

Localized Reduced Basis Additive Schwarz Methods

Martin J. Gander and Stephan Rave

1 Introduction

Reduced basis (RB) methods [9, 6] are a family of model order reduction schemes for parameterized PDEs, which can speed up the repeated solution of such equations by orders of magnitude. In the so-called *offline* phase, RB methods construct a problem-adapted low-dimensional approximation space by computing solutions of the PDE for selected *snapshot* parameters using a given high-fidelity discretization of the PDE. In the following *online* phase, the PDE is solved for arbitrary new parameters by computing the (Petrov-)Galerkin projection of its solution onto the precomputed reduced approximation space. While RB methods have been proven successful in various applications, for very large problems the computation of the solution snapshots in the offline phase may still be prohibitively expensive. To mitigate this issue, localized RB methods [7, 4] have been developed which construct the global approximation space from spatially localized less expensive problems. These local problems largely fall into two classes:

Training procedures construct local approximation spaces without knowledge of the global problem by, e.g., solving the equation on an enlarged subdomain with arbitrary boundary values and then restricting the solution to the domain of interest, or by solving related eigenvalue problems. As such, these training approaches have a strong connection with numerical multiscale methods and the construction of spectral coarse spaces in domain decomposition methods.

In this contribution, however, we will focus on the construction of local RB spaces via *online enrichment*, where these spaces are iteratively built by solving localized

Martin J. Gander

Section de Mathématiques, Université de Genève, 2-4 rue du Lièvre, CP 64, CH-1211, Genève, Suisse. e-mail: Martin.Gander@unige.ch

Stephan Rave

Mathematics Münster, University of Münster, Einsteinstrasse 62, 48149 Münster, Germany e-mail: Stephan.Rave@uni-muenster.de

corrector problems for the residual of the current reduced solution. In particular we mention the use of online enrichment in context of the LRBMS [8], GMsFEM [5] and ArbiLoMod [3] methods. These enrichment schemes share strong similarities with Schwarz methods, and it is the main goal of this contribution to shed some light on the connections between these methods. We will do so by introducing a simple localized RB additive Schwarz (LRBAS) method which is phrased in the language of the abstract Schwarz framework but incorporates the central ingredients of online adaptive localized RB methods. In particular, we hope that LRBAS will help the analysis of localized RB methods from the perspective of Schwarz methods. Following [3], we will consider arbitrary but localized changes of the problem instead of parametric variations. In Section 2.1 we will see that LRBAS can indeed be interpreted as a locally adaptive version of a multi-preconditioned CG method.

Compared to Schwarz methods, a distinctive feature of LRBAS is that updates are only computed in high-residual regions, which can lead to a significant reduction of the number of local updates and a concentration of the updates to a few regions affected by the localized changes (cf. Section 3). This property might be exploited for the reduction of the overall power consumption and to balance the computational load among a smaller amount of compute nodes, in particular in cloud environments, where additional computational resources can be easily allocated and deallocated again.

2 A Localized Reduced Basis Additive Schwarz Method

Our goal is to efficiently solve a sequence, indexed by k , of linear systems

$$A^{(k)}x^{(k)} = f \quad (1)$$

with $A^{(k)} \in \mathbb{R}^{n \times n}$ symmetric, positive definite and $x^{(k)}, f \in \mathbb{R}^n$, up to some fixed error tolerance ε . To this end, let $n \times n_i$ matrices R_i^T of rank n_i be given for $1 \leq i \leq I$ and $n \times n_0^{(k)}$ matrices $R_0^{(k)T}$ of rank $n_0^{(k)}$. Typically, R_1, \dots, R_I will be the restriction matrices corresponding to a finite element basis associated with an overlapping domain decomposition Ω_i of the computational domain Ω , and the columns of $R_0^{(k)T}$ contain a basis of a suitable coarse space for $A^{(k)}$. In particular we assume that each R_i is non-orthogonal to only a few neighboring spaces, i.e., there are a small constant C and index sets $O_i \subset \{1, \dots, I\}$ with $\#O_i \leq C \cdot I$ such that

$$R_j \cdot R_i^T = 0_{n_j \times n_i} \quad \text{whenever } j \notin O_i. \quad (2)$$

