

Efficient implementation of the localized orthogonal decomposition method

Christian Engwer, Patrick Henning, Axel Målqvist, Daniel Peterseim

Angaben zur Veröffentlichung / Publication details:

Engwer, Christian, Patrick Henning, Axel Målqvist, and Daniel Peterseim. 2019. "Efficient implementation of the localized orthogonal decomposition method." *Computer Methods in Applied Mechanics and Engineering* 350: 123–53. <https://doi.org/10.1016/j.cma.2019.02.040>.

Efficient implementation of the Localized Orthogonal Decomposition method

Christian Engwer¹, Patrick Henning², Axel Målqvist³, Daniel Peterseim⁴

February 21, 2019

Abstract

In this paper we present algorithms for an efficient implementation of the Localized Orthogonal Decomposition method (LOD). The LOD is a multiscale method for the numerical simulation of partial differential equations with a continuum of inseparable scales. We show how the method can be implemented in a fairly standard Finite Element framework and discuss its realization for different types of problems, such as linear elliptic problems with rough coefficients and linear eigenvalue problems.

1 Introduction

By now, the Localized Orthogonal Decomposition (LOD) of a subspace $V \subset H^1(\Omega)$ into a coarse space and a detail space is a well established method for the numerical homogenization of partial differential equations. So far it has been successfully applied to linear elliptic multiscale problems in the context of continuous finite elements [1, 2, 3], discontinuous finite elements [4, 5, 6], mixed finite elements [7, 8], partition of unity methods [9] and reduced basis simulations [10]. The range of applications covers linear and quadratic eigenvalue problems [11, 12], problems in perforated domains [13] and high-contrast media [14, 15], stochastic homogenization [16, 17], semilinear elliptic problems [18], the wave equation [19, 20, 21], parabolic and coupled problems [22, 23, 24], the Buckley-Leverett equation [25], fractional diffusion problems [26], Helmholtz problems [27, 28, 29, 30] and the simulation of Bose-Einstein condensates [31]. An introductory general overview is given in [32].

Initially inspired by the Variational Multiscale Method [33, 34, 35, 36, 37, 7], the LOD in its present form was first proposed and rigorously justified in [1]. Further basic modifications of the method were suggested in [2]. The LOD is constructed to handle discrete problems that involve a high-dimensional solution space (also referred to as the 'fine space'). This typically takes place in two steps. In the first step the full fine space is decomposed into a low-dimensional space with good approximation properties and a high-dimensional remainder space. In the second step, this decomposition is localized in the sense that the low-dimensional space is approximated by constructing suitable locally supported basis functions that are the solutions of small patch problems. Due to their size, the patch problems are cheap to solve. Furthermore, they can be solved independently from each other and are hence perfect for parallelization. This strategy is particularly useful to reduce/distribute the computational cost of solving large systems of equations (arising from finite element discretizations). The method can be linked to conceptually very different techniques of mathematical modeling and scientific computing, e.g., it recovers the mathematical theory of homogenization [38] in periodic diffusion problems and even bridges to the theory of iterative solvers and subspace decomposition methods [39, 40]. Moreover, the method may be interpreted as a stabilization technique that coincides with the streamline upwind Petrov-Galerkin method SUPG [35]. In the last 5 years it has inspired numerous new developments in

¹Institute for Computational and Applied Mathematics, University of Münster, Germany

²Department of Mathematics, KTH Royal Institute of Technology, Stockholm, Sweden

³Department of Mathematics, Chalmers University of Technology and University of Gothenburg, Sweden.

⁴Institute for Mathematics, University of Augsburg, Germany

the field of multiscale partial differential equations including rough polyharmonic splines [41], iterative numerical homogenization [42], and gamblets [43]. While previous works focused on the numerical analysis of the method, this paper aims at the detailed explanation of how the method can be algorithmically realized. We give detailed explanations on how the method works on an algebraic level. The results may as well be useful for implementing related multiscale methods.

2 Preliminaries

In this section we recall the Localized Orthogonal Decomposition (LOD) for finite element spaces. The decomposition is always with respect to a linear elliptic part of the differential operator.

2.1 Computational domain and boundary

For the rest of the paper, we consider a bounded polygonal domain $\Omega \subset \mathbb{R}^d$. The boundary $\partial\Omega$ is divided into two parts Γ_D and Γ_N . On Γ_D we prescribe a Dirichlet boundary condition and on Γ_N we prescribe a Neumann boundary condition. We have $\Gamma_D \cup \Gamma_N = \partial\Omega$ and we assume $\Gamma_D \neq \emptyset$. With that, we define the space $H_{\Gamma_D}^1(\Omega) := \{v \in H^1(\Omega) \mid v|_{\Gamma_D} = 0\}$, where $v|_{\Gamma_D} = 0$ is understood in the sense of traces.

2.2 Elliptic differential operator

Subsequently we consider the following differential operator. Let $\kappa \in L^\infty(\Omega, \mathbb{R}^{d \times d})$ denote a matrix-valued, symmetric, possibly highly varying and heterogeneous coefficient with uniform spectral bounds $\gamma_{\min} > 0$ and $\gamma_{\max} \geq \gamma_{\min}$,

$$\sigma(\kappa(x)) \subset [\gamma_{\min}, \gamma_{\max}] \quad \text{for almost all } x \in \Omega.$$

This coefficient defines a scalar product $\mathcal{A}(\cdot, \cdot)$ on $H_{\Gamma_D}^1(\Omega)$ that is given by

$$\mathcal{A}(v, w) := \int_{\Omega} \kappa \nabla v \cdot \nabla w \quad \text{for } v, w \in H_{\Gamma_D}^1(\Omega).$$

2.3 Meshes and spaces

We wish to discretize a problem that is associated with $\mathcal{A}(\cdot, \cdot)$. Then the discretization is constrained by the diffusion coefficient κ , in the sense that variations of κ must be resolved by the computational mesh. We call such a discretization a *fine scale discretization*. In addition to this, we have a second discretization on a *coarse scale*. The coarse mesh is arbitrary and no more related to $\mathcal{A}(\cdot, \cdot)$. It contains elements of maximum diameter $H > 0$. The fine mesh consists of elements of maximum diameter $h < H$. Let $\mathcal{T}_H, \mathcal{T}_h$ denote the corresponding simplicial or quadrilateral subdivisions of Ω into (closed) conforming shape regular simplicial elements or conforming shape regular quadrilateral elements, i.e., $\bar{\Omega} = \bigcup_{K_h \in \mathcal{T}_h} K_h = \bigcup_{K \in \mathcal{T}_H} K$. We assume that \mathcal{T}_h is a regular, possibly non-uniform, mesh refinement of \mathcal{T}_H . Furthermore we also assume that \mathcal{T}_H and \mathcal{T}_h are shape-regular in the sense that there exists a positive constant c_0 such that

$$\max \left\{ \max_{K_h \in \mathcal{T}_h} \frac{\text{diam}(K_h)^d}{|K_h|}, \max_{K \in \mathcal{T}_H} \frac{\text{diam}(K)^d}{|K|} \right\} \leq c_0$$

and regular in the sense that any two elements are either disjoint, share exactly one face, share exactly one edge, or share exactly one vertex. For $\mathcal{T} = \mathcal{T}_H, \mathcal{T}_h$, let

$$P_1(\mathcal{T}) = \{v \in C^0(\Omega) \mid \forall K \in \mathcal{T}, v|_K \text{ is a polynomial of total degree } \leq 1\} \quad \text{and} \quad (1a)$$

$$Q_1(\mathcal{T}) = \{v \in C^0(\Omega) \mid \forall K \in \mathcal{T}, v|_K \text{ is a polynomial of partial degree } \leq 1\} \quad (1b)$$

denote the typical $p1$ degree and bi- $p1$ degree Finite Element Spaces for triangular and quadrilateral partitions respectively. We set $V_h := P_1(\mathcal{T}_h)$ if \mathcal{T}_h is a triangulation and $V_h := Q_1(\mathcal{T}_h)$ if \mathcal{T}_h is a quadrilateralization. The 'coarse space' (i.e. low dimensional space) $V_H \subset V_h$ is defined analogously. Furthermore, we set $V_{h,\Gamma_D} := V_h \cap H_{\Gamma_D}^1(\Omega)$ and $V_{H,\Gamma_D} := V_H \cap H_{\Gamma_D}^1(\Omega)$. For simplicity we assume that V_{H,Γ_D} is aligned with Γ_D (in $2d$ this means that $\overline{\Gamma_D} \cap \overline{\Gamma_N}$ is a subset of coarse grid nodes). The full sets of fine nodes in V_h , respectively coarse nodes in V_H , are given by

$$\mathcal{N}_H = \{Z_i | 0 \leq i \leq N_H - 1\} \quad \text{and} \quad \mathcal{N}_h = \{z_j | 0 \leq j \leq N_h - 1\},$$

where N_H and N_h denote the number of vertices in the fine and the coarse mesh. Accordingly we introduce $N_{\mathcal{T}_H} = |\mathcal{T}_H|$ and $N_{\mathcal{T}_h} = |\mathcal{T}_h|$ as the number of cells in the mesh. The coarse nodal basis function that is associated with a node $Z_i \in \mathcal{N}_H$ shall be denoted by $\Phi_i \in V_H$ and the fine nodal basis function associated with $z_j \in \mathcal{N}_h$ shall be denoted by $\phi_j \in V_h$.

2.4 Two-scale decompositions

In order to introduce an \mathcal{A} -orthogonal decomposition of the space V_{h,Γ_D} , we require a projection $I_H : V_{h,\Gamma_D} \rightarrow V_{H,\Gamma_D}$ (i.e. $(I_H \circ I_H) = I_H$) that maps a fine-scale function into the coarse fine element space V_{H,Γ_D} . The chosen projection will help us to characterize the "details" in V_{h,Γ_D} and it is desirable that I_H is L^2 - and H^1 -stable. Before we introduce a decomposition based on I_H , we state examples of possible choices for I_H .

Remark 2.1. Examples for projections I_H that fulfill the desired stability properties on quasi-uniform meshes.

- The operator $I_H : V_{h,\Gamma_D} \rightarrow V_{H,\Gamma_D}$ can be chosen as the global L^2 -projection onto finite elements given by

$$(I_H(v_h), \Phi_H)_{L^2(\Omega)} = (v_h, \Phi_H)_{L^2(\Omega)} \quad \text{for all } \Phi_H \in V_{H,\Gamma_D}.$$

- The operator $I_H : V_{h,\Gamma_D} \rightarrow V_{H,\Gamma_D}$ can be also constructed from a local L^2 -projection. Given a coarse-node Z_i and corresponding nodal patch $\omega_i := \text{supp}(\Phi_i)$, we let P_{H,ω_i} denote the L^2 -projection onto the standard P_1 finite element space on ω_i . Exploiting this, we define I_H for $v_h \in V_{h,\Gamma_D}$ by $I_H(v_h) := \sum_{i=0}^{N_H-1} \alpha_i(v_h) \Phi_i$ where $\alpha_i(v_h) = 0$ if $Z_i \in \mathcal{N}_H \cap \Gamma_D$ and $\alpha_i(v_h) = P_{H,\omega_i}(v_h)(Z_i)$ in any other case.
- A similar construction is obtained by projecting locally into the space of discontinuous finite elements. Given $Z_i \in \mathcal{N}_H$ with corresponding nodal patch $\omega_i := \text{supp}(\Phi_i)$, we let \tilde{P}_{H,ω_i} denote the L^2 -projection onto the space of functions on ω_i that are affine on each coarse grid element (discontinuous P_1 finite elements on ω_i). For $v_h \in V_{h,\Gamma_D}$, we can now define $I_H(v_h) := \sum_{i=0}^{N_H-1} \alpha_i(v_h) \Phi_i$ where $\alpha_i(v_h) = |\omega_i|^{-1} \int_{\omega_i} \tilde{P}_{H,\omega_i}(v_h)$ for all active nodes and $\alpha_i(v_h) = 0$ if $Z_i \in \mathcal{N}_H \cap \Gamma_D$.
- An example for a projection $I_H : V_{h,\Gamma_D} \rightarrow V_{H,\Gamma_D}$ that is *not suitable* because it lacks the desired stability properties is the Lagrange (nodal) interpolation.

There are also many other choices for I_H , e.g., the orthogonal projection onto V_H with respect to the H^1 inner product and quasi-interpolation operators of Clément or Scott-Zhang type as they are well-established in the finite element community in the context of fast solvers and a posteriori error estimation [44, 45, 46, 47, 48]. For some problems, it can be advantageous to equip I_H with information about the problem, e.g., κ -weighted L^2 averaging for high-contrast problems [14]. As we see next, in practice we only require the kernel of the projection I_H for an implementation of

the method. This simplifies the computations significantly. For instance, for typical choices of I_H , there exist sets of functionals that can be used to decide if a function is in the kernel of I_H or not (cf. [35, Section 2.3]).

Once we decided for a suitable projection operator I_H , we can define the detail space

$$W_h := \{v_h \in V_{h,\Gamma_D} \mid I_H(v_h) = 0\}.$$

This detail space contains fine-scale functions in V_{h,Γ_D} that cannot be expressed in the coarse space V_{H,Γ_D} . In terms of the LOD we wish to correct classical nodal basis functions by an appropriate “detail function” from the space W_h . This can be achieved in a natural way by introducing the following elliptic decomposition of V_{h,Γ_D} . We refer to this decomposition as the \mathcal{A} -orthogonal splitting of V_{h,Γ_D} (cf. [1] for more details).

Definition 2.2 (\mathcal{A} -orthogonal decomposition). We define the $\mathcal{A}(\cdot, \cdot)$ -orthogonal complement of W_h in V_{h,Γ_D} by

$$V_{\text{LOD}} := \{v_h \in V_{h,\Gamma_D} \mid \mathcal{A}(v_h, w_h) = 0 \text{ for all } w_h \in W_h\}.$$

This is well-defined since $\mathcal{A}(\cdot, \cdot)$ is a scalar product on W_h . We obtain the (ideal) splitting

$$V_{h,\Gamma_D} = V_{\text{LOD}} \oplus W_h,$$

where $\dim(V_{\text{LOD}}) = \dim(V_{H,\Gamma_D})$.

We wish to use an approximation of V_{LOD} as a discrete solution space for Galerkin approximations. Observe that V_{LOD} is low dimensional, but practically expensive to assemble. Therefore we introduce a localized decomposition.

2.5 Localization to patches

To localize the splitting $V_{h,\Gamma_D} = V_{\text{LOD}} \oplus W_h$, we first need to localize the space W_h to patches $U(K) \subset \Omega$. We therefore introduce coarse-layer patches:

Definition 2.3 (Coarse-layer patch). For any positive $k \in \mathbb{N}$ and a coarse element $K \in \mathcal{T}_H$, we define patches $U_k(K)$ that consist of K itself and k -surrounding layers of coarse elements, i.e. we define $U_k(K)$ iteratively by

$$\begin{aligned} U_0(K) &:= K, \\ U_k(K) &:= \cup \{T \in \mathcal{T}_H \mid T \cap U_{k-1}(K) \neq \emptyset\} \quad k = 1, 2, \dots \end{aligned} \tag{2}$$

The restriction of W_h to a patch $U(K)$ is defined by $W_h(U(K)) := \{v_h \in W_h \mid v_h = 0 \text{ in } \Omega \setminus U(K)\}$. The localized decomposition can be now characterized using local correction operators.

Definition 2.4 (Correction Operators). For a given positive $k \in \mathbb{N}$ and for a coarse function $\Phi_H \in V_{H,\Gamma_D}$, the correction operator $\mathcal{Q}_h : V_{H,\Gamma_D} \rightarrow V_{h,\Gamma_D}$ is given by

$$\mathcal{Q}_h(\Phi_H) := \sum_{K \in \mathcal{T}_H} \mathcal{Q}_h^K(\Phi_H),$$

where $\mathcal{Q}_h^K(\Phi_H) \in W_h(U_k(K))$ (for $K \in \mathcal{T}_H$) is the solution of

$$\int_{U_k(K)} \kappa \nabla \mathcal{Q}_h^K(\Phi_H) \cdot \nabla w_h = - \int_K \kappa \nabla \Phi_H \cdot \nabla w_h \quad \text{for all } w_h \in W_h(U_k(K)). \tag{3}$$

We obtain the space $\{\Phi_H + \mathcal{Q}_h(\Phi_H) \mid \Phi_H \in V_{H,\Gamma_D}\}$ as a localized approximation of V_{LOD} . Here, localized is to be understood in the sense that there exists a local (nodal) basis $\{\Phi_i + \mathcal{Q}_h(\Phi_i) \mid Z_i \in \mathcal{N}_H \setminus \Gamma_D\}$ of V_{LOD} where the support of a basis function is restricted to $k+1$ layers of coarse elements around the corresponding node. The dimension of the new space is low (it is of the same dimension as V_{H,Γ_D}) and it can be constructed by solving the small problems (3), potentially in parallel. This generalized finite element space may be used in a Galerkin approximation of a prototypical linear elliptic model problem.

