of large-scale dynamical systems. Appl Numer Math 43(1–2):9–44
15. Balmes E (1996) Parametric families of reduced finite element models: Theory and applications. Mech Syst Signal Process 10(4):381–394
16. Balsa-Canto E, Alonso A, Banga J (2004) Reduced-order models for nonlinear distributed process systems and their application in dynamic optimization. Ind Eng Chem Res 43(13):3353–3363
17. Banks HT, Kunisch K (1989) Estimation techniques for distributed parameter systems. Systems & control: foundations & applications. Birkhäuser, Boston
18. Barrault M, Nguyen NC, Maday Y, Patera AT (2004) An “empirical interpolation” method: Application to efficient reduced-basis discretization of partial differential equations. C R Acad Sci Paris, Sér I 339:667–672
19. Barrett A, Reddien G (1995) On the reduced basis method. Z Angew Math Mech 75(7):543–549
20. Barsom JM, Rolfe ST (1999) Fracture and fatigue control in structures. American society for testing and metals. Butterworth, Stoneham
21. Bashir O, Willcox K, Ghattas O, van Bloemen Waanders B, Hill J (2008) Hessian-based model reduction for large-scale systems with initial condition inputs. Int J Numer Methods Eng 73(6):844–868
22. Bathe KJ (1996) Finite element procedures. Prentice Hall, Englewood Cliffs
23. Becker R, Rannacher R (1996) A feedback approach to error control in finite element method: Basic analysis and examples. East-West J Numer Math 4:237–264
24. Benner P, Mehrmann V, Sorensen D (eds) (2003) Dimension reduction of large-scale systems. Lecture notes in computational science and engineering. Springer, Heidelberg
25. Bensoussan A, Lions JL, Papanicolaou G (1978) Asymptotic analysis of periodic structures. North-Holland, Amsterdam
26. Boyaval S (2007) Application of reduced basis approximation and *a posteriori* error estimation to homogenization theory. Multiscale Model Simul (to appear)
27. Braess D (2001) Finite elements. Theory, fast solvers, and applications in solid mechanics. Cambridge University Press, Cambridge
28. Brezzi F (1974) On the existence, uniqueness, and approximation of saddle point problems arising from Lagrangian multipliers. RAIRO Anal Numer 2:129–151
29. Brezzi F, Fortin M (1991) Mixed and hybrid finite element methods. Springer series in computational mathematics, vol 15. Springer, Berlin
30. Brezzi F, Rappaz J, Raviart P (1980) Finite dimensional approximation of nonlinear problems. Part I: Branches of nonsingular solutions. Numer Math 36:1–25
31. Bui-Thanh T, Damodaran M, Willcox K (2003) Proper orthogonal decomposition extensions for parametric applications in transonic aerodynamics (AIAA Paper 2003-4213). In: Proceedings of the 15th AIAA computational fluid dynamics conference
32. Bui-Thanh T, Willcox K, Ghattas O (2008) Model reduction for large-scale systems with high-dimensional parametric input space. SIAM J Sci Comput (to appear)
33. Bui-Thanh T, Willcox K, Ghattas O (2007) Model reduction for large-scale systems with high-dimensional parametric input space (AIAA Paper 2007-2049). In: Proceedings of the 48th AIAA/ASME/ASCE/AHS/ASC structures, structural dynamics and material conference
34. Caloz G, Rappaz J (1997) Numerical analysis for nonlinear and bifurcation problems. In: Ciarlet P, Lions J (eds) Handbook of

- numerical analysis. Techniques of scientific computing (Part 2), vol V. Elsevier, Amsterdam, pp 487–637
35. Cancès E, Le Bris C, Maday Y, Turinici G (2002) Towards reduced basis approaches in *ab initio* electronic structure computations. *J Sci Comput* 17(1–4):461–469
 36. Cancès E, Le Bris C, Nguyen NC, Maday Y, Patera AT, Pau GSH (2007) Feasibility and competitiveness of a reduced basis approach for rapid electronic structure calculations in quantum chemistry. In: Proceedings of the workshop for high-dimensional partial differential equations in science and engineering (Montreal)
 37. Cazemier W (1997) Proper orthogonal decomposition and low dimensional models for turbulent flows. University of Groningen, Groningen