As usual, we define the local matrices

$$A_0^{(k)} := R_0^{(k)} A^{(k)} R_0^{(k)T} \quad \text{and} \quad A_i^{(k)} := R_i A^{(k)} R_i^T.$$

We are interested in the case where $A^{(k+1)}$ is obtained from $A^{(k)}$ by an arbitrary but local modification in the sense that

$$A_i^{(k+1)} = A_i^{(k)} \quad \text{for } i \notin C^{(k+1)}, \quad (3)$$

where the sets $C^{(k)}$ contain the indices of the spaces affected by the change, generally assuming that $\#C^{(k)} \ll I$.

Over the course of the computation of the solutions $x^{(k)}$ we will build local low-dimensional reduced bases $\tilde{R}_i^{(k,l)T} \in \mathbb{R}^{n_i \times N_i^{(k,l)}}$ for $i \geq 1$ such that there are local coefficients $\tilde{x}_i^{(k,l)} \in \mathbb{R}^{N_i^{(k,l)}}$ and $\tilde{x}_0^{(k,l)} \in \mathbb{R}^{n_0^{(k)}}$ such that

$$\tilde{x}^{(k,l)} := R_0^{(k)T} \tilde{x}_0^{(k,l)} + \sum_{i=1}^I R_i^T \tilde{R}_i^{(k,l)T} \tilde{x}_i^{(k,l)} \quad (4)$$

is a good approximation of x^k for sufficiently large l . We obtain such an approximation via Galerkin projection onto the global reduced basis space spanned by the images of $R_0^{(k)T}$ and all $R_i^T \tilde{R}_i^{(k,l)T}$, i.e., $\tilde{x}^{(k,l)}$ is determined by the $(n_0^{(k)} + \sum_{i=1}^I N_i^{(k,l)})$ -dimensional linear system

$$\begin{aligned} R_0^{(k)} A^{(k)} \tilde{x}^{(k,l)} &= R_0^{(k)} f, \\ \tilde{R}_i^{(k,l)} R_i A^{(k)} \tilde{x}^{(k,l)} &= \tilde{R}_i^{(k,l)} R_i f, \quad 1 \leq i \leq I. \end{aligned} \quad (5)$$

Thanks to the locality (2) of the space decomposition, the matrix of the system (5) has a block structure allowing us to efficiently assemble and solve it.

To build the local reduced bases $\tilde{R}_i^{(k,l)T}$ we use an iterative enrichment procedure where the basis is extended with local Schwarz corrections $y_i^{(k,l)} \in \mathbb{R}^{n_i}$ for the current residual,

$$A_i^{(k)} y_i^{(k,l)} = r_i^{(k,l)} := R_i(f - A^{(k)} \tilde{x}^{(k)}). \quad (6)$$

In view of (3), the corrections are only computed in subdomains i with large residual norm $\|r_i^{(k,l)}\|$. In particular, for finite-element discretizations of elliptic PDEs without high-conductivity channels, we expect that with increasing k the number of enriched bases will be of the same order as the cardinality of C^{k+1} . The exact definition of the enrichment scheme is given in Algorithm 1. There are various possibilities to choose the criterion for the localized enrichment in line 9 of Algorithm 1. In this work we simply select those reduced spaces for enrichment for which the quotient between the norm of the local residual and the norm of the global residual is larger than a fixed constant that scales with the number of the subdomains.