Example 2.5. Let $f \in L^2(\Omega)$ and let $u \in H_{\Gamma_D}^1(\Omega)$ solve

$$\mathcal{A}(u, v) = \int_{\Omega} f v \quad \text{for all } v \in H_{\Gamma_D}^1(\Omega).$$

Then the corresponding LOD approximation is given by $u_{\text{LOD}} = u_H + \mathcal{Q}_h(u_H)$, where $u_H \in V_{H,\Gamma_D}$ solves

$$\mathcal{A}(u_H + \mathcal{Q}_h(u_H), \Phi_H + \mathcal{Q}_h(\Phi_H)) = \int_{\Omega} f(\Phi_H + \mathcal{Q}_h(\Phi_H)) \quad \text{for all } \Phi_H \in V_{H,\Gamma_D}.$$

Remark 2.6 (Alternative iterative localization). There is an alternative characterization of localized correctors \mathcal{Q}_h via a preconditioned iterative solver that is based on a domain decomposition preconditioner as proposed in [39] (which in turn is based on [42]). To explain this alternative strategy we define nodal patches Ω_i as the union of all elements $T \in \mathcal{T}_H$ that share the vertex $Z_i \in \mathcal{N}_H$ and let

$$W_{h,i} := \{v - I_H v \mid v \in V_{h,\Gamma_D} : v|_{\Omega \setminus \Omega_i} \equiv 0\}. \quad (4)$$

The functions in $W_{h,i}$ are supported in a small neighbourhood of the vertex Z_i depending on the choice of I_H (typically within two layers of coarse elements). The $W_{h,i}$ are closed subspaces of the kernel W_h of I_H , see [39]. Let P_i be the \mathcal{A} -orthogonal projection from V_{h,Γ_D} to $W_{h,i}$, defined via the equation

$$\mathcal{A}(P_i v, w) = \mathcal{A}(v, w), \quad \forall w \in W_{h,i}. \quad (5)$$

With this, we introduce an operator P as

$$P = P_0 + P_1 + \dots + P_{N_H-1}. \quad (6)$$

The operator P is symmetric with respect to the bilinear form $\mathcal{A}(\cdot, \cdot)$ and in [42, 39] it is shown that it is a quasi-optimal preconditioner for the ideal corrector \mathcal{Q}_h^∞ , i.e. for the \mathcal{A} -orthogonal projection from V_{H,Γ_D} onto W_h . Starting from $\mathcal{Q}_h^0(\bullet) = 0$, localized approximations \mathcal{Q}_h^j of \mathcal{Q}_h^∞ can be defined, for any $\Phi_H \in V_{H,\Gamma_D}$, via the iteration

$$\mathcal{Q}_h^j(\Phi_H) = \mathcal{Q}_h^{j-1}(\Phi_H) + \vartheta P(\Phi_H - \mathcal{Q}_h^{j-1}(\Phi_H)), \quad j = 1, 2, \dots, k. \quad (7)$$

Note the information is spread at most by a fixed number of layers (typically two) of coarse elements in each step, that is, the support of $\mathcal{Q}_h^k(\Phi_H)$ is at most $\mathcal{O}(k)$ layers larger than the support of Φ_H . If the scaling factor $\vartheta > 0$ is chosen appropriately as discussed in [39, p. 2771], the iteration converges linearly, thereby producing a localized corrector $\mathcal{Q}_h = \mathcal{Q}_h^k$ that is close to the one given in Definition 2.4, not equal in general though. Please note that the parameter k has a slightly different meaning in the two variants. To avoid the educated guess of ϑ and to achieve a more accurate global approximation one may consider the enlarged coarse space

$$\{\Phi_H + \mathcal{Q}_h^j(\Phi_H) \mid \Phi_H \in V_{H,\Gamma_D}, j = 1, 2, \dots, k\}$$

which underlies the error analysis of [39]. Here, the choice $\vartheta = 1$ is appropriate. This space provides the sharpest global error bounds at the price that the dimension of the final multiscale

space is k -times larger compared to the dimension of the LOD multiscale space $(\text{Id} + \mathcal{Q}_h)(V_{H,\Gamma_D})$ used in Example 2.5.

In terms of computational cost, the iterative computation of the corrector(s) is comparable to the original variant. Since only local problems of type (5) need to be solved the possible degree of parallelism is slightly larger. However, in this paper we trade a slight reduction of offline efficiency for a better online performance which is achieved by the original variant as it typically produces more localized functions for given fixed accuracy.

We shall emphasize that this preference for the online efficiency pays off only if sufficiently many online computations are to be performed. If a source problem as given in Example 2.5 is to be solved for a very small number of right-hand sides the offline cost becomes relevant. In this case the iterative corrector computation seems favorable, in particular the variant introduced in [42]. Here, the correctors are computed on the fly during a preconditioned iterative solution of the full problem in the spirit of (7) (the corrector is only applied to the current approximation of the solution rather than precomputing it for all basis functions in the coarse space). For a more detailed discussion of this variant we refer to [42] and to [49] for a comparison to the original variant and some numerical results.

In the following sections we quantify the approximation properties of u_{LOD} depending on the choice of the localization parameter k .

3 The algebraic realization of the correctors \mathcal{Q}_h

Before we start to give a first example, we need to discuss how the local problems (3) can be assembled and solved practically. In particular we show how to interpret the corrector \mathcal{Q}_h on an algebraic level. We employ the notation of matrices as we believe it eases reading but want to point out that in an optimized implementation, most or all linear operators can be implemented matrix-free.

3.1 Analytic preliminaries

We start by introducing a general terminology that we use subsequently in the context of localization. Every patch $U_k(K_\ell)$ with $K_\ell \in \mathcal{T}_H$ and $k \in \mathbb{N}$ is in the following directly associated with the index ℓ (where $0 \leq \ell < N_{\mathcal{T}_H}$). In particular, we consider the localization parameter k to be fixed and hence drop it when defining $\mathcal{U}_\ell := U_k(K_\ell)$. For a given patch \mathcal{U}_ℓ we denote the sets of active coarse and fine nodes in \mathcal{U}_ℓ respectively by

$$\begin{aligned} \mathcal{N}_{\ell,H} &= \{Z_{\ell,i} \in \mathcal{N}_H \mid Z_{\ell,i} \in \overline{\mathcal{U}_\ell} \setminus \Gamma_D\} & \text{and} \\ \mathcal{N}_{\ell,h} &= \{z_{\ell,j} \in \mathcal{N}_h \mid z_{\ell,j} \in \overline{\mathcal{U}_\ell} \setminus (\partial\mathcal{U}_\ell \setminus \Gamma_N)\}. \end{aligned}$$

Furthermore, we set $N_{\ell,H} := |\mathcal{N}_{\ell,H}|$ and $N_{\ell,h} := |\mathcal{N}_{\ell,h}|$ the number of nodes. The corresponding coarse and fine Lagrange basis functions in the patch \mathcal{U}_ℓ are denoted respectively by $\Phi_{\ell,i}$ (i.e. $\Phi_{\ell,i}$ is coarse nodal basis function for node $Z_{\ell,i} \in \mathcal{N}_{\ell,H}$) and $\phi_{\ell,j}$ (i.e. $\phi_{\ell,j}$ is fine nodal basis function for node $z_{\ell,j} \in \mathcal{N}_{\ell,h}$). With this notation, we define the corresponding local basis function sets.

Definition 3.1 (Local basis sets). For each patch \mathcal{U}_ℓ we define

$$V_{h,\ell} := \{v_h \in V_h \mid v_h(z_{\ell,j}) = 0 \text{ for } z_{\ell,j} \in \mathcal{N}_h \setminus \mathcal{N}_{\ell,h}\}$$

and we let

$$\{\phi_{\ell,j} \mid 0 \leq j < N_{\ell,h}\} \subset V_{h,\ell}$$

be the set of fine Lagrange basis functions that belong to the active fine nodes in \mathcal{U}_ℓ (i.e. to $\mathcal{N}_{\ell,h}$). Accordingly we let

$$\{\Phi_{\ell,i} | 0 \leq i < N_{\ell,H}\} \subset V_H$$

denote the ordered set of all active coarse Lagrange basis functions in \mathcal{U}_ℓ , i.e. the coarse basis functions associated with the nodes in $\mathcal{N}_{\ell,H}$.

Since each $K_\ell \in \mathcal{T}_H$ contains c_d coarse nodes, we can order the global indices of these nodes by $p_0(\ell) < p_1(\ell) < \dots < p_{c_d-1}(\ell)$. This implies $\mathcal{Q}_h^{K_\ell}(\Phi_j) = 0$ for all $j \notin \{p_0(\ell), \dots, p_{c_d-1}(\ell)\}$, hence we only need to compute $\mathcal{Q}_h^{K_\ell}(\Phi_{p_i(\ell)})$ for $i \in \{0, \dots, c_d - 1\}$. For arbitrary $\Phi_H \in V_{H,\Gamma_D}$ we can hence write

$$\mathcal{Q}_h(\Phi_H) = \sum_{K_\ell \in \mathcal{T}_H} \sum_{i=0}^{c_d-1} \Phi_H(Z_{p_i(\ell)}) \mathcal{Q}_h^{K_\ell}(\Phi_{p_i(\ell)}). \quad (8)$$

3.2 Algebraic preliminaries

In this section, we introduce some algebraic preliminaries. As a general notation in this paper we denote for any matrix $\mathcal{M} \in \mathbb{R}^{m \times n}$ the transposed of the i 'th row of \mathcal{M} by $\mathcal{M}[i] \in \mathbb{R}^n$ (for $0 \leq i < m$), i.e.

$$\mathcal{M}[i] = \begin{pmatrix} \mathcal{M}_{i,0} \\ \vdots \\ \mathcal{M}_{i,n-1} \end{pmatrix}.$$

The entry at position (i, j) is denoted by $\mathcal{M}[i][j]$.

Definition 3.2 (Local-to-global-mapping). Let c_d denote the number of nodes in a grid element, i.e. $c_d = d + 1$ if \mathcal{T}_H and \mathcal{T}_h consist of simplicial elements and $c_d = 2^d$ if \mathcal{T}_H and \mathcal{T}_h consist of quadrilateral elements. Then we call

$$\sigma : \{0, \dots, N_{\mathcal{T}_h} - 1\} \times \{0, \dots, c_d - 1\} \rightarrow \{0, \dots, N_h - 1\}$$

the local-to-global-mapping for the fine grid, if it maps the local index of a node m in an element $T_t \in \mathcal{T}_h$ to its global index j . Here, $t \in \{0, \dots, N_{\mathcal{T}_h} - 1\}$ denotes the index of the element $T_t \in \mathcal{T}_h$. We write $\sigma(t, m) = j$. An example is given in Figure 1. For $T_t \in \mathcal{T}_h$ the algebraic version of $\sigma(t, \cdot)$ is given by the matrix $\boldsymbol{\sigma}_t \in \mathbb{R}^{N_h \times c_d}$ where for $0 \leq m < c_d$ and $0 \leq j < N_h$

$$\boldsymbol{\sigma}_t[m][j] := \begin{cases} 1 & \text{if } \sigma(t, m) = j, \\ 0 & \text{else.} \end{cases} \quad (9)$$

Definition 3.3 (Element stiffness and mass matrices). For $T_t \in \mathcal{T}_h$ with index t we define the (fine grid) element stiffness matrix $\mathbf{A}_t \in \mathbb{R}^{c_d \times c_d}$ by

$$\mathbf{A}_t[m][n] := \int_{T_t} \kappa \nabla \phi_{\sigma(t,n)} \cdot \nabla \phi_{\sigma(t,m)} \quad \text{for all } 0 \leq n, m < c_d$$

and the (fine grid) element mass matrix $\mathbf{M}_t \in \mathbb{R}^{c_d \times c_d}$ by

$$\mathbf{M}_t[m][n] := \int_{T_t} \phi_{\sigma(t,n)} \phi_{\sigma(t,m)} \quad \text{for all } 0 \leq n, m < c_d.$$

The corresponding global block matrices that store all these element contributions (i.e. stiffness and mass matrix in a discontinuous Galerkin discretization on the fine grid) shall be denoted by \mathbf{A}_{dc} and \mathbf{M}_{dc} respectively.




Figure 1: Example: the fine grid \mathcal{T}_h consists of 16 quadratic grid elements. We start counting with 0. For the 6'th element T_5 (i.e. $t = 5$), we find 4 global nodes with global indices 5, 6, 10 and 11. Each of these global indices j , can be mapped to a local index m and vice versa. This is done by σ . For instance, for $t = 5$, we have $\sigma(t, 0) = 5$, $\sigma(t, 1) = 6$, $\sigma(t, 2) = 11$ and $\sigma(t, 3) = 10$. For $m = 1$ and $j = 6$ as in the graphic, we have $\boldsymbol{\sigma}_t[m][j] = 1$, because $\sigma(t, m) = j$.

Definition 3.4 (Global stiffness and mass matrices). By $\mathbf{A}_h \in \mathbb{R}^{N_h \times N_h}$ we denote the global fine stiffness matrix with entries $\mathbf{A}_h[i][j] = (\kappa \nabla \phi_j, \nabla \phi_i)_{L^2(\Omega)}$ and by $\mathbf{M}_h \in \mathbb{R}^{N_h \times N_h}$ the global fine mass matrix with entries $\mathbf{M}_h[i][j] = (\phi_j, \phi_i)_{L^2(\Omega)}$. Analogously we denote by $\mathbf{A}_H \in \mathbb{R}^{N_H \times N_H}$ the coarse stiffness matrix and by $\mathbf{M}_H \in \mathbb{R}^{N_H \times N_H}$ the coarse mass matrix. Note that the above matrices are with respect to *all* nodes including the whole boundary.

Assuming that the local contributions $\mathbf{A}_t, \mathbf{M}_t \in \mathbb{R}^{c_d \times c_d}$ are computed for all $t \in \mathcal{T}_h$ and assuming that $\boldsymbol{\sigma}_t$ is available, we can assemble the global (fine) stiffness and mass matrix by

$$\mathbf{A}_h = \sum_{t \in \mathcal{T}_h} \boldsymbol{\sigma}_t \mathbf{A}_t \boldsymbol{\sigma}_t^\top \quad \text{and} \quad \mathbf{M}_h = \sum_{t \in \mathcal{T}_h} \boldsymbol{\sigma}_t \mathbf{M}_t \boldsymbol{\sigma}_t^\top. \quad (10)$$

Since V_h and V_H do still incorporate all boundary nodes on $\partial\Omega$, we require matrices that erases the unnecessary rows and columns in the system matrices (and load vectors) that are associated with the DOFs on the Dirichlet boundary part Γ_D .

Definition 3.5 (Boundary correction matrices). We define the boundary correction (or restriction) matrices by $\mathbf{B}^h \in \mathbb{R}^{N_h \times N_h}$ by

$$\mathbf{B}^h[i][j] = \begin{cases} 1 & \text{if } i = j \text{ and } z_i \in \mathcal{N}_h \setminus \Gamma_D \\ 0 & \text{else.} \end{cases}$$

and analogously $\mathbf{B}^H \in \mathbb{R}^{N_H \times N_H}$ by

$$\mathbf{B}^H[i][j] = \begin{cases} 1 & \text{if } i = j \text{ and } Z_i \in \mathcal{N}_H \setminus \Gamma_D \\ 0 & \text{else.} \end{cases}$$

As the coarse boundary condition vector is in general not available, we describe an easy way to compute it. We define a vertex map that receives the index of a coarse node and the index of a fine node. If the coordinates of the coarse node are identical to the coordinates of the fine node,

the vertex map is 1 (true). In any other case, the vertex map is 0 (false). In algebraic form, we describe the vertex map by the matrix $\mathbf{V}^h \in \mathbb{R}^{N_H \times N_h}$ that is given by

$$\mathbf{V}^h[i][j] := \begin{cases} 1 & \text{if } \text{coord}(Z_i) = \text{coord}(z_j), \quad \text{with } Z_i \in \mathcal{N}_H, z_j \in \mathcal{N}_h \\ 0 & \text{else.} \end{cases}$$

For instance, the matrix \mathbf{V}^h can be easily computed by using an interpolation matrix \mathbf{P}_h as defined in (14) below. If \mathbf{P}_h denotes such a matrix (expressing a function on the coarse grid in terms of fine degrees of freedom) we can define

$$\mathbf{V}^h[i][j] := (\mathbf{P}_h[i][j] \equiv 1.0).$$

3.3 Local restriction matrices

In order to localize computations to a patch \mathcal{U}_ℓ with $K_\ell \in \mathcal{T}_H$ we require a restriction operator $\mathcal{R}_\ell : V_h \rightarrow V_{h,\ell}$. For $v_h \in V_h$ we define the nodal interpolation $\mathcal{R}_\ell(v_h) \in V_{h,\ell}$ by

$$\mathcal{R}_\ell(v_h)(z) = \begin{cases} v_h(z) & \text{if } z \in \mathcal{N}_{\ell,h}, \\ 0 & \text{else.} \end{cases}$$

The algebraic version of the restriction operator is denoted by $\mathbf{R}_\ell^h \in \mathbb{R}^{N_{\ell,h} \times N_h}$ and defined by the entries

$$\mathbf{R}_\ell^h[i][j] = \begin{cases} 1 & \text{if } z_{\ell,i} = z_j \\ 0 & \text{else} \end{cases} \quad (11)$$

and where i is the index that corresponds to the fine node $z_i \in \mathcal{N}_h$. Hence, we get the local mass matrix \mathbf{M}_ℓ (respectively local stiffness matrix \mathbf{A}_ℓ) from the global mass matrix \mathbf{M}_h (respectively stiffness matrix \mathbf{A}_h) by matrix multiplication, i.e.

$$\begin{aligned} \mathbf{M}_\ell &= \mathbf{R}_\ell^h \mathbf{M}_h \mathbf{R}_\ell^{h\top} & \mathbf{M}_\ell \in \mathbb{R}^{N_{\ell,h} \times N_{\ell,h}} \text{ is local mass matrix,} \\ \mathbf{A}_\ell &= \mathbf{R}_\ell^h \mathbf{A}_h \mathbf{R}_\ell^{h\top} & \mathbf{A}_\ell \in \mathbb{R}^{N_{\ell,h} \times N_{\ell,h}} \text{ is local stiffness matrix.} \end{aligned}$$

Recall that the entries of \mathbf{A}_ℓ are given by

$$\mathbf{A}_\ell[i][j] = \mathcal{A}(\phi_{\ell,j}, \phi_{\ell,i}) \quad \text{for } 0 \leq i, j < N_{\ell,h}.$$

Besides the restriction \mathbf{R}_ℓ^h to fine grid nodes in \mathcal{U}_ℓ we also require a restriction \mathbf{R}_ℓ^H to active coarse grid nodes in \mathcal{U}_ℓ . We can define $\mathbf{R}_\ell^H \in \mathbb{R}^{N_{\ell,H} \times N_H}$ analogously by the entries

$$\mathbf{R}_\ell^H[i][j] := \begin{cases} 1 & \text{if } Z_{\ell,i} = Z_j \\ 0 & \text{else} \end{cases} \quad (12)$$

and where j is the global index that corresponds with the coarse node $Z_j \in \mathcal{N}_H$. The local restriction matrices \mathbf{R}_ℓ^H and \mathbf{R}_ℓ^h need to be stored only temporary. Both matrices (as well as \mathbf{A}_ℓ) can be deleted as soon as the ℓ 'th corrector matrix is computed (cf. Section 3.6).