 38. Chen J, Kang SM (2001) Model-order reduction of nonlinear mems devices through arclength-based Karhunen-Loève decomposition. In: Proceeding of the IEEE international symposium on circuits and systems, vol 2, pp 457–460
 39. Chen Y, White J (2000) A quadratic method for nonlinear model order reduction. In: Proceedings of the international conference on modeling and simulation of microsystems, pp 477–480
 40. Christensen E, Brøns M, Sørensen J (2000) Evaluation of proper orthogonal decomposition-based decomposition techniques applied to parameter-dependent nonturbulent flows. *SIAM J Sci Comput* 21(4):1419–1434
 41. Ciarlet PG (2002) The finite element method for elliptic problems. Classics in applied mathematics, vol 40. SIAM, Philadelphia
 42. Daniel L, Ong C, White J (2002) Geometrically parametrized interconnect performance models for interconnect synthesis. In: Proceedings of the 2002 international symposium on physical design. Assoc Comput Mach, New York, pp 202–207
 43. Dedè L (2008) Advanced numerical methods for the solution of optimal control problems described by pdes with environmental applications. PhD thesis, Politecnico di Milano
 44. Demmel JW (1997) Applied numerical linear algebra. SIAM, Philadelphia
 45. Deparis S (2008) Reduced basis error bound computation of parameter-dependent Navier-Stokes equations by the natural norm approach. *SIAM J Numer Anal* 46:2039
 46. Farle O, Hill V, Nickel P, Dyczij-Edlinger R (2006) Multivariate finite element model order reduction for permittivity or permeability estimation. *IEEE Trans Magn* 42:623–626
 47. Fink JP, Rheinboldt WC (1983) On the error behavior of the reduced basis technique for nonlinear finite element approximations. *Z Angew Math Mech* 63(1):21–28
 48. Fox RL, Miura H (1971) An approximate analysis technique for design calculations. *AIAA J* 9(1):177–179
 49. Ganapathysubramanian S, Zabaras N (2004) Design across length scales: a reduced-order model of polycrystal plasticity for the control of microstructure-sensitive material properties. *Comput Methods Appl Mech Eng* 193:5017–5034
 50. Girault V, Raviart P (1986) Finite element approximation of the Navier-Stokes equations. Springer, Berlin
 51. Goberna MA, Lopez MA (1998) Linear semi-infinite optimization. Wiley, New York
 52. Grepl M (2005) Reduced-basis approximations and *a posteriori* error estimation for parabolic partial differential equations. PhD thesis, Massachusetts Institute of Technology
 53. Grepl MA, Maday Y, Nguyen NC, Patera AT (2007) Efficient reduced-basis treatment of nonaffine and nonlinear partial differential equations. *Model Math Anal Numer*
 54. Grepl MA, Nguyen NC, Veroy K, Patera AT, Liu GR (2007) Certified rapid solution of partial differential equations for real-time parameter estimation and optimization. In: Biegler LT, Ghattas O, Heinkenschloss M, Keyes D, van Wandeers B (eds) Proceedings of the 2nd Sandia workshop of PDE-constrained optimization: Real-time PDE-constrained optimization. SIAM computational science and engineering book series. SIAM, Philadelphia, pp 197–216
 55. Grepl MA, Patera AT (2005) *A Posteriori* error bounds for reduced-basis approximations of parametrized parabolic partial differential equations. *Model Math Anal Numer* 39(1):157–181
 56. Gresho P, Sani R (1998) Incompressible flow and the finite element method: advection-diffusion and isothermal laminar flow. Wiley, New York
 57. Gunzburger MD (1989) Finite element methods for viscous incompressible flows. Academic Press, San Diego
 58. Gunzburger MD (2003) Perspectives in flow control and optimization. Advances in design and control. SIAM, Philadelphia
 59. Gunzburger MD, Peterson J, Shadid JN (2007) Reduced-order modeling of time-dependent PDEs with multiple parameters in the boundary data. *Comput Methods Appl Mech* 196:1030–1047
 60. Haasdonk B, Ohlberger M (2008) Reduced basis method for finite volume approximations of parametrized evolution equations. *Math Model Numer Anal* 42(2):277–302