Note that an important property of localized enrichment is that after an enrichment step only those blocks (i, j) of the matrix corresponding to (5) have to be updated for which either $\tilde{R}_i^{(k,l)}$ or $\tilde{R}_j^{(k,l)}$ have been enriched. Using reduced basis techniques [2] it is further possible to evaluate the residual norms $\|R_i r^{(k,l)}\|$ and $\|r^{(k,l)}\|$ using only reduced quantities, which again only have to be updated for local bases $\tilde{R}_i^{(k,l)T}$ affected by the enrichment. Thus, in a distributed computing environment only the main compute node solving (5) and those nodes associated with the enriched bases have to perform any operations, while the other compute node lay at rest.

Algorithm 1 Localized Reduced Basis Additive Schwarz method (LRBAS)

```

1: procedure LRBAS( $A^{(k)}, f, R_0^T, \tilde{R}_i^T, \varepsilon, \varepsilon_{\text{loc}}$ )
2:    $\tilde{R}_i^{(1,1)T} \leftarrow 0_{n_i \times 0}, \quad 1 \leq i \leq I$  ▷ initialize local bases
3:   for  $k \leftarrow 1, \dots, \infty$  do
4:      $\tilde{x}^{(k,1)}, \tilde{x}_i^{(k,1)} \leftarrow$  solutions of (4), (5) ▷ initial solution
5:      $r^{(k,1)} \leftarrow f - A^{(k)} \tilde{x}^{(k,1)}$  ▷ initial residual
6:      $l \leftarrow 1$ 
7:     while  $\|r^{(k,l)}\| / \|f\| > \varepsilon$  do ▷ loop until converged
8:       for  $i \leftarrow 1, \dots, I$  do ▷ enrichment procedure
9:         if  $\|\tilde{R}_i r^{(k,l)}\|^2 > \varepsilon_{\text{loc}} \cdot I^{-1} \cdot \|r^{(k,l)}\|^2$  then
10:           $y_i^{(k,l)} \leftarrow$  solution of (6)
11:           $\tilde{R}_i^{(k,l+1)T} \leftarrow [\tilde{R}_i^{(k,l)T} \ y_i^{(k,l)}]$ 
12:        else
13:           $\tilde{R}_i^{(k,l+1)} \leftarrow \tilde{R}_i^{(k,l)}$ 
14:        end if
15:      end for
16:       $\tilde{x}^{(k,l+1)}, \tilde{x}_i^{(k,l+1)} \leftarrow$  solutions of (4), (5) ▷ update solution
17:       $r^{(k,l+1)} \leftarrow f - A^{(k)} \tilde{x}^{(k,l+1)}$  ▷ update residual
18:       $l \leftarrow l + 1$ 
19:    end while
20:    for  $i \leftarrow 1, \dots, I$  do ▷ update bases for next problem
21:      if  $\tilde{R}_i^{(k,l)T} \neq \tilde{R}_i^{(k,1)T}$  then ▷ basis enriched at least once?
22:         $\tilde{R}_i^{(k+1,1)T} \leftarrow [\tilde{R}_i^{(k,1)T} \ \tilde{R}_i^{(k,l)T} \ \tilde{x}_i^{(k,l)}]$  ▷ only keep local solution in basis
23:      else
24:         $\tilde{R}_i^{(k+1,1)T} \leftarrow \tilde{R}_i^{(k,1)T}$ 
25:      end if
26:    end for
27:  end for
28: end procedure

```

We remark that several extensions to the LRBAS method are possible. In particular, we assumed for simplicity that all matrices $A^{(k)}$ are of the same dimension. This, for instance, is the case when coefficient functions of the PDE underlying (1) are modified, but the computational mesh remains unchanged. However, also local geometry changes that lead to remeshing can be handled by resetting all local bases that are supported on the changed geometry. In this context we note that, as another simplification, in the definition of LRBAS we have chosen to keep all basis vectors when transitioning from $A^{(k)}$ to $A^{(k+1)}$, including bases $\tilde{R}_i^{(k,l)T}$ affected by the change, even though these retained bases will generally not contribute to the convergence of the scheme. Finally, in many applications, a local or global parametric variation of $A^{(k)}$, e.g. the change of some material parameters, in addition to the considered non-parametric modifications may be of interest. In such cases, parametric model order reduction techniques such as greedy basis generation algorithms or offline/online decomposition of the reduced order system (5) can be incorporated into the scheme. In particular we refer to [3] where both additional parameterization of $A^{(k)}$ as well as the reinitialization of the local bases after non-parametric changes from $A^{(k)}$ to $A^{(k+1)}$ are discussed.