3.4 An algebraic characterization of the space W_h

Recall the notation

$$V_H = \text{span}\{\Phi_i | 0 \leq i \leq N_H - 1\} \quad \text{and} \quad V_h = \text{span}\{\phi_j | 0 \leq j \leq N_h - 1\}.$$

This subsection describes how we can characterize the kernel of the projection I_H . In order to illustrate our method, we restrict our considerations to the choice that I_H defines the L^2 -projection. It is obvious that a function $v_h \in V_{h,\Gamma_D}$ is in the kernel of the global L^2 -projection (i.e. $I_H(v_h) = 0$) if it holds

$$(v_h, \Phi_i)_{L^2(\Omega)} = 0 \quad \text{for all } \Phi_i \in \mathcal{N}_H \setminus \Gamma_D. \quad (13)$$

In order to handle this constraint, we first observe that any coarse basis function Φ_i can be easily expressed in terms of fine basis functions by

$$\Phi_i = \sum_{j=0}^{N_h-1} \Phi_i(z_j) \phi_j.$$

Hence we have

$$(\Phi_i, \Phi_j)_{L^2(\Omega)} = \sum_{k,\ell=0}^{N_h-1} \Phi_i(z_k) (\phi_k, \phi_\ell)_{L^2(\Omega)} \Phi_j(z_\ell).$$

Consequently, we can define the projection matrix $\mathbf{P}_h \in \mathbb{R}^{N_H \times N_h}$ from the coarse-mesh Lagrange space to the fine-mesh Lagrange space by

$$\mathbf{P}_h := \begin{pmatrix} \Phi_0(z_0) & \cdots & \Phi_0(z_{N_h-1}) \\ \vdots & \ddots & \vdots \\ \Phi_{N_H-1}(z_0) & \cdots & \Phi_{N_H-1}(z_{N_h-1}) \end{pmatrix} \quad (14)$$

and relate the coarse and the fine mass matrix via

$$\mathbf{M}_H = \mathbf{P}_h \mathbf{M}_h \mathbf{P}_h^\top.$$

With that we can see that the analytical constraint (13), i.e. $I_H(v) = 0$, can be equivalently expressed through the algebraic constraint $(\mathbf{B}^H \mathbf{P}_h \mathbf{M}_h) \mathbf{v} = 0$. Motivated by these considerations we define the global constraint matrix $\mathbf{C}_h \in \mathbb{R}^{N_H \times N_h}$ by

$$\mathbf{C}_h := \mathbf{P}_h \mathbf{M}_h.$$

Remark 3.6. The definition of $\mathbf{C}_h = \mathbf{P}_h \mathbf{M}_h$ might be surprising since it does not account for the distinction between Dirichlet-nodes and Neumann-nodes. In fact, the natural way is to define $\mathbf{C}_h := \mathbf{B}^H \mathbf{P}_h \mathbf{M}_h$, where the restriction (or boundary correction) matrix \mathbf{B}^H is used to remove the coarse basis functions associated with nodes on Γ_D . However, the boundary matrix \mathbf{B}^H causes that $\mathbf{B}^H \mathbf{P}_h \mathbf{M}_h$ has not a maximal rank and hence the arising saddle point problem would be singular. This would cause numerical issues for the method with patches $\mathcal{U}_\ell = \Omega$. For that reason, we define $\mathbf{C}_h = \mathbf{P}_h \mathbf{M}_h$ and note that the arising constraint would be stronger than condition (13). However, the smaller error that we make in the definition of \mathbf{C}_h is already corrected in the next step by using local restrictions \mathbf{R}_ℓ^H and \mathbf{R}_ℓ^h .

Recalling the definition of the local restriction matrix \mathbf{R}_ℓ^h given by (11), we can define the localization of \mathbf{C}_h to the patch \mathcal{U}_ℓ (and to the correct boundary nodes) by

$$\mathbf{C}_\ell = \mathbf{R}_\ell^H \mathbf{C}_h \mathbf{R}_\ell^{h\top} \quad \text{where } \mathbf{C}_\ell \in \mathbb{R}^{N_{\ell,H} \times N_{\ell,h}} \text{ is the local constraint matrix.}$$

Observe that if $\mathbf{C}_\ell \mathbf{v} = 0$ for some $\mathbf{v} \in \mathbb{R}^{N_{\ell,h}}$, then the represented function $v = \sum_{j=0}^{N_{\ell,h}-1} \mathbf{v}_j \phi_{\ell,j} \in V_{h,\ell}$ has the property

$$0 = \int_{\mathcal{U}_\ell} v \Phi_{\ell,i} = \int_{\Omega} v \Phi_{\ell,i} \quad \text{for all } 0 \leq i \leq N_{\ell,H} - 1,$$

i.e. L^2 -orthogonality for all coarse basis functions Φ_j that have a support intersecting $\mathcal{U}_\ell = \text{supp } v$. Consequently, the property holds for the whole space, $v \perp_{L^2} V_{H,\Gamma_D} = 0$, which means that v is in the kernel of L^2 -projection I_H as desired. The converse conclusion follows analogously.

Remark 3.7. The matrix \mathbf{C}_ℓ fully represents the local constraints and maps a function from the fine scale finite element space $V_{h,\ell}$ onto the coarse finite element space V_H (restricted to the local subdomain \mathcal{U}_ℓ). Since there are only constraints for coarse vertices, the dimension of the first component of \mathbf{C}_ℓ is small. If we use a localization parameter k with $k \simeq C|\log(H)|$ (as it will be suggested by Theorem 4.2 below), we obtain that $N_{\ell,h} = \mathcal{O}((H|\log(H)|)/h)^{-d}$ (which is the dimension of $V_{h,\ell}$) and that $N_{H,\ell}$ (the number of coarse nodes in \mathcal{U}_ℓ) grows proportional to $|\log(H)|^d$.

3.5 Assembling of a local load vector

Let us again fix some coarse element $K_\ell \in \mathcal{T}_H$. Beside assembling the local stiffness matrices and the local constraints matrix, we also need to compute the load vector that corresponds to the right hand side in (3), i.e. the term

$$- \int_{K_\ell} \kappa \nabla \Phi_i \cdot \nabla \phi_j$$

for every coarse and fine basis function $\Phi_i \in V_{H,\Gamma_D}$ and $\phi_j \in V_{h,\Gamma_D}$ with support on K_ℓ . We start with defining a suitable (algebraic) restriction operator for coarse grid nodes (in K_ℓ).

Definition 3.8 (Coarse-node-in-coarse-element restriction). Recall that each coarse element $K_\ell \in \mathcal{T}_H$ contains c_d coarse nodes. Let the global indices of these nodes be denoted by $p_0(\ell) < p_1(\ell) < \dots < p_{c_d-1}(\ell)$. Then the coarse element restriction matrix $\mathbf{T}_\ell^H \in \mathbb{R}^{c_d \times N_H}$ is given for $0 \leq i < c_d$ and $0 \leq j < N_H$ by

$$\mathbf{T}_\ell^H[i][j] = \begin{cases} 1 & \text{if } j = p_i(\ell) \\ 0 & \text{else.} \end{cases} \quad (15)$$

As for the global stiffness matrix \mathbf{A}_h in (10), we can obtain the local stiffness matrix on K_ℓ (i.e. with entries $(\kappa \nabla \phi_i, \nabla \phi_j)_{L^2(K_\ell)}$) from the element stiffness matrices \mathbf{A}_t . After that, we can restrict the resulting matrix to the fine basis functions that belong to fine nodes $z_i \in \mathcal{U}_\ell$ by using \mathbf{R}_ℓ^h (see (11)). We obtain for $0 \leq j < N_h$ and $0 \leq i < N_{\ell,h}$

$$\left(\left(\sum_{\substack{t \in \mathcal{T}_h \\ t \subset K_\ell}} \boldsymbol{\sigma}_t \mathbf{A}_t \boldsymbol{\sigma}_t^\top \right) \mathbf{R}_\ell^{h\top} \right) [j][i] = (\kappa \nabla \phi_{\ell,i}, \nabla \phi_j)_{L^2(K_\ell)}.$$

Consequently we can define the matrix $\mathbf{r}_\ell \in \mathbb{R}^{c_d \times N_{\ell,h}}$ that stores the c_d load vectors as its rows

$$\mathbf{r}_\ell := -\mathbf{T}_\ell^H \mathbf{B}^H \mathbf{P}_h \left(\sum_{\substack{t \in \mathcal{T}_h \\ t \subset K_\ell}} \boldsymbol{\sigma}_t \mathbf{A}_t \boldsymbol{\sigma}_t^\top \right) \mathbf{R}_\ell^{h^\top}, \quad (16)$$

i.e. \mathbf{r}_ℓ as defined above fulfills either $\mathbf{r}_\ell[i] = 0$ for $Z_{p_i(\ell)} \in \Gamma_D$ or else its rows transposed are given by

$$\mathbf{r}_\ell[i] = - \begin{pmatrix} (\kappa \nabla \Phi_{p_i(\ell)}, \nabla \phi_{\ell,0})_{L^2(K_\ell)}, \\ \vdots \\ (\kappa \nabla \Phi_{p_i(\ell)}, \nabla \phi_{\ell,M_{\ell,h}-1})_{L^2(K_\ell)} \end{pmatrix},$$

where $\Phi_{p_i(\ell)}$ is the $p_i(\ell)$ 'th coarse basis function (i.e. the global index is $p_i(\ell)$ and the local index in K_ℓ is i). If $\mathbf{r}_\ell[i] = 0$, no local problem has to be solved and the local corrector is zero.

3.6 Assembly and solution of a local problem

Observe that (3) must be solved for every $K_\ell \in \mathcal{T}_H$ and every coarse basis function Φ_m that has a support on K_ℓ (i.e. for $\Phi_{p_i(\ell)}$ with $0 \leq i < c_d$, except the ones that belong to nodes on Γ_D) and recall that the correct boundary condition on $\partial \mathcal{U}_\ell$ is already included in the local stiffness matrix \mathbf{A}_ℓ . Let us fix $\ell \in \{0, \dots, |\mathcal{T}_H| - 1\}$ and a coarse basis function $\Phi_{p_i(\ell)}$. In the light of the discussion in Section 3.4, we can formulate the local problem (3) in the following way.

Definition 3.9 (Continuous formulation of a local problem). Let $0 \leq i < c_d$ and let us denote $w_{\ell,i} := \mathcal{Q}_h^{K_\ell}(\Phi_{p_i(\ell)})$. Then $w_{\ell,i} \in V_h(\mathcal{U}_\ell)$ is characterized by the property $I_H(w_{\ell,i}) = 0$ and the property that it solves

$$\int_{\mathcal{U}_\ell} \kappa \nabla w_{\ell,i} \cdot \nabla w_h = - \int_{K_\ell} \kappa \nabla \Phi_{p_i(\ell)} \cdot \nabla w_h \quad (17)$$

for all $w_h \in V_h(\mathcal{U}_\ell)$ with $I_H(w_h) = 0$.

Problem (17) can be obviously interpreted as a saddle point problem. Hence, we obtain the following algebraic formulation using the notation from the previous subsections.

Definition 3.10 (Algebraic formulation of a local problem). Let $0 \leq i < c_d$. The algebraic version of problem (17) is the following saddle point problem. Find the tuple $(\mathbf{w}_\ell[i], \boldsymbol{\lambda}_\ell[i]) \in \mathbb{R}^{N_{\ell,h}} \times \mathbb{R}^{N_{\ell,H}}$; with

$$\begin{aligned} \mathbf{A}_\ell \mathbf{w}_\ell[i] + \mathbf{C}_\ell^\top \boldsymbol{\lambda}_\ell[i] &= \mathbf{r}_\ell[i] \\ \mathbf{C}_\ell \mathbf{w}_\ell[i] &= 0. \end{aligned} \quad (18)$$

Here, $\mathbf{w}_\ell[i]$ is the coefficient vector for the solution $w_{\ell,i} = \mathcal{Q}_h^{K_\ell}(\Phi_{p_i(\ell)})$ of (17), i.e.

$$w_{\ell,i} = \sum_{j=0}^{N_{\ell,h}-1} \mathbf{w}_\ell[i][j] \phi_{\ell,j},$$

and $\boldsymbol{\lambda}_\ell[i]$ is the corresponding Lagrange multiplier.

We can state this problem also in Schur complement formulation.

Remark 3.11 (Schur complement). The Schur complement matrix \mathbf{S}_ℓ associated with problem (18) is given by

$$\mathbf{S}_\ell := (\mathbf{C}_\ell \mathbf{A}_\ell^{-1} \mathbf{C}_\ell^\top). \quad (19)$$

Hence the solution $\mathbf{w}_\ell[i]$ of (18) can be written as

$$\mathbf{w}_\ell[i] = \mathbf{A}_\ell^{-1} \mathbf{r}_\ell[i] - (\mathbf{A}_\ell^{-1} \mathbf{C}_\ell^\top) \boldsymbol{\lambda}_\ell[i], \quad (20)$$

where $\boldsymbol{\lambda}_\ell[i] \in \mathbb{R}^{N_{\ell,H}}$ solves

$$\mathbf{S}_\ell \boldsymbol{\lambda}_\ell[i] = (\mathbf{C}_\ell \mathbf{A}_\ell^{-1}) \mathbf{r}_\ell[i]. \quad (21)$$

The common approach is to solve systems such as (19)-(21) iteratively with an approximate Schur complement matrix. As the system (21) is only of size $N_{\ell,H} \times N_{\ell,H}$ and must be solved c_d times (for different righthand sides corresponding to each coarse basis function with support in \mathcal{K}_ℓ) it is faster to compute the whole Schur-complement matrix, solve it directly and apply back-substitution for each right-hand-side vector. Solving the local problem (18) for all (transposed) rows of the matrix $\mathbf{r}_\ell \in \mathbb{R}^{c_d \times N_{\ell,h}}$ can be hence obtained in the following way. It involves a pre-processing step that is independent of \mathbf{r}_ℓ and a post-processing step that must be performed for each row of \mathbf{r}_ℓ .

Pre-processing steps.

1. Compute the matrix $\mathbf{Y}_\ell := \mathbf{A}_\ell^{-1} \mathbf{C}_\ell^\top$. This involves to solve $N_{\ell,H}$ problems of size $N_{\ell,h} \times N_{\ell,h}$, i.e. for all $0 \leq m < N_{\ell,H}$ we need to solve for $\mathbf{Y}_\ell[m] \in \mathbb{R}^{N_{\ell,h}}$ with

$$\mathbf{A}_\ell \mathbf{Y}_\ell[m] = \mathbf{C}_\ell^\top[m].$$

The $(1 \times N_{\ell,h})$ -matrix $(\mathbf{Y}_\ell[m])^\top$ forms the m 'th row of \mathbf{Y}_ℓ .

2. Assemble the Schur complement $\mathbf{S}_\ell = \mathbf{C}_\ell \mathbf{Y}_\ell$ by matrix multiplication and compute \mathbf{S}_ℓ^{-1} . Since \mathbf{S}_ℓ is only a $(N_{\ell,H} \times N_{\ell,H})$ -matrix its inversion is cheap.

Post-processing steps for all $0 \leq i < c_d$.

1. If $\mathbf{r}_\ell[i] \neq 0$, solve for $\mathbf{q}_\ell[i] \in \mathbb{R}^{N_{\ell,h}}$ with $\mathbf{A}_\ell \mathbf{q}_\ell[i] = \mathbf{r}_\ell[i]$.
2. Since \mathbf{S}_ℓ^{-1} is precomputed, we obtain $\boldsymbol{\lambda}_\ell[i]$ from $\mathbf{q}_\ell[i]$ via equation (21), i.e. set

$$\boldsymbol{\lambda}_\ell[i] = \mathbf{S}_\ell^{-1} \mathbf{C}_\ell \mathbf{q}_\ell[i].$$

3. Using the precomputed matrix $\mathbf{Y}_\ell = \mathbf{A}_\ell^{-1} \mathbf{C}_\ell^\top$ and inserting $\boldsymbol{\lambda}_\ell[i]$ in (20) we obtain

$$\mathbf{w}_\ell[i] = \mathbf{q}_\ell[i] - \mathbf{Y}_\ell \boldsymbol{\lambda}_\ell[i].$$

When the post-processing step is concluded for the ℓ 'th local problem, a term of the form $(\mathbf{T}_\ell^H)^\top \mathbf{w}_\ell \mathbf{R}_\ell^h$ needs to be stored in a global corrector matrix \mathbf{Q}_h . Once this is done, all local matrices involved in the pre- and post-processing steps above are no longer required and can be deleted.

Remark 3.12 (Cost). Recall that $N_{\ell,H} \approx |\log(H)|^d$ is typically a small number (see also Remark 3.7 above). The pre-processing step requires to solve $N_{\ell,H}$ equations of size $N_{\ell,h} \times N_{\ell,h}$ and to invert one matrix of size $N_{\ell,H} \times N_{\ell,H}$ (cost $\mathcal{O}(N_{\ell,H}^3)$). And in the post-processing step, for each $i \in \{0, \dots, c_d - 1\}$, we only need to solve one additional problem. In total, for one patch \mathcal{U}_ℓ , the procedure involves to solve $(c_d + N_{\ell,H})$ equations of dimension $N_{\ell,h} \times N_{\ell,h}$ and $N_{\ell,H}$ equations of dimension $N_{\ell,H} \times N_{\ell,H}$. This also justifies why we solve the saddle point problem (18) with a direct inversion of the Schur complement instead of using an iterative solver like the Uzawa solver. Roughly speaking, if the average number of iterations of an iterative solver is larger than $(c_d + N_{\ell,H})/c_d$, then the direct inversion above is the cheaper approach. This is in most cases fulfilled.