 61. Huynh DBP (2007) Reduced-basis approximation and application to fracture and inverse problems. PhD thesis, Singapore-MIT Alliance, National University of Singapore
 62. Huynh DBP, Patera AT (2007) Reduced-basis approximation and *a posteriori* error estimation for stress intensity factors. *Int J Numer Methods Eng* 72(10):1219–1259
 63. Huynh DBP, Peraire J, Patera AT, Liu GR (2007) Reduced basis approximation and *a posteriori* error estimation for stress intensity factors: Application to failure analysis. In: Singapore-MIT alliance symposium
 64. Huynh DBP, Rozza G, Sen S, Patera AT (2007) A successive constraint linear optimization method for lower bounds of parametric coercivity and inf-sup stability constants. *C R Acad Sci Paris Ser I* 345:473–478
 65. Isaacson E, Keller HB (1994) Computation of eigenvalues and eigenvectors, analysis of numerical methods. Dover, New York
 66. Ito K, Ravindran SS (1998) A reduced basis method for control problems governed by PDEs. In: Desch W, Kappel F, Kunisch K (eds) Control and estimation of distributed parameter systems. Birkhäuser, Boston, pp 153–168
 67. Ito K, Ravindran SS (1998) A reduced-order method for simulation and control of fluid flows. *J Comput Phys* 143(2):403–425
 68. Ito K, Ravindran SS (2001) Reduced basis method for optimal control of unsteady viscous flows. *Int J Comput Fluid Dyn* 15(2):97–113
 69. Ito K, Schroeter JD (2001) Reduced order feedback synthesis for viscous incompressible flows. *Math Comput Model* 33(1–3):173–192
 70. Jabbar M, Azeman A (2004) Fast optimization of electromagnetic-problems: the reduced-basis finite element approach. *IEEE Trans Magn* 40(4):2161–2163
 71. Johnson CR (1989) A Gershgorin-type lower bound for the smallest singular value. *Linear Algebra Appl* 112:1–7
 72. Karhunen K (1946) Zur spektraltheorie stochastischer prozesse. *Ann Acad Sci Fenn* 37
 73. Kunisch K, Volkwein S (2002) Galerkin proper orthogonal decomposition methods for a general equation in fluid dynamics. *SIAM J Numer Anal* 40(2):492–515
 74. Kwang ATY (2006) Reduced basis methods for 2nd order wave equation: Application to one dimensional seismic problem. Master's thesis, Singapore-MIT Alliance, Computation for Design and Optimization
 75. Le Bris C (2006) Private communication. MIT
 76. Lee MYL (1991) Estimation of the error in the reduced-basis method solution of differential algebraic equations. *SIAM J Numer Anal* 28:512–528

77. LeGresley PA, Alonso JJ (2000) Airfoil design optimization using reduced order models based on proper orthogonal decomposition. In: Fluids 2000 conference and exhibit. Denver, CO (2000). Paper 2000-2545
78. Loeve MM (1955) Probability theory. Van Nostrand, Princeton
79. Løvgrén AE, Maday Y, Rønquist EM (2006) A reduced basis element method for complex flow systems. In: Wesseling P, Oñate E, Periaux J (eds) ECCOMAS CFD 2006 proceedings. TU Delft, Delft
80. Løvgrén AE, Maday Y, Rønquist EM (2006) A reduced basis element method for the steady Stokes problem. *Math Model Numer Anal* 40(3):529–552
81. Løvgrén AE, Maday Y, Rønquist EM (2006) A reduced basis element method for the steady Stokes problem: Application to hierarchical flow systems. *Model Identif Control* 27(2):79–94
82. Løvgrén AE, Maday Y, Rønquist EM (2007) The reduced basis element method for fluid flows. In: Analysis and simulation of fluid dynamics. Advances in mathematical fluid mechanics. Birkhäuser, Boston, pp 129–154
83. Ly H, Tran H (2001) Modeling and control of physical processes using proper orthogonal decomposition. *Math Comput Model* 33:223–236
84. Machiels L, Maday Y, Oliveira IB, Patera A, Rovas D (2000) Output bounds for reduced-basis approximations of symmetric positive definite eigenvalue problems. *C R Acad Sci Paris, Sér I* 331(2):153–158
85. Maday Y, Patera A, Turinici G (2002) *A Priori* convergence theory for reduced-basis approximations of single-parameter elliptic partial differential equations. *J Sci Comput* 17(1–4):437–446