2.1 LRBAS as an additive-Schwarz multi-preconditioned CG method

Consider the solution of the systems (1) with the preconditioned conjugate gradient (PCG) algorithm, where we choose as preconditioner the additive Schwarz operator $(M^{(k)})^{-1} := R_0^{(k)T} (A_0^{(k)})^{-1} R_0^{(k)} + \sum_{i=1}^I R_i^T (A_i^{(k)})^{-1} R_i$. Let $x_{\text{pcg}}^{(k,l)}$ denote the l -th iterate of the PCG algorithm, starting with $x^{(k,0)} = 0$ as the initial guess. Then it is well known that $x_{\text{pcg}}^{(k,l)}$ lies in the search space $\mathcal{S}_{\text{pcg}}^{(k,l)}$ given by the Krylov space $\mathcal{K}^l \left((M^{(k)})^{-1} A^{(k)}, (M^{(k)})^{-1} f \right)$ and that the error $x^{(k)} - x_{\text{pcg}}^{(k,l)}$ is $A^{(k)}$ -orthogonal to this space. Denoting by $r_{\text{pcg}}^{(k,l)} := f - A^{(k)} x_{\text{pcg}}^{(k,l)}$ the l -th residual, one readily checks that $\mathcal{S}_{\text{pcg}}^{(k,l)}$ is equivalently given by

$$\mathcal{S}_{\text{pcg}}^{(k,l)} := \text{span} \left\{ (M^{(k)})^{-1} r_{\text{pcg}}^{(k,0)}, \dots, (M^{(k)})^{-1} r_{\text{pcg}}^{(k,l-1)} \right\},$$

i.e., in each iteration the search space is extended by the vector obtained from the application of the preconditioner to the current residual. The idea of multi-preconditioning [1] is to enlarge this search space by including each local preconditioner $(A_i^{(k)})^{-1}$ application into the search space individually, leading to

$$\begin{aligned} \mathcal{S}_{\text{mpcg}}^{(k,l)} := \text{span} \left(\left\{ R_0^{(k)T} (A_0^{(k)})^{-1} R_0^{(k)} r_{\text{mpcg}}^{(k,t)} \mid 0 \leq t \leq l-1 \right\} \right. \\ \left. \cup \left\{ R_i^T (A_i^{(k)})^{-1} R_i r_{\text{mpcg}}^{(k,t)} \mid 1 \leq i \leq I, 0 \leq t \leq l-1 \right\} \right), \end{aligned}$$

with $r_{\text{mpcg}}^{(k,l)}$ denoting the multi-preconditioned CG residuals. Conversely, we easily see from (5) and (6) that for $\varepsilon_{\text{loc}} = 0$ the LRBAS iterates $\tilde{x}^{(k,l)}$ lie within the search space

$$\begin{aligned} \mathcal{S}_{\text{lrbas},0}^{(k,l)} := \text{Im} \left(\begin{bmatrix} R_0^{(k)T} & R_1^T \tilde{R}_1^{(k,1)} & \dots & R_I^T \tilde{R}_I^{(k,1)} \end{bmatrix} \right) \\ + \text{span} \left\{ R_i^T (A_i^{(k)})^{-1} R_i r^{(k,t)} \mid 1 \leq i \leq I, 1 \leq t \leq l-1 \right\}, \end{aligned}$$

and that the error $x^{(k)} - \tilde{x}^{(k,l)}$ is $A^{(k)}$ -orthogonal to this space. Hence, LRBAS with $\varepsilon_{\text{loc}} = 0$ can be seen as a projected multi-preconditioned CG method for solving (5), where the projection space is given by the span of the coarse space and the initial local reduced bases and where the new solution iterate $\tilde{x}^{(k,l)}$ is obtained by direct solution of the reduced system (5) instead of an incremental update in order to preserve the locality of the reduced bases.