Remark 3.13. Note that we can practically use the fact that the Lagrange basis functions of V_H have a partition of unity property, which implies that it is only required to solve the local corrector problem (3) $d \cdot |\mathcal{T}_H|$ times in the case of a triangulation and $(d + 1) \cdot |\mathcal{T}_H|$ times in the case of a quadrilateral. We do not consider this in the algorithms. However, a corresponding modification is straightforward.

Remark 3.14. The algorithm can also be formulated with the boundary matrix \mathbf{B}^H . In this case, there is a small overhead in terms of the number of local problems to be solved, i.e. we solve problems for right hand sides $\mathbf{r}_\ell[i]$ that correspond to inactive coarse basis functions (basis functions belonging to Dirichlet-nodes).

3.7 The global corrector matrix

To store the information that we obtained from the solutions of the local problems, we introduce the global corrector matrix $\mathbf{Q}_h \in \mathbb{R}^{N_H \times N_h}$.

Definition 3.15 (Global corrector matrix \mathbf{Q}_h). Recall the matrix $\mathbf{w}_\ell \in \mathbb{R}^{c_d \times N_{\ell,h}}$ introduced in Definition 3.10 and recall that it is related to the correctors $\mathcal{Q}_h^{K_\ell}(\Phi_{p_i(\ell)})$ (solving equation (17)) via

$$\mathcal{Q}_h^{K_\ell}(\Phi_{p_i(\ell)}) = \sum_{j=0}^{N_{\ell,h}-1} \mathbf{w}_\ell[i][j] \phi_{\ell,j}.$$

With (8) and the previously defined local restriction matrices $\mathbf{R}_\ell^h \in \mathbb{R}^{N_{\ell,h} \times N_h}$ and $\mathbf{T}_\ell^H \in \mathbb{R}^{c_d \times N_H}$ we get the corrector matrix $\mathbf{Q}_h \in \mathbb{R}^{N_H \times N_h}$ via

$$\mathbf{Q}_h := \sum_{K_\ell \in \mathcal{T}_H} (\mathbf{T}_\ell^H)^\top \mathbf{w}_\ell \mathbf{R}_\ell^h.$$

Hence, for any coarse function $\Phi_H \in V_{H,\Gamma_D}$, we can compute $Q_h(\Phi_H) \in V_{h,\Gamma_D}$ easily from \mathbf{Q}_h . For instance, let $\Phi \in \mathbb{R}^{N_H}$ be the vector with entries $\Phi[i] = \Phi_H(Z_i)$. Then we have

$$Q_h(\Phi_H) = \sum_{i=1}^{N_H-1} \left(\mathbf{Q}_h^\top \Phi \right)[i] \Phi_i.$$

The complete assembly of the global corrector matrix \mathbf{Q}_h is summarized in Algorithm 1. Note that the global corrector matrix \mathbf{Q}_h and the global stiffness matrix \mathbf{A}_h are the only relevant matrices that need to be stored at this point. All other matrices are no longer required (or can be recomputed cheaply).

Algorithm 1: Computation of global corrector matrix \mathbf{Q}_h .

```

1 data:  $N_h, N_H, N_{\mathcal{T}_H}$ ,           # fine space / coarse space / coarse grid size
2        $\mathbf{A}_{dc}, \mathbf{M}_{dc}$ ,           # element stiffness / mass matrix, cf. Def. 3.3
3        $\mathbf{B}^H$ ,                        #  $N_H \times N_H$  boundary correction matrix
4        $\mathbf{P}_h$                         #  $N_H \times N_h$  projection matrix in (14)
5        $\mathbf{R}_\ell^h, \mathbf{R}_\ell^H, \mathbf{T}_\ell^H$  # local restrictions in (11),(12),(15)
6        $\boldsymbol{\sigma}$ ,                  # per elem. to conforming map in (9)
7 def computeCorrections:
8     compute  $\mathbf{A}_h = \boldsymbol{\sigma} \mathbf{A}_{dc} \boldsymbol{\sigma}^\top$            #  $N_h \times N_h$  Stiffness matrix
9     compute  $\mathbf{M}_h = \boldsymbol{\sigma} \mathbf{M}_{dc} \boldsymbol{\sigma}^\top$            #  $N_h \times N_h$  Mass matrix
10    compute  $\mathbf{C}_h := \mathbf{P}_h \mathbf{M}_h$                         #  $N_H \times N_h$  constraints matrix
11     $\mathbf{Q}_h := \text{Matrix}(N_H, N_h)$                         #  $N_H \times N_h$  Corrector matrix
12    ## -- foreach  $K_\ell \in \mathcal{T}_H$  --
13    for  $0 \leq \ell < N_{\mathcal{T}_H}$ :
14         $N_{\ell,h} := \text{rows}(\mathbf{R}_\ell^h)$                     # local fine space size
15         $N_{\ell,H} := \text{rows}(\mathbf{R}_\ell^H)$                     # local coarse space size
16         $\mathbf{A}_\ell := \mathbf{R}_\ell^h \mathbf{A}_h \mathbf{R}_\ell^{h\top}$                 #  $N_{\ell,h} \times N_{\ell,h}$  Stiffness matrix
17         $\mathbf{C}_\ell := \mathbf{R}_\ell^H \mathbf{C}_h \mathbf{R}_\ell^{h\top}$                 #  $N_{\ell,H} \times N_{\ell,h}$  constraints matrix
18         $\mathbf{r}_\ell := -\mathbf{T}_\ell^H \mathbf{B}^H \mathbf{P}_h (\boldsymbol{\sigma}_\ell \mathbf{A}_t \boldsymbol{\sigma}_\ell^\top \mathbf{R}_\ell^{h\top})$  #  $c_d \times N_{\ell,h}$  load vec. matrix
19        ## -- compute inverse operator --
20         $\mathbf{A}_\ell^{\text{inv}} := \mathbf{A}_\ell^{-1}$                         # e.g. using sparse LU
21        ## -- precomputations related to the operator --
22         $\mathbf{Y}_\ell := \text{Matrix}(N_{\ell,H}, N_{\ell,h})$                 #  $N_{\ell,H} \times N_{\ell,h}$  matrix
23        for  $0 \leq m < N_{\ell,H}$ :
24             $\mathbf{Y}_\ell[m] := \mathbf{A}_\ell^{\text{inv}} (\mathbf{C}_\ell^\top[m])$ 
25        ## -- compute inverse Schur complement --
26         $\mathbf{S}_\ell^{\text{inv}} := (\mathbf{C}_\ell \mathbf{Y}_\ell)^{-1}$                 #  $N_{\ell,H} \times N_{\ell,H}$  matrix
27        ## -- compute correction for each coarse space function --
28        ## -- which has a support on  $K_\ell$  --
29         $c_d := \text{rows}(\mathbf{T}_\ell^H)$ 
30         $\mathbf{w}_\ell := \text{Matrix}(c_d, N_{\ell,h})$                 #  $c_d \times N_{\ell,h}$  matrix
31        for  $0 \leq i < c_d$ :
32            ## compute  $\mathbf{w}_\ell[i] = \mathbf{A}_\ell^{-1} \mathbf{r}_\ell[i] - (\mathbf{A}_\ell^{-1} \mathbf{C}_\ell^\top) \mathbf{S}_\ell^{-1} (\mathbf{C}_\ell \mathbf{A}_\ell^{-1}) \mathbf{r}_\ell[i]$ 
33            compute  $\mathbf{q}_i = \mathbf{A}_\ell^{\text{inv}} (\mathbf{r}_\ell[i])$                 #  $\mathbf{q}_i := \mathbf{A}_\ell^{-1} \mathbf{r}_\ell[i]$ 
34            compute  $\lambda_i = \mathbf{S}_\ell^{\text{inv}} (\mathbf{C}_\ell \mathbf{q}_i)$                 #  $\lambda_\ell[i] := \mathbf{S}_\ell^{-1} \mathbf{C}_\ell \mathbf{q}_\ell[i]$ 
35            compute  $\mathbf{w}_\ell[i] = \mathbf{q}_i - \mathbf{Y}_\ell \lambda_i$                 #  $\mathbf{w}_\ell[i] := \mathbf{q}_\ell[i] - \mathbf{Y}_\ell \lambda_\ell[i]$ 
36        ## update correction
37         $\mathbf{Q}_h := \mathbf{Q}_h + (\mathbf{T}_\ell^H)^\top \mathbf{w}_\ell \mathbf{R}_\ell^h$ 
38    return  $\mathbf{Q}_h$ 

```

4 The LOD for linear elliptic problems

We are prepared to state the first full example for an application of the LOD. Given $f \in L^2(\Omega)$, we seek the weak solution of

$$\begin{aligned} -\nabla \cdot \kappa \nabla u &= f & \text{in } \Omega, \\ u &= 0 & \text{on } \Gamma_D, \\ \kappa \nabla u \cdot n &= 0 & \text{on } \Gamma_N, \end{aligned}$$

i.e., we seek $u \in H_{\Gamma_D}^1(\Omega)$ that satisfies

$$\mathcal{A}(u, v) = \int_{\Omega} \kappa \nabla u \cdot \nabla v = \int_{\Omega} f v =: \mathcal{F}(v) \quad \text{for all } v \in H_{\Gamma_D}^1(\Omega). \quad (22)$$

4.1 Method and convergence results

With the definitions from Section 2, we can state the Local Orthogonal Decomposition method (LOD) for model problem (7.1).

Definition 4.1 (LOD approximation for problem (22)). Recall Definition 2.4 for a given localization parameter $k \in \mathbb{N}$. If $u_H \in V_{H, \Gamma_D}$ solves

$$\int_{\Omega} \kappa \nabla (u_H + \mathcal{Q}_h(u_H)) \cdot \nabla (\Phi_H + \mathcal{Q}_h(\Phi_H)) = \int_{\Omega} f (\Phi_H + \mathcal{Q}_h(\Phi_H)) \quad \text{for all } \Phi_H \in V_{H, \Gamma_D} \quad (23)$$

the final LOD approximation is given by $u_{\text{LOD}} = u_H + \mathcal{Q}_h(u_H)$.

The Galerkin solution $u_h \in V_{h, \Gamma_D}$ which satisfies

$$\mathcal{A}(u_h, v) = \mathcal{F}(v) \quad \text{for all } v \in V_{h, \Gamma_D} \quad (24)$$

can be considered as a reference solution in the sense that $u_{\text{LOD}} = u_H + \mathcal{Q}_h(u_H)$ is constructed to approximate u_h with a desired accuracy of at least $\mathcal{O}(H)$. This approximation quality can be quantified:

Theorem 4.2 (A priori error estimate). *Assume that the localization parameter fulfills $k \gtrsim m |\log(H)|$ for some $m \in \mathbb{N}$. Then, there exists a positive constant C that depends on the space dimension d , on Ω , γ_{\min} , γ_{\max} and interior angles of the partitions, but not on the mesh sizes H and h , such that*

$$\begin{aligned} \|u_h - (u_H + \mathcal{Q}_h(u_H))\|_{L^2(\Omega)} &\leq C(H + H^{rm})^2 \quad \text{and} \\ \|u_h - (u_H + \mathcal{Q}_h(u_H))\|_{H^1(\Omega)} + \|u_h - u_H\|_{L^2(\Omega)} &\leq C(H + H^{rm}), \end{aligned}$$

for some constant $r > 0$ that depends linearly on the square root of the contrast.

The theorem was proved in [1, 2, 3]. Practically, numerical experiments indicate that the choice $m \in \{1, 2, 3\}$ typically yields good results even for high contrast cases [14]. We refer to the numerical experiments in [3, 1].

4.2 Assembly and solution of the global problem

4.2.1 Formal description

Assume that all local problems are solved (i.e. solved for every $K_\ell \in \mathcal{T}_H$ and every coarse basis function $\Phi_{p_i(\ell)}$ with support on K_ℓ , where $0 \leq i < c_d$) so that $\mathcal{Q}_h^{K_\ell}(\Phi_{p_i(\ell)})$ is available and we can write

$$\mathcal{Q}_h(\Phi_H) = \sum_{K_\ell \in \mathcal{T}_H} \sum_{i=0}^{c_d-1} \Phi_H(Z_{p_i(\ell)}) \mathcal{Q}_h^{K_\ell}(\Phi_{p_i(\ell)})$$

for any $\Phi_H \in V_{H,\Gamma_D}$. Consequently we can assemble the (global) LOD stiffness matrix $\mathbf{A}_H^{\text{LOD}} \in \mathbb{R}^{N_H \times N_H}$ that is given by the entries

$$\mathbf{A}_H^{\text{LOD}}[m][n] := \begin{cases} \mathcal{A}(\Phi_n + \mathcal{Q}_h(\Phi_n), \Phi_m + \mathcal{Q}_h(\Phi_m)) & \text{for } Z_m, Z_n \in \mathcal{N}_H \setminus \Gamma_D, \\ 0 & \text{else} \end{cases}$$

and LOD load vector $\mathbf{f}_H \in \mathbb{R}^{N_H}$ given by

$$\mathbf{f}_H[m] := \begin{cases} (f, \Phi_m + \mathcal{Q}_h(\Phi_m))_{L^2(\Omega)} & \text{for } Z_m \in \mathcal{N}_H \setminus \Gamma_D, \\ 0 & \text{else.} \end{cases}$$

With that, the algebraic version of (23) hence reads: find $\mathbf{u}_H^{\text{LOD}} \in \mathbb{R}^{N_H}$ with

$$\mathbf{A}_H^{\text{LOD}} \mathbf{u}_H^{\text{LOD}} = \mathbf{f}_H.$$

Once this is solved, the final LOD approximation is given by

$$u_{\text{LOD}} = \sum_{m=0}^{N_H-1} \mathbf{u}_H^{\text{LOD}}[m] (\Phi_m + \mathcal{Q}_h(\Phi_m)) \in V_{h,\Gamma_D}.$$

The condition number of the LOD system matrix $\mathbf{A}_H^{\text{LOD}}$ is of order $1/H^2$, i.e. of the same order as the condition number for standard finite elements on the coarse scale. A corresponding estimate for the condition number of $\mathbf{A}_H^{\text{LOD}}$ is for instance given in [20, Lemma 1]. Depending on the coarse space size one can employ a direct solver or a standard iterative solver, like the multigrid method.

4.2.2 Algorithmic realization

Next we describe an efficient algorithmic realization of how to assemble and solve the global problem. When all local problems are solved, the global corrector matrix $\mathbf{Q}_h \in \mathbb{R}^{N_H \times N_h}$ is available (cf. Definition 3.15). From \mathbf{Q}_h , \mathbf{B}^H and the projection matrix \mathbf{P}_h we hence obtain the global LOD system matrix by matrix multiplication

$$\mathbf{A}_H^{\text{LOD}} = \mathbf{B}^H (\mathbf{P}_h + \mathbf{Q}_h) \mathbf{A}_h (\mathbf{P}_h + \mathbf{Q}_h)^\top \mathbf{B}^H.$$

Similarly, we get the load vector \mathbf{f}_H by

$$\mathbf{f}_H = \mathbf{B}^H (\mathbf{P}_h + \mathbf{Q}_h) \mathbf{f}_h,$$

where $\mathbf{f}_h \in \mathbb{R}^{N_h}$ denotes the classical FEM load vector with entries $\mathbf{f}_h[i] = (f, \phi_i)_{L^2(\Omega)}$ for $0 \leq i < N_h$. Now we can solve for $\mathbf{u}_H^{\text{LOD}} \in \mathbb{R}^{N_H}$ with

$$\mathbf{A}_H^{\text{LOD}} \mathbf{u}_H^{\text{LOD}} = \mathbf{f}_H$$

and obtain the final coefficient vector $\mathbf{u}_h^{\text{LOD}} \in \mathbb{R}^{N_h}$ of our LOD approximation by

$$\mathbf{u}_h^{\text{LOD}} := (\mathbf{P}_h + \mathbf{Q}_h)^\top \mathbf{u}_H^{\text{LOD}}.$$

The procedure is summarized in Algorithm 2.

Algorithm 2: Computation of the final LOD approximation $\mathbf{u}_h^{\text{LOD}}$.

```

1 data:  $\mathcal{T}_h, \mathcal{T}_H,$                                 # fine / coarse mesh
2          $N_h, N_H,$                                 # fine / coarse space size
3          $\mathbf{P}_h$                                     #  $N_H \times N_h$  interp. matrix in (14)
4          $\mathbf{A}_h$                                     #  $N_h \times N_h$  stiffness matrix
5          $\mathbf{f}_h$                                     #  $N_h \times 1$  (fine) load vector
6          $\mathbf{B}^H$                                     #  $N_H \times N_H$  boundary correction matrix
7          $\mathbf{Q}_h$                                     #  $N_H \times N_h$  global corrector matrix
8 def solveLODSystem:
9     compute  $\mathbf{A}_H^{\text{LOD}} = \mathbf{B}^H(\mathbf{P}_h + \mathbf{Q}_h)\mathbf{A}_h(\mathbf{P}_h + \mathbf{Q}_h)^\top \mathbf{B}^H$     #  $N_H \times N_H$  LOD sys.mat.
10    compute  $\mathbf{f}_H = \mathbf{B}^H(\mathbf{P}_h + \mathbf{Q}_h)\mathbf{f}_h$                                 #  $N_H \times 1$  LOD rhs.vec.
11    solve  $\mathbf{A}_H^{\text{LOD}} \mathbf{u}_H^{\text{LOD}} = \mathbf{f}_H$                                     # solve LOD system
12    compute  $\mathbf{u}_h^{\text{LOD}} := (\mathbf{P}_h + \mathbf{Q}_h)^\top \mathbf{u}_H^{\text{LOD}}$                         #  $N_h \times 1$  LOD solution
13    return  $\mathbf{u}_h^{\text{LOD}}$ 

```

Remark 4.3. Typically it is possible to replace the right hand side in (23) by $(f, \Phi_H)_{L^2(\Omega)}$ without a significant loss in accuracy. In particular if f is a slow variable. Consequently, we do not require the correctors any longer to compute the load vector in the global LOD system. This can turn out to be an immense computational advantage if the LOD system has to be solved for several source terms f . In this case, we can fully reuse $\mathbf{A}_H^{\text{LOD}}$ and quickly assemble \mathbf{f}_H with entries $\mathbf{f}_H[i] = (f, \Phi_i)_{L^2(\Omega)}$ (only involving coarse basis functions).