86. Maday Y, Patera AT, Rovas DV (2002) A blackbox reduced-basis output bound method for noncoercive linear problems. In: Cioranescu D, Lions JL (eds) Nonlinear partial differential equations and their applications, Collège de France Seminar, vol XIV. Elsevier, Amsterdam, pp 533–569
87. Maday Y, Patera AT, Turinici G (2002) Global *a priori* convergence theory for reduced-basis approximation of single-parameter symmetric coercive elliptic partial differential equations. *C R Acad Sci Paris, Sér I* 335(3):289–294
88. Meyer CD (2000) Matrix analysis and applied linear algebra. SIAM, Philadelphia
89. Meyer M, Matthies HG (2003) Efficient model reduction in nonlinear dynamics using the Karhunen-Loève expansion and dual-weighted-residual methods. *Comput Mech* 31(1–2):179–191
90. Mortenson ME (1990) Computer graphics handbook. Industrial Press, New York
91. Murakami Y (2001) Stress intensity factors handbook. Elsevier, Amsterdam
92. Nagy DA (1979) Modal representation of geometrically nonlinear behaviour by the finite element method. *Comput Struct* 10:683–688
93. Newman AJ (1996) Model reduction via the Karhunen-Loeve expansion part i: an exposition. Technical report, Institute for System Research University of Maryland, pp 96–322
94. Newman JN (1977) Marine hydrodynamics. MIT Press, Cambridge
95. Nguyen NC (2005) Reduced-basis approximation and *a posteriori* error bounds for nonaffine and nonlinear partial differential equations: Application to inverse analysis. PhD thesis, Singapore-MIT Alliance, National University of Singapore
96. Nguyen NC, Patera AT (2007) Efficient and reliable parameter estimation in heat conduction using Bayesian inference and a reduced basis method (in preparation)
97. Nguyen NC, Veroy K, Patera AT (2005) Certified real-time solution of parametrized partial differential equations. In: Yip S (ed) Handbook of materials modeling. Springer, Berlin, pp 1523–1558
98. Noor AK (1981) Recent advances in reduction methods for nonlinear problems. *Comput Struct* 13:31–44
99. Noor AK (1982) On making large nonlinear problems small. *Comput Methods Appl Mech Eng* 34:955–985
100. Noor AK, Balch CD, Shibut MA (1984) Reduction methods for non-linear steady-state thermal analysis. *Int J Numer Methods Eng* 20:1323–1348
101. Noor AK, Peters JM (1980) Reduced basis technique for nonlinear analysis of structures. *AIAA J* 18(4):455–462
102. Noor AK, Peters JM (1981) Bifurcation and post-buckling analysis of laminated composite plates via reduced basis techniques. *Comput Methods Appl Mech Eng* 29:271–295
103. Noor AK, Peters JM (1981) Tracing post-limit-point paths with reduced basis technique. *Comput Methods Appl Mech Eng* 28:217–240
104. Noor AK, Peters JM (1983) Multiple-parameter reduced basis technique for bifurcation and post-buckling analysis of composite plates. *Int J Numer Methods Eng* 19:1783–1803
105. Noor AK, Peters JM (1983) Recent advances in reduction methods for instability analysis of structures. *Comput Struct* 16:67–80
106. Noor AK, Peters JM, Andersen CM (1984) Mixed models and reduction techniques for large-rotation nonlinear problems. *Comput Methods Appl Mech Eng* 44:67–89
107. Oliveira I, Patera AT (2007) Reduced-basis techniques for rapid reliable optimization of systems described by affinely parametrized coercive elliptic partial differential equations. *Optim Eng* 8:43–65
108. Paraschivoiu M, Peraire J, Maday Y, Patera AT (1998) Fast bounds for outputs of partial differential equations. In: Borggaard J, Burns J, Cliff E, Schreck S (eds) Computational methods for optimal design and control. Birkhäuser, Boston, pp 323–360
109. Parks DM (1974) A stiffness derivative finite element technique for determination of crack tip stress intensity factors. *Int J Fract* 10(4):487–502
110. Parlett BN (1998) The symmetric eigenvalue problem. SIAM, Philadelphia
111. Patera AT, Rønquist EM (2007) Reduced basis approximations and *a posteriori* error estimation for a Boltzmann model. *Comput Methods Appl Mech Eng* 196:2925–2942