For $\varepsilon_{\text{loc}} > 0$ we arrive at an adaptive version of multi-preconditioning similar to [10]. However, in contrast to [10] where either all local search directions or their global sum are added to the search space, LRBAS is locally adaptive in the sense that only those local search directions are computed and included where a large local residual has to be corrected.

3 Numerical Experiment

We consider the test case from [3] and solve a sequence of five elliptic problems

$$\begin{aligned}
 \nabla \cdot (-\sigma^{(k)}(x, y)\nabla u^{(k)}(x, y)) &= 0, & x, y \in (0, 1), \\
 u^{(k)}(0, y) &= 1, & y \in (0, 1), \\
 u^{(k)}(1, y) &= -1, & y \in (0, 1), \\
 -\sigma^{(k)}(x, y)\nabla u^{(k)}(x, y) \cdot \mathbf{n}(x, y) &= 0, & x \in (0, 1), y \in \{0, 1\},
 \end{aligned}
 \tag{7}$$

where the coefficient $\sigma^{(k)}(x)$ is given as in Fig. 1. The problem is discretized using bilinear finite elements over a uniform 200×200 mesh. The resulting solutions are visualized in Fig. 2. We decompose the computational domain uniformly into 10×10 subdomains with an overlap of 4 mesh elements. For $R_0^{(k)}$ we choose GenEO [11] basis functions with eigenvalues below 0.5, yielding between two and five functions per subdomain. When connecting or disconnecting the high-conductivity channels, we expect enrichment to be required along the subdomains adjacent to the channels, whereas the other subdomains should be largely unaffected by the local change.

In Table 1 we compare the total number of iterations for all five problems and the total number of Schwarz corrections (6) required to reach a relative error tolerance of $\varepsilon = 10^{-6}$ for the following solution strategies: 1. the additive Schwarz preconditioned CG method with zero initial guess or with a localized RB solution as initial guess, where the localized basis is obtained from the linear span of previous solutions $x^{(k)}$ decomposed using the GenEO partition of unity; 2. LRBAS with and without local adaptivity ($\varepsilon_{\text{loc}} = 0.25$ or 0); 3. a version of LRBAS where the entire bases $\tilde{R}_i^{(k,l)T}$ are preserved when transitioning to $k + 1$ instead of only the final solution $\tilde{R}_i^{(k,l)T} \tilde{x}_i^{(k,l)}$. As we see, LRBAS with locally adaptive enrichment significantly outperforms the PCG method with or without initial guess, both regarding the number of required iterations as well as the number of Schwarz corrections. Compared to non-adaptive




Fig. 1: Definition of the coefficient functions $\sigma^{(k)}$ for the numerical test case (7); left: function $\sigma^{(0)}$, taking the values $10^5 + 1$ inside the high-conductivity regions and 1 elsewhere; right: $\sigma^{(k)}$ is obtained from $\sigma^{(0)}$ by connecting the three channels to the boundary regions at the marked locations.