Remark 4.4. In the case that the coefficient κ has certain structural properties (such as periodicity), it might be possible to only assemble some of the local correctors \mathcal{Q}_h^K and reuse them on the different location in Ω . This is possible if a corrector can be expressed as a rotation and translation of another corrector. With that the computational complexity can be decreased significantly. This has been exploited in the context of acoustic scattering in [28].

4.3 Petrov-Galerkin version of the method

In some cases it can happen that the fine space V_h is so large that the storing of the full system matrix \mathbf{A}_h (respectively the storing of the corrector matrix \mathbf{Q}_h) becomes too memory demanding. In such cases we cannot afford the multiplication of $N_h \times N_h$ matrices as frequently done in Algorithm 1 and 2. To overcome the issue that the size of V_h exceeds the computational resources, a Petrov-Galerkin (PG) formulation of the LOD can be used. This method allows an on-the-fly assembling of the LOD system matrix $\mathbf{A}_H^{\text{LOD}}$ on the expense that we lose the symmetry. Let us start with describing the Petrov-Galerkin LOD from the analytical point of view.

4.3.1 Description and properties of the PG-LOD

We use the notation introduced earlier in this section.

Definition 4.5 (PG LOD approximation). Let $k \in \mathbb{N}$ be fixed and let $\mathcal{Q}_h : V_{H,\Gamma_D} \rightarrow V_{h,\Gamma_D}$ denote the corresponding correction operator as in Definition 2.4. If $u_H^{\text{PG}} \in V_{H,\Gamma_D}$ solves

$$\int_{\Omega} \kappa \nabla(u_H^{\text{PG}} + \mathcal{Q}_h(u_H^{\text{PG}})) \cdot \nabla \Phi_H = \int_{\Omega} f \Phi_H \quad \text{for all } \Phi_H \in V_{H,\Gamma_D}, \quad (25)$$

the final Petrov-Galerkin LOD approximation is given by $u^{\text{PG-LOD}} := u_H^{\text{PG}} + \mathcal{Q}_h(u_H^{\text{PG}})$.

Obviously, the standard formulation of the LOD only differs from the PG formulation by the choice of test functions in (25). In particular, the solving of the local corrector problems is identical for both methods. From the analytical point of view, the change of test functions in (23) does not have a crucial influence. We still have well-posedness of the PG-LOD solution and the obtained convergence rates are the same as for the original method. We summarize the corresponding main result in the following theorem, which is proved in [25].

Theorem 4.6 (A priori error estimate for the PG-LOD). *Assume that $k \gtrsim m|\log(H)|$ for some $m \in \mathbb{N}$. Furthermore, let the positive constants C and r be as in Theorem 4.2. Then the left side of (25) represents a coercive bilinear form on V_{H,Γ_D} , i.e.*

$$\int_{\Omega} \kappa \nabla(\Phi_H + \mathcal{Q}_h(\Phi_H)) \cdot \nabla \Phi_H \geq C(\alpha - CH^{rm}) \|\Phi_H\|_{H^1(\Omega)}^2. \quad (26)$$

Consequently, problem (25) is well-posed and the PG-LOD approximation fulfills the same error estimates as the standard LOD approximation, i.e. we have

$$\begin{aligned} \|u_h - u_H^{\text{PG}} - \mathcal{Q}_h(u_H^{\text{PG}})\|_{L^2(\Omega)} &\leq C(H + H^{rm})^2 \quad \text{and} \\ \|u_h - u_H^{\text{PG}} - \mathcal{Q}_h(u_H^{\text{PG}})\|_{H^1(\Omega)} + \|u_h - u_H^{\text{PG}}\|_{L^2(\Omega)} &\leq C(H + H^{rm}). \end{aligned}$$

For even sharper results in L^2 we refer to [16].

Remark 4.7 (Relevance of L^2 -approximations). Theorem 4.6 contains L^2 - and H^1 error estimate for the full RG-LOD approximation $u_H^{\text{PG}} + \mathcal{Q}_h(u_H^{\text{PG}})$. However, to compute it from the coarse part u_H^{PG} , we need to know the operator \mathcal{Q}_h . But recall that the algebraic version of \mathcal{Q}_h is represented by the corrector matrix $\mathbf{Q}_h \in \mathbb{R}^{N_h \times N_h}$, which is of the same size (and even less sparse) than the global stiffness matrix $\mathbf{A}_h \in \mathbb{R}^{N_h \times N_h}$. So if we do not have the capacities to store \mathbf{A}_h , neither do we have the capacities to store \mathbf{Q}_h . Consequently, even though u_H^{PG} might be available, $\mathcal{Q}_h(u_H^{\text{PG}})$ is typically not. Hence, our final approximation u_H^{PG} is only an L^2 -approximation, instead of a full H^1 -approximation (as e.g. $u_H + \mathcal{Q}_h(u_H)$). Hence, the relevant estimate that remains from Theorem 4.6 is the L^2 -error estimate

$$\|u_h - u_H^{\text{PG}}\|_{L^2(\Omega)} \leq C(H + H^{rm}).$$

However, note that given the RG-LOD solution u_H we could go easily back to the local problems and recompute with known right hand side to form the full fine scale solution without having to store the correctors $\mathcal{Q}_h(u_H)$. This strategy allows to obtain H^1 -approximations through local post-processing.

4.3.2 Computational advantages and disadvantages

Let us now describe the computational advantages and disadvantages of the Petrov-Galerkin formulation. We start with the advantages to demonstrate how the PG formulation overcomes the capacity issues.

Advantages. The basic advantage of the PG-LOD is that matrices of size $N_h \times N_h$ have to be handled at no point. Operations either involve $N_{\ell,h} \times N_{\ell,h}$ -matrices ($N_{\ell,h}$ is the number of fine nodes in the patch $U_k(K_\ell)$) or they involve $N_H \times N_H$ -matrices (where N_H denotes the size of V_H). The reason why this is possible is that no corrector-to-corrector communication is required for the PG-LOD. For instance, in order to assemble the system matrix that is associated with standard LOD (cf. (23)), we need to compute entries such as

$$\int_{\Omega} \kappa \nabla(\Phi_i + \mathcal{Q}_h(\Phi_i)) \cdot \nabla(\Phi_j + \mathcal{Q}_h(\Phi_j))$$

for two coarse basis functions Φ_i and Φ_j . It is impossible to compute this entry without knowing both $\mathcal{Q}_h(\Phi_i)$ and $\mathcal{Q}_h(\Phi_j)$ at the same time. Consequently correctors must be stored so that they can communicate with each other. For the PG-LOD, system matrix entries are always of the structure

$$\int_{\Omega} \kappa \nabla (\Phi_i + \mathcal{Q}_h(\Phi_i)) \cdot \nabla \Phi_j = \sum_{\substack{K \in \mathcal{T}_H \\ K \subset \text{supp}(\Phi_i)}} \int_{U_k(K)} \kappa \nabla (\Phi_i + \mathcal{Q}_h^K(\Phi_i)) \cdot \nabla \Phi_j,$$

which can be assembled (respectively updated) after a corrector $\mathcal{Q}_h^K(\Phi_i)$ is computed. If desired, $\mathcal{Q}_h^K(\Phi_i)$ can be immediately deleted after this. Global matrices such as the corrector matrix \mathbf{Q}_h or the stiffness matrix \mathbf{A}_h do neither have to be stored nor explicitly computed. Only a sparse global system matrix $\mathbf{A}_H^{\text{PG-LOD}}$ of size $N_H \times N_H$ is required. Consequently, the storage requirements are significantly lower for the Petrov-Galerkin version.

Trade-offs. In comparison to (23), we observe that the PG-LOD system given by (25) can no more be represented by a symmetric matrix, which formally excludes the usage of certain efficient algebraic solvers that rely on symmetry. However, having a closer look, we see that the method only suffers from a mild loss of symmetry in the sense that the PG-LOD is still symmetric if there is no localization and that the lack of symmetry can be hence quantified by the exponential decay property. A symmetric approximation can be for instance obtained by using $\frac{1}{2} \mathbf{A}_H^{\text{PG-LOD}} + \frac{1}{2} (\mathbf{A}_H^{\text{PG-LOD}})^{\top}$ for the system matrix. The second trade-off is rather subtle. Theorem 4.6 predicts a coercivity constant that can be disturbed by a term of order $\mathcal{O}(H^{rm})$. Even though this seems to be mostly unproblematic for small H , there is formally no guarantee that $(\alpha - CH^{rm})$ is always positive. This can only be guaranteed by a numerical investigation of the eigenvalues. However, we also note that non-positivity has never been observed in numerical experiments. So far it seems that the result (26) is not yet optimal and the coercivity appears to be always fulfilled in practical applications. Another trade-off was already mentioned in Remark 4.7. If we are in a scenario where the PG-LOD is used to decrease the memory demand, then the correctors \mathcal{Q}_h will not be stored. Hence we will not be able to compute $\mathcal{Q}_h(u_H^{\text{PG}})$ from u_H^{PG} and have to be content with an L^2 -approximation of u_h . In many applications this is enough. Especially when considering that u_H^{PG} can be stored with significantly lower costs than the full fine scale approximation $u_H^{\text{PG}} + \mathcal{Q}_h(u_H^{\text{PG}})$. If local fine scale information is required afterwards by a user, it is possible to perform a local "real time" post-processing where the missing fine-scale information is only computed in the relevant region.

4.3.3 Realization

Even though the (algebraic) realization differs only slightly from the realization of the classical LOD, these differences are essential. The local corrector problems are computed in the same way as before, however, instead of storing their solutions in a global corrector matrix, their contributions are directly added to the global PG-LOD system matrix and can be immediately deleted afterwards.

To summarize the basic procedure, let us fix a coarse element $K_{\ell} \in \mathcal{T}_H$ and a corresponding coarse layer patch $\mathcal{U}_{\ell} = U_k(K_{\ell})$. The following step has to be repeated for every K_{ℓ} . For every coarse basis function $\Phi_{p_i(\ell)}$ (with $0 \leq i < c_d$ and only if $Z_{p_i(\ell)} \notin \Gamma_D$) we solve for the local corrector for $\mathcal{Q}_h^{K_{\ell}}(\Phi_{p_i(\ell)})$ according to (17). After $\mathcal{Q}_h^{K_{\ell}}(\Phi_{p_i(\ell)})$ is computed, we update the system matrix $\mathbf{A}_H^{\text{PG-LOD}}$. To do that we visit every coarse basis function $\Phi_m \in V_{H, \Gamma_D}$ with $\text{supp}(\Phi_m) \cap \mathcal{U}_{\ell} \neq \emptyset$. For each such Φ_m we can make the update

$$\mathbf{A}_H^{\text{PG-LOD}}[m][p_i(\ell)] := \mathbf{A}_H^{\text{PG-LOD}}[m][p_i(\ell)] + \int_{\mathcal{U}_{\ell}} \kappa \nabla (\Phi_{p_i(\ell)} + \mathcal{Q}_h^{K_{\ell}}(\Phi_{p_i(\ell)})) \cdot \nabla \Phi_m.$$

When all loops have terminated we can incorporate the homogenous boundary condition by multiplying $\mathbf{A}_H^{\text{PG-LOD}}$ at both sides with the boundary correction matrix \mathbf{B}^H , i.e. $\mathbf{A}_H^{\text{PG-LOD}} := \mathbf{B}^H \mathbf{A}_H^{\text{PG-LOD}} \mathbf{B}^H$. In total, we computed the correct PG-LOD system matrix with entries

$$\begin{aligned} \mathbf{A}_H^{\text{PG-LOD}}[m][n] &= \int_{\Omega} \kappa \nabla (\Phi_n + \mathcal{Q}_h(\Phi_n)) \cdot \nabla \Phi_m \\ &= \sum_{\substack{K_\ell \in \mathcal{T}_H \\ K_\ell \subset \text{supp}(\Phi_n)}} \int_{\mathcal{U}_\ell} \kappa \nabla (\Phi_n + \mathcal{Q}_h^{K_\ell}(\Phi_n)) \cdot \nabla \Phi_m. \end{aligned}$$

Observe that we could generate a new local fine mesh for every \mathcal{U}_ℓ . Basically, there is no need for a global fine mesh \mathcal{T}_h . The remaining procedure is straightforward. Since the right hand side of (25) only involves standard coarse functions we can set $\mathbf{f}_H^{\text{PG-LOD}} := \mathbf{B}^H \mathbf{f}_H$, where $\mathbf{f}_H \in \mathbb{R}^{N_H}$ denotes the standard coarse load vector with entries $\mathbf{f}_H[m] = \int_{\Omega} f \Phi_m$. Consequently it only remains to solve for $\mathbf{u}_H^{\text{PG-LOD}} \in \mathbb{R}^{N_H}$ with

$$\mathbf{A}_H^{\text{PG-LOD}} \mathbf{u}_H^{\text{PG-LOD}} = \mathbf{f}_H^{\text{PG-LOD}}.$$

The final approximation is given by

$$u_H^{\text{PG}} = \sum_{m=0}^{N_H-1} \mathbf{u}_H^{\text{PG-LOD}}[m] \Phi_m.$$

Accordingly modified algorithms can be formulated analogously to the algorithms presented for the standard (symmetric) LOD.

5 The treatment of rough boundary data and sources

In this section we discuss how we can incorporate nonhomogeneous boundary conditions in the LOD approximations and how we can treat regions in which the source term f becomes close to singular.

5.1 Model problem and discretization

We consider the following problem with mixed Dirichlet and Neumann boundary conditions. Find u with

$$\begin{aligned} -\nabla \cdot \kappa \nabla u &= f && \text{in } \Omega, \\ u &= g && \text{on } \Gamma_D, \\ \kappa \nabla u \cdot n &= q && \text{on } \Gamma_N. \end{aligned}$$

In addition to the previous assumptions on Ω and κ we assume that the Dirichlet boundary values fulfill $g \in H^{\frac{1}{2}}(\Gamma_D)$ and that the Neumann boundary values fulfill $q \in L^2(\Gamma_N)$. The weak formulation of problem (5.1) reads: find $u \in H^1(\Omega)$, with $T_D(u) = g$, such that

$$\int_{\Omega} \kappa \nabla u \cdot \nabla v = \int_{\Omega} f v + \int_{\Gamma_N} q v \quad \text{for all } v \in H_{\Gamma_D}^1(\Omega).$$

Assume that g is sufficiently regular so that point evaluations are possible. Then we can define $g_H \in V_H$ as the function that is uniquely determined by the nodal values $g_H(Z) = g(Z)$ for all $Z \in \mathcal{N}_H \cap \Gamma_D$ and $g_H(Z) = 0$ for all $Z \in \mathcal{N}_H \setminus \Gamma_D$. Using this, we define the (fine scale) Dirichlet extension $g_h \in V_h$ uniquely by the nodal values $g_h(z) = g(z)$ for all $z \in \mathcal{N}_h \cap \Gamma_D$ and

$g_h(z) = g_H(z)$ for all $z \in \mathcal{N}_h \setminus \Gamma_D$. With this, we avoid degeneracy of g_h for h tending to zero. Note that the extension g_h needs to be explicitly constructed. The reference problem reads: find $U_h \in V_{h,\Gamma_D}$ with

$$\mathcal{A}(U_h, v_h) = \int_{\Omega} f v_h - \int_{\Omega} \kappa \nabla g_h \cdot \nabla v_h + \int_{\Gamma_N} q v_h \quad \text{for all } v_h \in V_{h,\Gamma_D}. \quad (27)$$

The final fine scale approximation is then given by $u_h := U_h + g_h \in V_h$. Now observe that problem (27) is basically of the same structure as the homogenous problem (24). This suggest to apply the same methodology as before. Unfortunately, the correctors introduced in Definition 2.4 might not be sufficient to construct accurate approximations, if e.g. the Dirichlet boundary condition is highly oscillatory. The slight difference that the right hand side is no longer purely represented by an L^2 -function f (but by a less regular functional which is only in the dual space of V_{h,Γ_D}) makes it necessary to introduce additional correctors to preserve the previous convergence rates. We call these new correctors *source correctors*.