112. Patera AT, Rozza G (2008) Reduced basis approximation and a posteriori error estimation for parametrized partial differential equations. Copyright MIT (2006–2007). MIT Pappalardo monographs in mechanical engineering (to appear)
113. Pau GSH (2007) Reduced-basis method for quantum models of periodic solids. PhD thesis, Massachusetts Institute of Technology
114. Peterson JS (1989) The reduced basis method for incompressible viscous flow calculations. *SIAM J Sci Stat Comput* 10(4):777–786
115. Phillips JR (2000) Projection frameworks for model reduction of weakly nonlinear systems. In: Proceeding of the 37th ACM/IEEE Design Automation Conference, pp 184–189
116. Phillips JR (2003) Projection-based approaches for model reduction of weakly nonlinear systems, time-varying systems. *IEEE Trans Comput Aided Des Integr Circuits Syst* 22:171–187
117. Pierce N, Giles MB (2000) Adjoint recovery of superconvergent functionals from PDE approximations. *SIAM Rev* 42(2):247–264
118. Pironneau O (2006) Calibration of barrier options. In: Fitzgibbon W, Hoppe R, Periaux J, Pironneau O, Vassilevski Y (eds) Advances in numerical mathematics. Moscow/Houston, Russian Academy of Sciences/University of Houston, pp 183–192
119. Porsching TA (1985) Estimation of the error in the reduced basis method solution of nonlinear equations. *Math Comput* 45(172):487–496

120. Porsching TA, Lee MYL (1987) The reduced-basis method for initial value problems. *SIAM J Numer Anal* 24:1277–1287
121. Prud'homme C, Rovas D, Veroy K, Maday Y, Patera A, Turinici G (2002) Reliable real-time solution of parametrized partial differential equations: Reduced-basis output bounds methods. *J Fluids Eng* 124(1):70–80
122. Prud'homme C, Rovas D, Veroy K, Patera AT (2002) A mathematical and computational framework for reliable real-time solution of parametrized partial differential equations. *Model Math Anal Numer* 36(5):747–771
123. Quarteroni A, Rozza G (2007) Numerical solution of parametrized Navier-Stokes equations by reduced basis method. *Numer Methods Partial Differ Equ* 23:923–948
124. Quarteroni A, Rozza G, Quaini A (2006) Reduced basis method for optimal control of advection-diffusion processes. In: Fitzgibbon W, Hoppe R, Periaux J, Pironneau O, Vassilevski Y (eds) *Advances in numerical mathematics*. Russian Academy of Sciences/University of Houston, Moscow/Houston, pp 193–216
125. Quarteroni A, Sacco R, Saleri F (2000) *Numerical mathematics. Texts in applied mathematics*, vol 37. Springer, New York
126. Quarteroni A, Valli A (1997) *Numerical approximation of partial differential equations*, 2nd edn. Springer, Berlin
127. Ravindran SS (2000) Reduced-order adaptive controllers for fluid flows using pod. *J Sci Comput* 15(4):457–478
128. Ravindran SS (2000) A reduced order approach to optimal control of fluids flow using proper orthogonal decomposition. *Int J Numer Methods Fluids* 34(5):425–448
129. Ravindran SS (2002) Adaptive reduced-order controllers for a thermal flow system using proper orthogonal decomposition. *SIAM J Sci Comput* 23(6):1924–1942
130. Rewienski M, White J (2003) A trajectory piecewise-linear approach to model order reduction and fast simulation of nonlinear circuits and micromachined devices. *IEEE Trans Comput-Aided Des Integr Circuits Syst* 22:155–170
131. Rheinboldt WC (1981) Numerical analysis of continuation methods for nonlinear structural problems. *Comput Struct* 13(1–3):103–113
132. Rheinboldt WC (1993) On the theory and error estimation of the reduced basis method for multi-parameter problems. *Nonlinear Anal Theory, Methods Appl* 21(11):849–858
133. Rovas D, Machiels L, Maday Y (2005) Reduced basis output bounds methods for parabolic problems. *IMA J Appl Math*
134. Rovas DV (2002) Reduced-basis output bound methods for parametrized partial differential equations. PhD thesis, Massachusetts Institute of Technology, Cambridge, MA