Fig. 2: Solutions of the test problem (7) for $k = 1, 2, 3$ (top row) and $k = 4, 5$ (bottom row).

	iterations	local enrichments (6)
PCG	107	10700
PCG + LRB solution as initial value	63	6300
LRBAS ($\varepsilon_{\text{loc}} = 0$)	33	3300
LRBAS ($\varepsilon_{\text{loc}} = 0.25$)	39	1386
LRBAS ($\varepsilon_{\text{loc}} = 0, \tilde{R}_i^{(k+1,1)} := \tilde{R}_i^{(k,l)}$)	28	2800
LRBAS ($\varepsilon_{\text{loc}} = 0.25, \tilde{R}_i^{(k+1,1)} := \tilde{R}_i^{(k,l)}$)	34	1335

Table 1: Total number of iterations and local Schwarz corrections (6) required to reach a relative error tolerance $\varepsilon = 10^{-6}$ for the test problem (7).

multi-preconditioning, i.e. LRBAS with $\varepsilon_{\text{loc}} = 0$, the number of local corrections is more than halved at the expense of a slightly increased number of iterations. Keeping all of $\tilde{R}_i^{(k,l)}$ improves the convergence of the method only slightly. Finally, in Fig. 3 we depict the number of required Schwarz corrections per subdomain for each k . We observe a good localization of the computational work among the subdomains most affected by the local changes.

Acknowledgements Funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under Germany’s Excellence Strategy EXC 2044 –390685587, Mathematics Münster: Dynamics–Geometry–Structure.




Fig. 3: Number of local Schwarz corrections (6) required by the LRBAS method with $\varepsilon_{\text{loc}} = 0.25$ to solve the five test problems (7) up to a relative error tolerance of $\varepsilon = 10^{-6}$.

References

1. R. Bridson and C. Greif. A multipreconditioned conjugate gradient algorithm. *SIAM Journal on Matrix Analysis and Applications*, 27(4):1056–1068, jan 2006.
2. A. Buhr, C. Engwer, M. Ohlberger, and S. Rave. A Numerically Stable A Posteriori Error Estimator for Reduced Basis Approximations of Elliptic Equations. In E. Oñate, X. Oliver, and A. Huerta, editors, *11th. World Congress on Computational Mechanics*, pages 4094–4102. International Center for Numerical Methods in Engineering, Barcelona, 2014.
3. A. Buhr, C. Engwer, M. Ohlberger, and S. Rave. ArbiLoMod, a simulation technique designed for arbitrary local modifications. *SIAM J. Sci. Comput.*, 39(4):A1435–A1465, 2017.
4. A. Buhr, L. Iapichino, M. Ohlberger, S. Rave, F. Schindler, and K. Smetana. Localized model reduction for parameterized problems. In P. Benner, S. Grivet-Talocia, A. Quarteroni, G. Rozza, W. Schilders, and L. M. Silveira, editors, *Model Order Reduction (Volume 2)*. De Gruyter, Berlin, Boston, 2021.
5. E. T. Chung, Y. Efendiev, and W. T. Leung. Residual-driven online generalized multiscale finite element methods. *Journal of Computational Physics*, 302:176–190, 2015.
6. J. S. Hesthaven, G. Rozza, and B. Stamm. *Certified Reduced Basis Methods for Parametrized Partial Differential Equations*. SpringerBriefs in Mathematics. Springer International Publishing, 1 edition, 2016.
7. Y. Maday and E. M. Rønquist. A reduced-basis element method. *Comptes Rendus Mathématique*, 335(2):195–200, 2002.
8. M. Ohlberger and F. Schindler. Error control for the localized reduced basis multiscale method with adaptive on-line enrichment. *SIAM J. Sci. Comput.*, 37(6):A2865–A2895, 2015.
9. A. Quarteroni, A. Manzoni, and F. Negri. *Reduced Basis Methods for Partial Differential Equations*, volume 92 of *La Matematica per il 3+2*. Springer International Publishing, 2016.
10. N. Spillane. An adaptive MultiPreconditioned conjugate gradient algorithm. *SIAM Journal on Scientific Computing*, 38(3):A1896–A1918, jan 2016.
11. N. Spillane, V. Dolean, P. Hauret, F. Nataf, C. Pechstein, and R. Scheichl. Abstract robust coarse spaces for systems of PDEs via generalized eigenproblems in the overlaps. *Numerische Mathematik*, 126(4):741–770, aug 2013.