5.2 Source correctors

In this section we introduce *source term correctors*. They are defined analogously to the correctors \mathcal{Q}_h . Their purpose is to capture oscillatory effects that are produced by a general source. For that purpose, we split the right hand side of (27) into two parts. One part (we shall denote by \mathcal{F}) that has basically a coarse scale structure and that can be considered as harmless if ignored by the fine grid, and a second part (we shall denote by \mathcal{F}^s) which might have a considerable influence on the oscillations of u_h . Hence we let $\mathcal{F} : H^1(\Omega) \rightarrow \mathbb{R}$ and $\mathcal{F}^s : H^1(\Omega) \rightarrow \mathbb{R}$ be source functionals such that

$$\mathcal{F}(v_h) + \mathcal{F}^s(v_h) = \int_{\Omega} f v_h - \int_{\Omega} \kappa \nabla g_h \cdot \nabla v_h + \int_{\Gamma_N} q v_h \quad \text{for all } v_h \in V_{h,\Gamma_D}.$$

We only wish to introduce additional correctors for the \mathcal{F}^s -contribution. It can incorporate source terms and boundary conditions and we assume that it is of the structure

$$\mathcal{F}^s(v) = \int_{\Omega} \eta_1 v + \kappa \nabla \eta_2 \cdot \nabla v + \int_{\partial\Omega} \eta_3 v$$

with some given $\eta_1 \in L^2(\Omega)$, $\eta_2 \in H^1(\Omega)$ and $\eta \in H^{1/2}(\partial\Omega)$. Typical choices would be

$$\begin{aligned} \mathcal{F}^s(v) &= \int_{\Omega} \kappa \nabla g_h \cdot \nabla v + \int_{\Gamma_N} q v \quad (\text{boundary source}) \text{ or,} \\ \mathcal{F}^s(v) &= \int_{\Omega} f v + \kappa \nabla g_h \cdot \nabla v + \int_{\Gamma_N} q v \quad (\text{total source}). \end{aligned}$$

We define the localization of \mathcal{F}^s to a coarse element $K \in \mathcal{T}_H$ by

$$\mathcal{F}_K^s(v) := \int_K \eta_1 v + \kappa \nabla \eta_2 \cdot \nabla v + \int_{\partial\Omega \cap K} \eta_3 v.$$

With that, we can define local source correctors.

Definition 5.1 (Source term correctors). Let \mathcal{F}^s be fixed according to the previous discussion. For a given positive $k \in \mathbb{N}$ and $K \in \mathcal{T}_H$ we define the local source corrector $\mathcal{Q}_{\mathcal{F}^s,h}^K \in W_h(U_k(K))$ as the solution of

$$\int_{U_k(K)} \kappa \nabla \mathcal{Q}_{\mathcal{F}^s,h}^K \cdot \nabla w_h = -\mathcal{F}_K^s(w_h) \quad \text{for all } w_h \in W_h(U_k(K)). \quad (28)$$

The corresponding global corrector is given by

$$\mathcal{Q}_{\mathcal{F}^s, h} := \sum_{K \in \mathcal{T}_H} \mathcal{Q}_{\mathcal{F}^s, h}^K.$$

Note that it is desirable that \mathcal{F}^s only contains locally supported sources, i.e. that $\mathcal{F}_K^s(w_h) = 0$ for most of the coarse elements K . The more elements with $\mathcal{F}_K^s(w_h) \neq 0$, the more local problems to solve. Therefore \mathcal{F}^s typically only contains boundary terms or parts of f in a small region where f might become close to singular.

5.3 Formulation of the method and error estimates

Using Definition 2.4 and 5.1 we propose the following LOD approximation.

Definition 5.2 (LOD approximation for boundary value problems).

For fixed \mathcal{F}^s and fixed $k \in \mathbb{N}$ the LOD approximation to problem (27) is given by $u_{\text{LOD}} := U_H + \mathcal{Q}_h(U_H) - B_h$, where $U_H \in V_{H, \Gamma_D}$ solves:

$$\begin{aligned} & \int_{\Omega} \kappa \nabla (\text{Id} + \mathcal{Q}_h)(U_H) \cdot \nabla (\text{Id} + \mathcal{Q}_h)(\Phi_H) \\ &= (\mathcal{F} + \mathcal{F}^s)((\text{Id} + \mathcal{Q}_h)(\Phi_H)) + \int_{\Omega} \kappa \nabla \mathcal{Q}_{\mathcal{F}^s, h} \cdot \nabla (\text{Id} + \mathcal{Q}_h)(\Phi_H) \end{aligned} \quad (29)$$

for all $\Phi_H \in V_{H, \Gamma_D}$.

The following a priori error estimate is proved in [3].

Theorem 5.3. Assume that $k \gtrsim m |\log(H)|$ for some $m \in \mathbb{N}$. Furthermore, let

$$\mathcal{F}^s(v) = \int_{\Omega} \kappa \nabla g_h \cdot \nabla v + \int_{\Gamma_N} qv \quad \text{and} \quad \mathcal{F}(v) = \int_{\Omega} f v.$$

Recall that we compute the source corrector $\mathcal{Q}_{\mathcal{F}^s, h}$ only with respect to \mathcal{F}^s . Then the LOD approximation u_{LOD} introduced in Definition 5.2 fulfills the estimates

$$\|u_h - u_{\text{LOD}}\|_{L^2(\Omega)} \leq C(H + H^{rm})^2 \quad \text{and} \quad \|u_h - u_{\text{LOD}}\|_{H^1(\Omega)} \leq C(H + H^{rm}),$$

where C and $r > 0$ are as in Theorem 4.2.

5.4 Algebraic realization

The algebraic realization is straightforward following the ideas presented in Section 4. We only need to solve one additional linear elliptic problem more for each coarse element $K_{\ell} \in \mathcal{T}_H$ with $\mathcal{F}_{K_{\ell}}^s \neq 0$ and we need to assemble an additional vector that stores the entries $\int_{\Omega} \kappa \nabla \mathcal{Q}_{\mathcal{F}^s, h} \cdot \nabla (\text{Id} + \mathcal{Q}_h)(\Phi_j)$.

5.4.1 Solving of the additional local problem

Let us fix a coarse element $K_{\ell} \in \mathcal{T}_H$. First recall that for every coarse basis function $\Phi_{p_i(\ell)}$ ($0 \leq i < c_d$, see Section 3.5) we need to solve for $w_{\ell, i} := \mathcal{Q}_h^{K_{\ell}}(\Phi_{p_i(\ell)}) \in V_{h, \ell}$ with $I_H(w_{\ell, i}) = 0$ and

$$\int_{\mathcal{U}_{\ell}} \kappa \nabla w_{\ell, i} \cdot \nabla w_h = - \int_{K_{\ell}} \kappa \nabla \Phi_{p_i(\ell)} \cdot \nabla w_h$$

for all $w_h \in V_{h,\ell}$ with $I_H(w_h) = 0$. Now, if $\mathcal{F}_{K_\ell}^s \neq 0$ we also need to solve for $\hat{w}_\ell := \mathcal{Q}_{\mathcal{F}^s,h}^{K_\ell} \in V_{h,\ell}$ with $I_H(\hat{w}_\ell) = 0$ and

$$\int_{\mathcal{U}_\ell} \kappa \nabla \hat{w}_\ell \cdot \nabla w_h = -\mathcal{F}_K^s(w_h) \quad (30)$$

for all $w_h \in V_{h,\ell}$ with $I_H(w_h) = 0$. However, both problems only differ in their source terms, where the inverse of the Schur-complement matrix \mathbf{S}_ℓ^{-1} is already precomputed for solving the original corrector problems. Let us introduce a notation for the algebraic version of (30).

Definition 5.4 (Algebraic formulation of (30)). Let the load vector $\hat{\mathbf{r}}_\ell \in \mathbb{R}^{N_{\ell,h}}$ be given by the entries $\hat{\mathbf{r}}_\ell[j] := -\mathcal{F}_K^s(\phi_{\ell,j})$ for $0 \leq j < N_{\ell,h}$. The algebraic version of problem (30) is the following saddle point problem. Find the tuple $(\hat{\mathbf{w}}_\ell, \hat{\boldsymbol{\lambda}}_\ell) \in \mathbb{R}^{N_{\ell,h}} \times \mathbb{R}^{N_{\ell,H}}$ with

$$\begin{aligned} \mathbf{A}_\ell \hat{\mathbf{w}}_\ell + \mathbf{C}_\ell^\top \hat{\boldsymbol{\lambda}}_\ell &= \hat{\mathbf{r}}_\ell \\ \mathbf{C}_\ell \hat{\mathbf{w}}_\ell &= 0. \end{aligned} \quad (31)$$

Here, $\hat{\mathbf{w}}_\ell$ is the coefficient vector for the solution $\hat{w}_\ell = \mathcal{Q}_{\mathcal{F}^s,h}^{K_\ell}$.

Exploiting the notation from Section 3.6, we obtain $\hat{\mathbf{w}}_\ell$ in three steps:

1. Solve for $\hat{\mathbf{q}}_\ell \in \mathbb{R}^{N_{\ell,h}}$ with $\mathbf{A}_\ell \hat{\mathbf{q}}_\ell = \hat{\mathbf{r}}_\ell$.
2. Since \mathbf{S}_ℓ^{-1} is precomputed, we obtain $\hat{\boldsymbol{\lambda}}_\ell$ from $\hat{\mathbf{q}}_\ell$ via $\hat{\boldsymbol{\lambda}}_\ell = \mathbf{S}_\ell^{-1} \mathbf{C}_\ell \hat{\mathbf{q}}_\ell$.
3. Using the precomputed matrix $\mathbf{Y}_\ell = \mathbf{A}_\ell^{-1} \mathbf{C}_\ell^\top$ we obtain $\hat{\mathbf{w}}_\ell = \hat{\mathbf{q}}_\ell - \mathbf{Y}_\ell \hat{\boldsymbol{\lambda}}_\ell$.

Observe that this procedure only involves one single (low dimensional) system of equations to solve.

5.4.2 Assembly and solution of the global problem

The procedure is basically analogous to the case of a homogenous boundary condition. The global corrector matrix \mathbf{Q}_h is assembled identically as before. The same holds for the interpolation matrix \mathbf{P}_h . With that, we obtain the LOD stiffness matrix $\mathbf{A}_H^{\text{LOD}} \in \mathbb{R}^{N_H \times N_H}$ (associated with the left hand side of (29)) by

$$\mathbf{A}_H^{\text{LOD}} = \mathbf{B}^H (\mathbf{P}_h + \mathbf{Q}_h) \mathbf{A}_h (\mathbf{P}_h + \mathbf{Q}_h)^\top \mathbf{B}^H.$$

In order to assemble the LOD load vector $\mathbf{f}_H \in \mathbb{R}^{N_H}$, we first need to assemble the vector $\hat{\mathbf{f}}_h \in \mathbb{R}^{N_h}$ that stores the information gained from source correctors. For a given fine basis function ϕ_j the corresponding entry of $\hat{\mathbf{f}}_h$ is given by

$$\begin{aligned} \hat{\mathbf{f}}_h[j] &:= \int_{\Omega} \kappa \nabla \mathcal{Q}_{\mathcal{F}^s,h} \cdot \nabla \phi_j = \sum_{\ell=0}^{N_H-1} \int_{\mathcal{U}_\ell} \kappa \nabla \mathcal{Q}_{\mathcal{F}^s,h}^{K_\ell} \cdot \nabla \phi_j \\ &= \sum_{\ell=0}^{N_H-1} \sum_{i=0}^{N_{\ell,h}-1} \hat{\mathbf{w}}_\ell[i] \int_{\mathcal{U}_\ell} \kappa \nabla \phi_{\ell,i} \cdot \nabla \phi_j. \end{aligned}$$

Consequently we obtain $\hat{\mathbf{f}}_h$ by matrix multiplication and summation as

$$\hat{\mathbf{f}}_h = -\mathbf{A}_h \hat{\mathbf{w}}_h, \quad \text{where } \hat{\mathbf{w}}_h := - \sum_{\ell=0}^{N_H-1} \mathbf{R}_\ell^{h^\top} \hat{\mathbf{w}}_\ell.$$

The standard load vector $\mathbf{f}_h \in \mathbb{R}^{N_h}$ associated with a classical fine element method on the fine grid \mathcal{T}_h is given by

$$\mathbf{f}_h[i] := (\mathcal{F} + \mathcal{F}^s)(\phi_i).$$

In total, we obtain the LOD load vector as $\mathbf{f}_H = \mathbf{B}^H(\mathbf{P}_h + \mathbf{Q}_h)(\mathbf{f}_h + \hat{\mathbf{f}}_h) \in \mathbb{R}^{N_H}$ (i.e. the vector associated with the right hand side of (29)). Using this, we can solve for $\mathbf{U}_H^{\text{LOD}} \in \mathbb{R}^{N_H}$ with

$$\mathbf{A}_H^{\text{LOD}} \mathbf{U}_H^{\text{LOD}} = \mathbf{B}^H(\mathbf{P}_h + \mathbf{Q}_h)(\mathbf{f}_h - \mathbf{A}_h \hat{\mathbf{w}}_h)$$

and obtain the final solution vector $\mathbf{u}_h^{\text{LOD}} \in \mathbb{R}^{N_h}$ of our LOD approximation by

$$\mathbf{u}_h^{\text{LOD}} := (\mathbf{P}_h + \mathbf{Q}_h)^\top \mathbf{u}_H^{\text{LOD}} + \hat{\mathbf{w}}_h,$$

respectively $u_{\text{LOD}} := \sum_{i=0}^{N_h-1} \mathbf{u}_h^{\text{LOD}}[i] \phi_i$. With these changes, Algorithm 1 and 2 can be modified in an obvious way.

6 Linear elliptic eigenvalue problems

In this section we describe how the LOD can be applied to solve eigenvalue problems. Eigenvalue problems have a distinct status since the LOD is particularly efficient for tackling them, even if the diffusion coefficient has no multiscale character. We consider the following linear eigenvalue problem with a homogenous Dirichlet boundary condition, i.e. $\Gamma_D = \partial\Omega$, $g = 0$ and $\Gamma_N = \emptyset$. We seek tuples $(\lambda^{(n)}, u^{(n)}) \in \mathbb{R} \times H_0^1(\Omega)$ with

$$\mathcal{A}(u^{(n)}, v) = \lambda^{(n)} \int_{\Omega} u^{(n)} v \quad \text{for all } v \in H_0^1(\Omega). \quad (32)$$

We assume that the eigenvalues are ordered and that the eigenvalues are L^2 -normalized, i.e. we have $\lambda^{(n)} \leq \lambda^{(n+1)}$ and $\|u^{(n)}\|_{L^2(\Omega)} = 1$ for all $n \in \mathbb{N}$. Note that we always have $\lambda^{(0)} > 0$. The corresponding fine-scale reference solutions are given by the tuples $(\lambda_h^{(n)}, u_h^{(n)}) \in \mathbb{R} \times V_{h,\Gamma_D}$ with $\|u_h^{(n)}\|_{L^2(\Omega)} = 1$ and

$$\mathcal{A}(u_h^{(n)}, v_h) = \lambda_h^{(n)} \int_{\Omega} u_h^{(n)} v_h \quad \text{for all } v_h \in V_{h,\Gamma_D} \quad (33)$$

and where $0 < \lambda_h^{(n)} \leq \lambda_h^{(n+1)}$.

6.1 LOD approximation of a linear eigenvalue problem

The LOD method for approximating the eigenpair $(\lambda^{(n)}, u^{(n)})$ of (32) can be implemented in a straightforward way. After solving all local problems as described in Section 3 we can assemble the LOD system matrix $\mathbf{A}_H^{\text{LOD}}$ as described in Section 4.2. This has to be done only once and can be seen as a one-time preprocessing step. The local solutions and the LOD system matrix are stored and can be reused for all eigenvalues that we are interested in. We define the LOD approximation as follows.

Definition 6.1 (Eigenvalue LOD approximation). Let $k \in \mathbb{N}$ be fixed and let $\mathcal{Q}_h : V_{H,\Gamma_D} \rightarrow V_{h,\Gamma_D}$ denote the corresponding correction operator as in Definition 2.4. We seek tuples $(\lambda_H^{(n)}, u_H^{(n)}) \in \mathbb{R} \times V_{H,\Gamma_D}$ with $\|u_H^{(n)} + \mathcal{Q}_h(u_H^{(n)})\|_{L^2(\Omega)} = 1$ such that for all $\Phi_H \in V_{H,\Gamma_D}$

$$\mathcal{A}(u_H^{(n)} + \mathcal{Q}_h(u_H^{(n)}), \Phi_H + \mathcal{Q}_h(\Phi_H)) = \lambda_H^{(n)} \int_{\Omega} (u_H^{(n)} + \mathcal{Q}_h(u_H^{(n)})) (\Phi_H + \mathcal{Q}_h(\Phi_H)) \quad (34)$$

and $0 < \lambda_H^{(n)} \leq \lambda_H^{(n+1)}$ for $0 \leq n < N_H$. We denote the arising LOD approximations of the eigenvectors by $u_{\text{LOD}}^{(n)} := (\text{Id} + \mathcal{Q}_h)(u_H^{(n)})$.

The eigenvalue problem (34) can be solved with any favorite solver. The main cost (for computing $\mathbf{A}_H^{\text{LOD}}$ and \mathbf{Q}_h) arise only once. Hence the computational advantage of the method becomes bigger the more eigenvalues we want to compute. Furthermore, the obtained convergence rates in H surpass the classical convergence rates. The following result was proved in [11].

Theorem 6.2. *Recall the reference problem (33). By C we denote a generic constant as in Theorem 6.1. Let $U_k(K) = \Omega$ for all $K \in \mathcal{T}_H$ and assume that $H \leq \left(\sqrt{n}\lambda_h^{(n)}\right)^{-1/2}$. Then it holds*

$$\frac{\lambda_H^{(n)} - \lambda_h^{(n)}}{\lambda_h^{(n)}} \leq Cn(\lambda_h^{(n)})^2 H^4 \quad (35)$$

for all $0 \leq n < N_H$.

Besides the fourth order convergence rates for the eigenvalues, it can be also shown that the corresponding eigenfunctions converge with higher rates (cf. [11]). Concerning the H^1 -error, the convergence is of quadratic order (in H) and concerning the L^2 -error of cubic order. These rates are higher than the rates that we obtained for the LOD for standard linear elliptic problems. Observe that these high convergence rates allow for much coarser grids and hence for a reduced computational complexity.

Remark 6.3 (Truncation). If the localization parameter k is chosen such that $k \geq m|\log(H)|$, the truncation error will be of order $\mathcal{O}(H^r)$ (with r as in Theorem 4.2). Consequently, for properly chosen m , the convergence rates in Theorem 6.2 remain valid even for the localized method.