135. Rozza G (2005) Real-time reduced basis techniques for arterial bypass geometries. In: Bathe K (ed) *Proceedings of the third MIT conference on computational fluid and solid mechanics*, June 14–17, 2005. Computational fluid and solid mechanics. Elsevier, Amsterdam, pp 1283–1287
136. Rozza G (2005) Reduced-basis methods for elliptic equations in sub-domains with *a posteriori* error bounds and adaptivity. *Appl Numer Math* 55(4):403–424
137. Rozza G (2005) Shape design by optimal flow control and reduced basis techniques: Applications to bypass configurations in haemodynamics. PhD thesis, EPFL, Ecole Polytechnique Federale de Lausanne
138. Rozza G (2008) Reduced basis method for Stokes equations in domains with non-affine parametric dependence. *Comput Vis Sci* 11(4). doi:10.1007/s00791-006-0044-7
139. Rozza G, Veroy K (2007) On the stability of reduced basis method for Stokes equations in parametrized domains. *Comput Methods Appl Mech Eng* 196:1244–1260
140. Schiesser WE, Silebi CA (1997) *Computational transport phenomena: numerical methods for the solution of transport problems*. Cambridge University Press, Cambridge
141. Sen S (2007) Reduced-basis approximation and *a posteriori* error estimation for non-coercive elliptic problems: Application to acoustics. PhD thesis, Massachusetts Institute of Technology
142. Sen S, Veroy K, Huynh DBP, Deparis S, Nguyen NC, Patera AT (2006) “Natural norm” *a posteriori* error estimators for reduced basis approximations. *J Comput Phys* 217:37–62
143. Shi G, Shi CJR (2004) Parametric model order reduction for interconnect analysis. In: *Proceedings of the 2004 conference on Asia South Pacific design automation: electronic design and solution fair*. IEEE Press, New York, pp 774–779
144. Sirisup S, Xiu D, Karniadakis G (2005) Equation-free/Galerkin-free POD-assisted computation of incompressible flows. *J Comput Phys* 207:617–642
145. Sirovich L (1987) Turbulence and the dynamics of coherent structures, part 1: Coherent structures. *Q Appl Math* 45(3):561–571
146. Strang G (2003) *Introduction to linear algebra*. Wellesley-Cambridge, Wellesley
147. Strang G, Fix GJ (1973) *An analysis of the finite element method*. Prentice-Hall, New York
148. Tonn T, Urban K (2006) A reduced-basis method for solving parameter-dependent convection-diffusion problems around rigid bodies. In: Wesseling P, Oñate E, Periaux J (eds) *ECCOMAS CFD 2006 proceedings*. TU Delft, Delft
149. Trefethen L, III DB (1997) *Numerical linear algebra*. SIAM, Philadelphia
150. Veroy K (2003) Reduced-basis methods applied to problems in elasticity: Analysis and applications. PhD thesis, Massachusetts Institute of Technology
151. Veroy K, Patera AT (2005) Certified real-time solution of the parametrized steady incompressible Navier-Stokes equations; Rigorous reduced-basis *a posteriori* error bounds. *Int J Numer Methods Fluids* 47:773–788
152. Veroy K, Prud'homme C, Patera AT (2003) Reduced-basis approximation of the viscous Burgers equation: Rigorous *a posteriori* error bounds. *C R Acad Sci Paris, Sér I* 337(9):619–624
153. Veroy K, Prud'homme C, Rovas DV, Patera AT (2003) *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*. Paper 2003-3847
154. Wang J, Zabarav N (2005) Using Bayesian statistics in the estimation of heat source in radiation. *Int J Heat Mass Transfer* 48:15–29
155. Weile DS, Michielssen E (2001) Analysis of frequency selective surfaces using two-parameter generalized rational Krylov model-order reduction. *IEEE Trans Antennas Propag* 49(11):1539–1549
156. Weile DS, Michielssen E, Gallivan K (2001) Reduced-order modeling of multiscreen frequency-selective surfaces using Krylov-based rational interpolation. *IEEE Trans Antennas Propag* 49(5):801–813
157. Willcox K, Peraire J (2002) Balanced model reduction via the proper orthogonal decomposition. *AIAA J* 40(11):2323–2330
158. Zienkiewicz O, Taylor R (2000) *Finite element method. The basis*, vol 1. Butterworth-Heinemann, London