6.2 Two-grid post-processing

The high convergence rates depicted in Theorem 6.2 can be even improved by applying a two-grid post-processing technique as initially suggested in [50]. The post-processing technique can be applied, if it is affordable to solve global (linear elliptic, non-eigenvalue) problems in the full fine space V_{h,Γ_D} . We define the post-processed LOD approximation as follows.

Definition 6.4 (Post-processed eigenvalue LOD approximation). Let $(\lambda_H^{(n)}, u_{\text{LOD}}^{(n)}) \in \mathbb{R} \times V_{h,\Gamma_D}$ (with $0 \leq n < N_H$) denote the eigenpair approximations obtained with the LOD as stated in Definition 6.1. We call $(\lambda_{H,\text{post}}^{(n)}, u_{\text{LOD},\text{post}}^{(n)}) \in \mathbb{R} \times V_{h,\Gamma_D}$ the corresponding post-processed approximations if $u_{\text{LOD},\text{post}}^{(n)} \in V_{h,\Gamma_D}$ solves

$$\mathcal{A}(u_{\text{LOD},\text{post}}^{(n)}, v_h) = \lambda_H^{(n)} \int_{\Omega} u_{\text{LOD}}^{(n)} v_h \quad \text{for all } v_h \in V_{h,\Gamma_D}$$

and where we define

$$\lambda_{H,\text{post}}^{(n)} := \frac{\mathcal{A}(u_{\text{LOD},\text{post}}^{(n)}, u_{\text{LOD},\text{post}}^{(n)})}{\|u_{\text{LOD},\text{post}}^{(n)}\|_{L^2(\Omega)}^2}.$$

Observe that the post-processing step involves to solve an additional linear elliptic problem in the full fine space V_{h,Γ_D} . Before discussing the feasibility of this step, we present a corresponding a priori error estimate.

Theorem 6.5. *Let $(\lambda_H^{(n)}, u_{\text{LOD}}^{(n)}) \in \mathbb{R} \times V_{h,\Gamma_D}$ (with $0 \leq n < N_H$) denote the LOD approximations as in Definition 6.1 and let $(\lambda_{H,\text{post}}^{(n)}, u_{\text{LOD},\text{post}}^{(n)}) \in \mathbb{R} \times V_{h,\Gamma_D}$ denote the post-processed LOD*

approximations as in Definition 6.4. Then it holds

$$\begin{aligned} |\lambda_{H,post}^{(n)} - \lambda_h^{(n)}| &\leq C(\lambda_H^{(n)} - \lambda_h^{(n)})^2 + C(\lambda_h^{(n)})^2 \|u_{LOD}^{(n)} - u_h^{(n)}\|_{L^2(\Omega)}^2 \quad \text{and} \\ \|u_{LOD,post}^{(n)} - u_h^{(n)}\|_{H^1(\Omega)} &\leq C|\lambda_H^{(n)} - \lambda_h^{(n)}| + C\lambda_h^{(n)} \|u_{LOD}^{(n)} - u_h^{(n)}\|_{L^2(\Omega)}. \end{aligned}$$

Consequently, we obtain that the eigenvalue $\lambda_{H,post}^{(n)}$ converges at least with order $\mathcal{O}(H^8)$ to $\lambda_h^{(n)}$ and that the H^1 -error between $u_{LOD,post}^{(n)}$ and $u_h^{(n)}$ converges to zero at least with order $\mathcal{O}(H^4)$.

In view of Theorem 6.5 we can see that the LOD can be a powerful tool to tackle linear eigenvalue problems even if κ has no multiscale character. If the cost for solving a full linear system on the fine scale is still feasible, it can be highly efficient to apply the LOD with pre- and post-processing. The extremely high convergence rates in H (at least $\mathcal{O}(H^8)$ if we are interested in the eigenvalues) allow to choose a very coarse grid \mathcal{T}_H . Depending on how coarse we choose \mathcal{T}_H , the truncation might be even skipped completely. First, we solve the corrector problems in a preprocessing step. This involves a number of linear elliptic fine scale problems that can be solved in parallel. After the correctors are available, we assemble the global LOD stiffness matrix \mathbf{A}_H^{LOD} , which has a very low dimension. All operations of the chosen algebraic eigenvalue solver, only involve \mathbf{A}_H^{LOD} and can be hence performed quickly and sequential. Once the eigenpairs are computed, we apply the postprocessing step, which only involves one global fine scale problem per eigenpair. Again, we can do this in parallel. With this approach, we can decrease the computational complexity. The implementation of pre- and post-processing is obvious, following the ideas presented in Section 4.

7 Numerical Examples

The algorithms present in this paper are available as a prototype implementation ¹. In this section we use this implementation to solve an elliptic boundary value problem and an eigenvalue problem. As the proposed algorithmic approach to the LOD is purely algebraic, the prototype is implemented in `python 2.7` using the modules `numpy` [51] and `scipy` [52]. The numerical example is implemented in `src/eigenvalues.py` and the actual LOD algorithm is implemented in the `compute_correction` function of the `lod` module. As input it requires different matrices, most of them directly related to the fine scale model:

- **Adc**: Sparse matrix with element wise contributions to the stiffness matrix, i.e. each diagonal block contains a particular local stiffness matrix \mathbf{A}_t , see definition 3.3. Basically this is the conforming fine scale operator assembled with discontinuous Galerkin shape functions.
- **fdc**: Vector with element wise contributions to the right-hand-side. Similarly defined as \mathbf{A}_t , it is necessary to compute the source corrector. In this implementation we do not distinguish between different right-hand-side contributions, but just compute a full corrector. Note, that one may save a bit, if the correct only needs to be computed for the boundary and not for the source term.
- **BH**: Coarse-Mesh boundary correction \mathbf{B}^H , see definition 3.5. This matrix has exactly one 1 in each row, every other entry is 0. It removes all constraint unknowns.
- **Mh**: Fine-mesh Mass-matrix, see definition 3.4. In many cases it can be convenient to also compute the mass matrix element wise, see 3.3, and then compute the global matrix as given in (10).
- **Ph**: Projection matrix \mathbf{P}_h from the coarse-mesh Lagrange space to the fine-mesh Lagrange space, see equation (14).

¹<https://gitlab.dune-project.org/christi/dune-py-lod/>




Figure 2: Model Problem 7.1. Left: Multiscale coefficient κ for $\varepsilon = 2^{-5}$. Right: Reference solution u obtained on a fine grid with resolution $h = 2^{-7}$. As a remark, the LOD solution for $H = 2^{-3}$ is visually indistinguishable from this reference solution.

- **Sigma**: Index-mapping from local to global dofs of the fine scale discretization, a block matrix containing all σ_t , see definition 3.2.
- **SubInfo** provides a list with different restriction operators for each sub-domain (actually our implementation provides this information via a generator). Each entry has to provide the following details:
 - **SubInfo.R** Restriction operator mapping from the fine space restriction to the patch, see equation (11).
 - **SubInfo.RH** Restriction operator mapping from the coarse space restriction to the patch, see equation (12).
 - **SubInfo.TH** Restriction operator from the coarse space to a coarse cell, see equation (15).

In our examples this fine scale model is implemented using the DUNE [53, 54] framework, a modern C++ library for grid based methods. It requires the DUNE core modules in version 2.4 [55] and DUNE-PDELab [56] in the 2.4 compatible version. For the sub-domain information we provide some additional infrastructure, so that it can be computed in `python`, eventhough it will usually be faster to assemble this information also in the framework.

The examples presented in the following are run on an AMD Epic 7501 server. Due to the global interpreter lock, python does not immediately allow for thread parallelization. Surely there are a range of different approaches for python to work around the limitations in multi-threading, but this would exceed the scope of an illustrative prototype implementations. We therefor decided to enforce all computations to be run sequentially on a single core. In particular for the eigenvalue solves, ARPACK's OpenMP parallelization was limited to a single thread to make timings comparable.

Furthermore, all our numerical experiments are performed for bilinear finite elements (cf. 1b) on uniform quadrilateral meshes with square elements.

7.1 Elliptic Problem

We first want to exemplify the algorithm and the basic performance using a standard elliptic model problem on the unit square $\Omega = (0, 1)^2$, with a heterogenous coefficient κ and a constant source term $f = 1$:

$$\begin{aligned} -\nabla \cdot \kappa \nabla u &= 1 & \text{in } \Omega, \\ u &= 0 & \text{on } \partial\Omega. \end{aligned}$$




Figure 3: Model Problem 7.1. Left: L^2 -errors for different coarse resolutions H . We compare the error $err_{coarse,H}$ for the standard approach with the LOD-error $err_{LOD,H}$ on the same mesh with overlap parameter $k \approx -0.9 \log_2(H)$. In addition, we list the times for solving the coarse problem $t_{coarse,H}$, the macroscopic LOD problem $t_{LOD,H}$, computing the correctors t_{corr} and for solving the full fine grid problem t_{fine} . Right: Performance comparison for solving the standard coarse problem compared to the macroscopic LOD-problem.

The diffusion coefficient

$$\kappa(x) = 1 + 10^{-8} + \frac{1}{2} \sin(\lfloor x_1 + x_2 \rfloor + \lfloor \frac{x_1}{\varepsilon} \rfloor + \lfloor \frac{x_2}{\varepsilon} \rfloor) + \frac{1}{2} \cos(\lfloor x_2 - x_1 \rfloor + \lfloor \frac{x_1}{\varepsilon} \rfloor + \lfloor \frac{x_2}{\varepsilon} \rfloor)$$

exhibits a multiscale structure. The fine scale reference solution is computed on level 7 (of quadrilateral mesh), i.e. $h = 2^{-7}$ and ε is chosen to be $\varepsilon = 2^{-5} = 4h$. Figure 2 shows the coefficient and the aforementioned fine scale solution used as a reference for the computation of errors. The same fine mesh is also used in all LOD computations. As for the LOD algorithm, we follow the descriptions presented in Section 5, which means that we included a source corrector $\mathcal{Q}_{1,h}$ for constant inputs in order to improve the accuracy. For optimal convergence orders, the overlap parameter k needs to be proportional to $\log(H)$, which is why we chose $k \approx -0.9 \log_2(H)$ in this numerical experiment. Timings for the different phases are presented in Figure 3. It should be noted that the implementation is to illustrate some central properties of the LOD in terms of performance, which is why it is not tuned for efficiency. In practice there are a couple of possible improvements as we will discuss later.

In Figure 3, $t_{coarse,H}$ denotes the time for solving the elliptic test problem using standard Q_1 -FEM on the coarse mesh \mathcal{T}_H ; $t_{LOD,H}$ denotes the time for solving the macroscopic LOD problem (29) on \mathcal{T}_H ; t_{corr} denotes the time for computing all the local correctors and for assembling the LOD-stiffness matrix associated with (29); finally, t_{fine} denotes the time for solving the test problem on the full fine mesh \mathcal{T}_h with resolution $h = 2^{-7}$. The relative errors are computed as

$$err_{coarse,H} = \frac{\|u - u_H\|_{L^2(\Omega)}}{\|u\|_{L^2(\Omega)}} \quad \text{and} \quad err_{LOD,H} = \frac{\|u - u_{LOD}\|_{L^2(\Omega)}}{\|u\|_{L^2(\Omega)}},$$

where u is the fine scale reference solution, u_H the standard finite element approximation on the coarse mesh and u_{LOD} the LOD approximation for a given coarse mesh \mathcal{T}_H , fine mesh \mathcal{T}_h and overlap parameter $k \approx -0.9 \log_2(H)$. From the graphs in Figure 3 we can see that the main cost of the LOD account for the assembly of the correctors which roughly grows (in a sequential implementation) with the rate $H^{-3/2}$, whereas the cost for solving the global LOD system are negligible in a one-shot simulation. We can also clearly see that the method becomes inefficient

if the fine mesh \mathcal{T}_h is only slightly finer than the coarse mesh \mathcal{T}_H . In practical situations, the fine mesh \mathcal{T}_h needs to be significantly finer than \mathcal{T}_H , where the regime $h \simeq H^2$ is often reasonable. We can also see that the costs for solving the global LOD system grow faster compared to the costs for solving the standard coarse system of the same dimension. The reason is that the number of nonzeros in the system matrix grows by the factor k^d , where k is the localization parameter. This pollution of the sparsity structure causes an increased computational complexity. However, comparing these increased costs with the improvement of the accuracy in the right graph of Figure 3, we can see that the LOD is significantly more efficient if we just look at the macroscopic solve. This is an important observation since the one-time costs for computing the correctors become negligible if there is either a high capacity for parallelization or if the calculations have to be repeated for many source terms. In the next model problem we will face such a situation where the overhead caused by the computation of the correctors is compensated by its repeated usage and where the LOD outperforms the standard approach on the fine grid.

7.2 Eigenvalue Problem

In the previous model we demonstrated the basic performance of the algorithm for one fixed source term. However, the big advantage of the LOD is getting more pronounced, if computations have to be repeated with different source terms. To pick a natural setting where this is the case, our next model is an eigenvalue problem. Here we follow the LOD algorithm described in Section 6. We consider computing the 20 smallest eigenvalues for the stationary linear Schrödinger equation with a discontinuous potential V . The problem is inspired by the Kronig-Penney model (cf. [57]) and exhibits fine scale heterogeneities and high contrast of order γ . Here we seek eigenfunctions u for the 20 smallest eigenvalues λ such that

$$-\Delta u + V(x)u = \lambda u \quad \text{in } \Omega, \quad (36a)$$

$$u = 0 \quad \text{on } \partial\Omega, \quad (36b)$$

with

$$V(x) = \gamma \left[\cos(\pi\nu(x_1 + 0.1)) \cos(\pi\nu x_2) \right]. \quad (36c)$$

As the `python` implementation is not parallelized, we enforced sequential computation of the eigenvalue problem. Different resolutions for the coarse space and different overlap sizes are considered. We solve for the 20 smallest eigenvalues and eigenmodes and compare the obtained eigenvalues for the coarse discretization, the full discretization on a fine mesh with resolution $h = 2^{-8}$ and the LOD. Overall the obtained timings are convincing and the results of the LOD yield good accuracy. All eigenvalue problems are solved using the `ARPACK` library [58] with the shift-inverse method. For the inner solve (in the inverse power iteration) we used a CG Krylov solver with the `pyamg`[59] algebraic multigrid as a preconditioner, which then leads to reproducible robust results. In our numerical experiments we used the following parameters, if not indicated differently for a particular experiment. The domain $\Omega = (0, 2) \times (0, 3)$, the wave-number $\nu = 20$, and the ratio $\gamma = 2 \cdot 10^4$. Figure 4 shows four selected eigenmodes for a fine scale resolution of $h = 2^{-8}$ and a coarse resolution of $H = 2^{-3}$. We compare the coarse simulation (which is too coarse to actually pick up the fine scale structure of the solution), the full simulation, using the fine scale discretization and two LOD simulations, with an overlap of H and $2H$. We measured the computational time and the error of different setup. Figure 5 shows the relative error and the costs (relative error / time) for different coarse mesh sizes and a fixed overlap of H . From the left graph in Figure 5 we can see that the LOD error $\text{err}_{\text{LOD},H}$ remains small for all coarse mesh sizes, where the total LOD computing time $t_{\text{LOD},H} + t_{\text{corr}}$ is smaller than the computing time t_{fine}




Figure 4: Model Problem 7.2 with $\nu = 20$ and $\gamma = 2 \cdot 10^4$. Comparison of the 1st, 6th, 11th, and 20th eigenmode computed for problem (36a) using the coarse ($H = 2^{-3}$), the full ($h = 2^{-8}$) and the LOD discretization ($H = 2^{-3}$, overlap H and $2H$, i.e. $k = 1$ and $k = 2$).




Figure 5: Model Problem 7.2 with $\nu = 20$ and $\gamma = 2 \cdot 10^4$. Timing and accuracy for different resolution H of the coarse mesh (left) and performance as error/time (right).

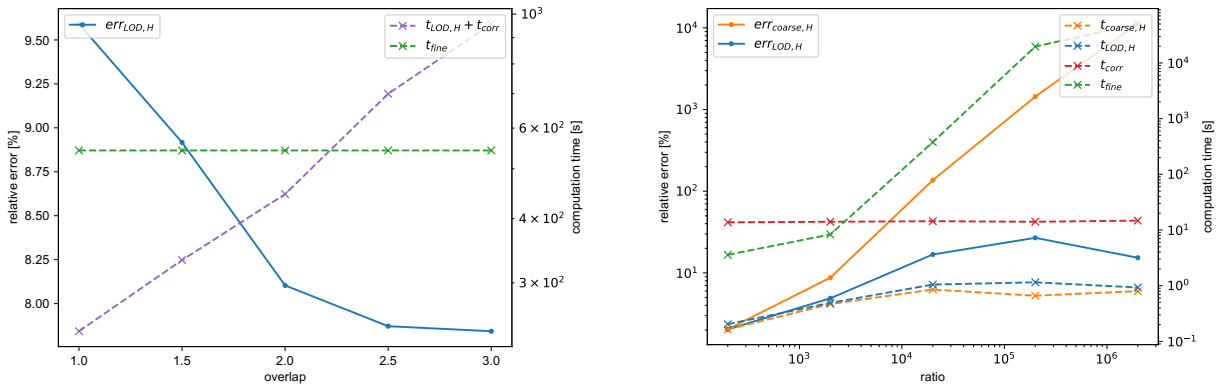


Figure 6: Model Problem 7.2. Timing and accuracy of the LOD for (left) different overlap sizes ($H = 2^{-3}$, $h = 2^{-8}$, $\gamma = 2 \cdot 10^4$) and (right) different ratios γ ($H = 2^{-3}$, $h = 2^{-8}$, overlap H).

for the full fine scale discretization, which is a remarkable observation. We can also see that the “offline costs” t_{corr} (i.e. the time for computing the corrections) still amounts for the major costs of the LOD, at least for reasonably coarse mesh sizes H .

Even more interesting results are depicted in Figure 6, which shows the relative error for overlap sizes and different coefficient ratios γ . The error was measured with respect to the eigenvalues of the fine scale solution. Given the vectors $\Lambda_{\text{coarse},H}$, $\Lambda_{\text{LOD},H}$ and Λ_{fine} of the first 20 eigenvalues, the relative errors are computed as

$$\text{err}_{\text{coarse},H} = \|\Lambda_{\text{coarse},H} - \Lambda_{\text{fine}}\|_{\infty}, \quad \text{err}_{\text{LOD},H} = \|\Lambda_{\text{LOD},H} - \Lambda_{\text{fine}}\|_{\infty}.$$

For the computation time, we consider the time t_{fine} for computing the 20 smallest eigenvectors and their eigenmodes using the fine scale discretization, the time $t_{\text{coarse},H}$ for the eigenvalue solve of the coarse system and $t_{\text{LOD},H}$ for the eigenvalue solve of the LOD system. In addition the LOD requires the usual preprocessing, in order to compute the correction of coarse space basis. The time necessary for this setup phase is denoted by t_{corr} as before. The left graph of Figure 6 shows that even for a high contrast parameter γ , an overlap of two coarse elements is sufficient so that truncation error becomes negligible. The right figure shows the efficiency of the LOD compared to a regular fine-scale computation (using AMG), in particular for large γ . For $\gamma \approx 5 \cdot 10^4$ both approaches show an equal performance. For $\gamma \geq 5 \cdot 10^4$ the efficiency of the LOD compared to the standard implementation is continuously increasing. In this regime, we observe that even for a sequential implementation of the LOD method, the total run time $t_{\text{LOD},H} + t_{\text{corr}}$ is significantly below the run times $t_{\text{coarse}}/t_{\text{fine}}$ of the multigrid alternative. Here the LOD was up to 1000 times faster than the multigrid alternative, which shows its enormous potential for even larger problems.

7.3 Efficiency

These examples just consider the actual solve and we didn’t compute the additional post-processing described in section 6.2 and the timings exclude the assembly of the different operators. All computation done in DUNE are very fast and negligible compared to the overall solving time. In practice, one would use several improvements, in particular, the different sub-domain problems (algorithm 1, line 14) are completely independent and can be solved in parallel, which yields a perfect speedup and is easy to implement for modern many-core systems. Additionally the necessary solves for different right-hand-sides in algorithm 1, line 33 allows for a slight reformulation of the algorithm and using vector instructions (e.g. SSE, AVX, Neon) to compute all updates in a single run.

8 Conclusions

In this contribution we presented an efficient implementation of the Localized Orthogonal Decomposition (LOD), including several applications and variations of the methodology. The efficiency of the algorithms is verified in numerical experiments, where we demonstrated that the approach can be even very powerful in its sequential implementation. This aspect is specifically stressed by Figure 6 where we compare the CPU times for the LOD with the CPU times of an efficient algebraic multigrid solver (AMG) on the fine scale depending on γ . It can be observed that the computational complexity of the LOD is independent of the contrast parameter γ , whereas the reference solver is not robust. For $\gamma = 10^6$, the run times using algebraic multigrid were of order 10^3 times higher than for the LOD. Furthermore, we observed that the AMG implementation suffered from “false eigenvalues” caused by numerical rounding errors. Ratios of order $\gamma > \mathcal{O}(10^6)$ could no longer be handled by the AMG, whereas the LOD was still performing well and with the same run times as for small values of γ .

Acknowledgements

This work was supported by: the German Research Foundation (DFG) through the Priority Programme ‘Software for Exascale Computing (SPP 1648)’ (grant EN-1042/2-1 and EN-1042/2-2) [supporting C. Engwer], the Swedish Research Council (grant 2016-03339) [supporting P. Henning] and (grant 2015-04964) [supporting A. Målqvist], and the *Hausdorff Institute for Mathematics* in Bonn for the kind hospitality of during the trimester program on multiscale problems [supporting P. Henning, A. Målqvist and D. Peterseim].

References

- [1] A. Målqvist, D. Peterseim, Localization of elliptic multiscale problems, *Math. Comp.* 83 (290) (2014) 2583–2603. doi:10.1090/S0025-5718-2014-02868-8.
- [2] P. Henning, D. Peterseim, Oversampling for the Multiscale Finite Element Method, *SIAM Multiscale Model. Simul.* 11 (4) (2013) 1149–1175. doi:10.1137/120900332.
- [3] P. Henning, A. Målqvist, Localized orthogonal decomposition techniques for boundary value problems, *SIAM J. Sci. Comput.* 36 (4) (2014) A1609–A1634. doi:10.1137/130933198.
- [4] D. Elfverson, E. H. Georgoulis, A. Målqvist, D. Peterseim, Convergence of a discontinuous Galerkin multiscale method, *SIAM J. Numer. Anal.* 51 (6) (2013) 3351–3372. doi:10.1137/120900113.
- [5] D. Elfverson, E. H. Georgoulis, A. Målqvist, An adaptive discontinuous Galerkin multiscale method for elliptic problems, *Multiscale Model. Simul.* 11 (3) (2013) 747–765. doi:10.1137/120863162.
- [6] D. Elfverson, A discontinuous Galerkin multiscale method for convection-diffusion problems, *Uppsala University Preprint* (2014).
- [7] A. Målqvist, Multiscale methods for elliptic problems, *Multiscale Model. Simul.* 9 (3) (2011) 1064–1086. doi:10.1137/090775592.
- [8] F. Hellman, P. Henning, A. Målqvist, Multiscale mixed finite elements, *Discrete Contin. Dyn. Syst. Ser. S* 9 (5) (2016) 1269–1298. doi:10.3934/dcdss.2016051.
- [9] P. Henning, P. Morgenstern, D. Peterseim, Multiscale partition of unity, in: M. Griebel, M. A. Schweitzer (Eds.), *Meshfree Methods for Partial Differential Equations VII*, Vol. 100 of *Lecture Notes in Computational Science and Engineering*, Springer International Publishing, 2015, pp. 185–204. doi:10.1007/978-3-319-06898-5_10.
- [10] A. Abdulle, P. Henning, A reduced basis localized orthogonal decomposition, *J. Comput. Phys.* 295 (2015) 379–401. doi:10.1016/j.jcp.2015.04.016.
- [11] A. Målqvist, D. Peterseim, Computation of eigenvalues by numerical upscaling, *Numer. Math.* 130 (2) (2015) 337–361. doi:10.1007/s00211-014-0665-6.
- [12] A. Målqvist, D. Peterseim, Generalized finite element methods for quadratic eigenvalue problems, *ESAIM Math. Model. Numer. Anal.* 51 (1) (2017) 147–163. doi:10.1051/m2an/2016019.
- [13] D. L. Brown, D. Peterseim, A multiscale method for porous microstructures, *Multiscale Model. Simul.* 14 (3) (2016) 1123–1152. doi:10.1137/140995210.

- [14] D. Peterseim, R. Scheichl, Robust numerical upscaling of elliptic multiscale problems at high contrast, *Comput. Methods Appl. Math.* 16 (4) (2016) 579–603. doi:10.1515/cmam-2016-0022.
- [15] F. Hellman, A. Målqvist, Contrast independent localization of multiscale problems, *Multiscale Model. Simul.* 15 (4) (2017) 1325–1355. doi:10.1137/16M1100460.
- [16] D. Gallistl, D. Peterseim, Numerical stochastic homogenization by quasilocal effective diffusion tensors, *ArXiv e-prints* arXiv:1702.08858.
- [17] M. Feischl, D. Peterseim, Sparse Compression of Expected Solution Operators, *arXiv e-prints* (2018) arXiv:1807.01741 arXiv:1807.01741.
- [18] P. Henning, A. Målqvist, D. Peterseim, A localized orthogonal decomposition method for semi-linear elliptic problems, *M2AN Math. Model. Numer. Anal.* 48 (5) (2014) 1331–1349. doi:10.1051/m2an/2013141.
- [19] A. Abdulle, P. Henning, Localized orthogonal decomposition method for the wave equation with a continuum of scales, *Math. Comp.* 86 (304) (2017) 549–587. doi:10.1090/mcom/3114.
- [20] D. Peterseim, M. Schedensack, Relaxing the CFL condition for the wave equation on adaptive meshes, *J. Sci. Comput.* 72 (3) (2017) 1196–1213.
- [21] R. Maier, D. Peterseim, Explicit computational wave propagation in micro-heterogeneous media, *BIT Numerical Mathematics* doi:10.1007/s10543-018-0735-8. URL <https://doi.org/10.1007/s10543-018-0735-8>
- [22] A. Målqvist, A. Persson, Multiscale techniques for parabolic equations, *Numer. Math.* 138 (1) (2018) 191–217. doi:10.1007/s00211-017-0905-7.
- [23] Målqvist, Axel, Persson, Anna, A generalized finite element method for linear thermoelasticity, *ESAIM: M2AN* 51 (4) (2017) 1145–1171. doi:10.1051/m2an/2016054. URL <https://doi.org/10.1051/m2an/2016054>
- [24] R. Altmann, E. Chung, R. Maier, D. Peterseim, S.-M. Pun, Computational multiscale methods for linear heterogeneous poroelasticity, *arXiv e-prints* (2018) arXiv:1801.00615 arXiv:1801.00615.
- [25] D. Elfverson, V. Ginting, P. Henning, On multiscale methods in petrov–galerkin formulation, *Numerische Mathematik* 131 (4) (2015) 643–682. doi:10.1007/s00211-015-0703-z. URL <http://dx.doi.org/10.1007/s00211-015-0703-z>
- [26] D. Brown, J. Gedicke, D. Peterseim, Numerical homogenization of heterogeneous fractional laplacians, *Multiscale Modeling & Simulation* 16 (3) (2018) 1305–1332. arXiv:https://doi.org/10.1137/17M1147305, doi:10.1137/17M1147305. URL <https://doi.org/10.1137/17M1147305>
- [27] D. Peterseim, Eliminating the pollution effect in Helmholtz problems by local subscale correction, *Math. Comp.* 86 (305) (2017) 1005–1036. doi:10.1090/mcom/3156.
- [28] D. Gallistl, D. Peterseim, Stable multiscale Petrov-Galerkin finite element method for high frequency acoustic scattering, *Comp. Meth. Appl. Mech. Eng.* 295 (2015) 1–17. doi:10.1016/j.cma.2015.06.017.

- [29] D. Brown, D. Gallistl, D. Peterseim, Multiscale Petrov-Galerkin method for high-frequency heterogeneous Helmholtz equations, in: M. Griebel, M. A. Schweitzer (Eds.), *Meshfree Methods for Partial Differential Equations VII*, Lecture Notes in Computational Science and Engineering, Springer, 2017, pp. 85–115.
- [30] D. L. Brown, D. Gallistl, Multiscale Sub-grid Correction Method for Time-Harmonic High-Frequency Elastodynamics with Wavenumber Explicit Bounds, arXiv e-prints (2016) arXiv:1608.04243arXiv:1608.04243.
- [31] P. Henning, A. Målqvist, D. Peterseim, Two-Level Discretization Techniques for Ground State Computations of Bose-Einstein Condensates, *SIAM J. Numer. Anal.* 52 (4) (2014) 1525–1550. doi:10.1137/130921520.
- [32] D. Peterseim, Variational multiscale stabilization and the exponential decay of fine-scale correctors, in: *Building bridges: connections and challenges in modern approaches to numerical partial differential equations*, Vol. 114 of *Lect. Notes Comput. Sci. Eng.*, Springer, [Cham], 2016, pp. 341–367.
- [33] T. J. R. Hughes, Multiscale phenomena: Green’s functions, the Dirichlet-to-Neumann formulation, subgrid scale models, bubbles and the origins of stabilized methods, *Comput. Methods Appl. Mech. Engrg.* 127 (1-4) (1995) 387–401. doi:10.1016/0045-7825(95)00844-9.
- [34] T. J. R. Hughes, G. R. Feijóo, L. Mazzei, J.-B. Quincy, The variational multiscale method—a paradigm for computational mechanics, *Comput. Methods Appl. Mech. Engrg.* 166 (1-2) (1998) 3–24. doi:10.1016/S0045-7825(98)00079-6.
- [35] T. J. R. Hughes, G. Sangalli, Variational multiscale analysis: the fine-scale Green’s function, projection, optimization, localization, and stabilized methods, *SIAM J. Numer. Anal.* 45 (2) (2007) 539–557. doi:10.1137/050645646.
- [36] M. G. Larson, A. Målqvist, Adaptive variational multiscale methods based on a posteriori error estimation: energy norm estimates for elliptic problems, *Comput. Methods Appl. Mech. Engrg.* 196 (21-24) (2007) 2313–2324. doi:10.1016/j.cma.2006.08.019.
- [37] M. G. Larson, A. Målqvist, An adaptive variational multiscale method for convection-diffusion problems, *Comm. Numer. Methods Engrg.* 25 (1) (2009) 65–79. doi:10.1002/cnm.1106.
- [38] D. Gallistl, D. Peterseim, Computation of quasi-local effective diffusion tensors and connections to the mathematical theory of homogenization, *Multiscale Model. Simul.* 15 (4) (2017) 1530–1552. doi:10.1137/16M1088533.
- [39] R. Kornhuber, D. Peterseim, H. Yserentant, An analysis of a class of variational multiscale methods based on subspace decomposition, *Math. Comp.* published electronically. doi:10.1090/mcom/3302.
- [40] D. Peterseim, D. Varga, B. Verfürth, From Domain Decomposition to Homogenization Theory, arXiv e-prints (2018) arXiv:1811.06319arXiv:1811.06319.
- [41] H. Owhadi, L. Zhang, L. Berlyand, Polyharmonic homogenization, rough polyharmonic splines and sparse super-localization, *ESAIM Math. Model. Numer. Anal.* 48 (2) (2014) 517–552. doi:10.1051/m2an/2013118.

- [42] R. Kornhuber, H. Yserentant, Numerical homogenization of elliptic multiscale problems by subspace decomposition, *Multiscale Model. Simul.* 14 (3) (2016) 1017–1036. doi:10.1137/15M1028510.
- [43] H. Owhadi, Multigrid with rough coefficients and multiresolution operator decomposition from hierarchical information games, *SIAM Rev.* 59 (1) (2017) 99–149. doi:10.1137/15M1013894.
- [44] C. Carstensen, Quasi-interpolation and a posteriori error analysis in finite element methods, *M2AN Math. Model. Numer. Anal.* 33 (6) (1999) 1187–1202. doi:10.1051/m2an:1999140.
- [45] P. Clément, Approximation by finite element functions using local regularization, *Rev. Française Automat. Informat. Recherche Opérationnelle Sér. RAIRO Analyse Numérique* 9 (R-2) (1975) 77–84.
- [46] L. R. Scott, S. Zhang, Finite element interpolation of nonsmooth functions satisfying boundary conditions, *Math. Comp.* 54 (190) (1990) 483–493. doi:10.2307/2008497.
- [47] C. Carstensen, R. Verfürth, Edge residuals dominate a posteriori error estimates for low order finite element methods, *SIAM J. Numer. Anal.* 36 (5) (1999) 1571–1587. doi:10.1137/S003614299732334X.
- [48] A. Ern, J.-L. Guermond, Finite element quasi-interpolation and best approximation, *ESAIM Math. Model. Numer. Anal.* 51 (4) (2017) 1367–1385.
- [49] R. Kornhuber, J. Podlesny, H. Yserentant, Direct and iterative methods for numerical homogenization, in: C.-O. Lee, X.-C. Cai, D. E. Keyes, H. H. Kim, A. Klawonn, E.-J. Park, O. B. Widlund (Eds.), *Domain Decomposition Methods in Science and Engineering XXIII*, Springer International Publishing, Cham, 2017, pp. 217–225.
- [50] J. Xu, A. Zhou, A two-grid discretization scheme for eigenvalue problems, *Math. Comp.* 70 (233) (2001) 17–25. doi:10.1090/S0025-5718-99-01180-1.
- [51] S. v. d. Walt, S. C. Colbert, G. Varoquaux, The numpy array: a structure for efficient numerical computation, *Computing in Science & Engineering* 13 (2) (2011) 22–30.
- [52] K. J. Millman, M. Aivazis, Python for scientists and engineers, *Computing in Science & Engineering* 13 (2) (2011) 9–12.
- [53] P. Bastian, M. Blatt, A. Dedner, C. Engwer, R. K. , M. Ohlberger, O. Sander, A generic grid interface for parallel and adaptive scientific computing. Part I: Abstract framework, *Computing* 82 (2–3) (2008) 103–119. doi:10.1007/s00607-008-0003-x.
- [54] P. Bastian, M. Blatt, A. Dedner, C. Engwer, R. K. , R. Kornhuber, M. Ohlberger, O. Sander, A generic grid interface for parallel and adaptive scientific computing. Part II: Implementation and tests in DUNE, *Computing* 82 (2–3) (2008) 121–138. doi:10.1007/s00607-008-0004-9.
- [55] M. Blatt, A. Burchardt, A. Dedner, C. Engwer, J. Fahlke, B. Flemisch, C. Gersbacher, C. Gräser, F. Gruber, C. Grüninger, et al., The distributed and unified numerics environment, version 2.4, *Archive of Numerical Software* 4 (100) (2016) 13–29.
- [56] P. Bastian, F. Heimann, S. Marnach, Generic implementation of finite element methods in the distributed and unified numerics environment (dune), *Kybernetika* 46 (2) (2010) 294–315.

- [57] K. Varga, J. A. Driscoll, Computational Nanoscience: Applications for Molecules, Clusters, and Solids, 1st Edition, Cambridge University Press, New York, NY, USA, 2011.
- [58] R. B. Lehoucq, D. C. Sorensen, C. Yang, ARPACK users' guide: solution of large-scale eigenvalue problems with implicitly restarted Arnoldi methods, SIAM, 1998.
- [59] W. N. Bell, L. N. Olson, J. Schroder, PyAMG: Algebraic Multigrid Solvers in Python, version 2.1 (2013).
URL <http://www.pyamg.org>