

Inexact linear solvers for control volume discretizations in porous media

Eirik Keilegavlen · Jan M. Nordbotten

Received: 26 June 2013 / Accepted: 17 October 2014
© Springer International Publishing Switzerland 2014

Abstract We discuss the construction of multi-level inexact linear solvers for control volume discretizations for porous media. The methodology forms a contrast to standard iterative solvers by utilizing an algebraic hierarchy of approximations which preserve the conservative structure of the underlying control volume. Our main result is the generalization of multiscale control volume methods as multi-level inexact linear solvers for conservative discretizations through the design of a particular class of preconditioners. This construction thereby bridges the gap between multiscale approximation and linear solvers. The resulting approximation sequence is referred to as inexact solvers. We seek a conservative solution, in the sense of control-volume discretizations, within a prescribed accuracy. To this end, we give an abstract guaranteed a posteriori error bound relating the accuracy of the linear solver to the underlying discretization. These error bounds are explicitly computable for the grids considered herein. The afore-mentioned hierarchy of conservative approximations can also be considered in the context of multi-level upscaling, and this perspective is highlighted in the text as appropriate. The new construction is supported by numerical examples highlighting the performance of the inexact linear solver realized in both a multi- and two-level context for two- and three-dimensional heterogeneous problems defined on structured and unstructured grids. The numerical examples assess the performance of the approach both as an inexact solver, as well in comparison to standard algebraic multigrid methods.

Keywords Linear solvers · Inexact solvers · Multiscale · Multigrid · Control volume methods

1 Introduction

Control volume discretizations are prevailing for flow calculations in geological porous media, including important applications such as oil and gas recovery and geological CO₂ storage [6]. The motivation for using control volume discretizations stems from the ease of guaranteeing exact flux conservation, which is important for the stability of the fluid flow calculations [2]. Low-order control volume methods also adapt well to the strong material heterogeneities of geological media. However, these heterogeneous structures are commonly described in a level of detail that makes converged numerical simulations infeasible by traditional approaches. This has led to the development of multiscale approximation methods that aim to capture the main features of the solution to a relatively low computational cost [17]. The fine-scale features that are not resolved in the approximated solution are commonly of lesser importance; furthermore, the data uncertainty that is inherent for the subsurface warrants an inexact solution procedure. However, while multiscale approximations lead to promising results on model problems, they have struggled with robustness issues in more general situations [21]. To remedy the robustness issues, both adaptive [30] and iterative [28] approaches have been suggested.

When multiscale methods have been extended to allow for iterations and adaptivity, they resemble preconditioners for classical iterative methods such as domain decomposition and multigrid. A key difference is that multiscale methods commonly are constructed to provide solutions that share crucial properties such as conservation with the

E. Keilegavlen (✉) · J. M. Nordbotten
Department of Mathematics, University of Bergen,
Bergen, Norway
e-mail: Eirik.Keilegavlen@math.uib.no

exact solution. Thus, in contrast with traditional domain decomposition and multigrid methods, an inexact solution from a multiscale can be used for transport purposes without leading to conservation errors. However, the desire to preserve the conservation structure has mostly limited the multiscale methods to two scales, with only a few exceptions, e.g., [25],[23]. For a two-scale approximation that is conservative on the coarse scale, a conservative approximation can be obtained also on the fine-scale through local post-processing, a key motivation for choosing control volume methods [17]. Unfortunately, the two-level post-processing does not extend trivially to multiple levels [35], and thus previous multiscale methods have mostly been limited to two-scale applications, although a multilevel method for structured grids was recently reported in [23].

In this paper, we present a novel framework for conservative multiscale methods. Following the spirit of [28], we develop the multiscale methods from the perspective of conservative preconditioners for iterative solvers. Our work makes the following extensions over classical multiscale approximations

1. We formulate a truly multilevel conservative preconditioner, where the exact conservation structure can be identified at all scales.
2. At all scales, we retain explicit local flux expressions, allowing the preconditioner to be interpreted as a multilevel control volume discretization.
3. We present a new hierarchical post-processing strategy for multilevel conservative approximations that is applicable for more than two levels.
4. Our framework allows for general grids, thus opening the possibility for algebraic aggregation of coarse levels.

The fully multilevel control volume approach presented herein bridges gaps between multilevel linear solvers, discretization methods, and upscaling. As such, it will benefit from the developments and understanding of each of these various disciplines. We highlight some connections of relevance below.

To arrive at solvers with the desired properties, we are interested in methods that preserve the same structures and thus are also of the same form as the underlying discretization methods. This is a natural thought in the setting of linear solvers for finite element setting and is the foundation of various developments such as geometric multigrid [44], through recent works such as algebraic multilevel methods preserving de Rahm complexes [31]. However, due to much fewer available theoretical results for control volume methods, and the added complexity of dealing with separate shape and test functions, these ideas have received less attention in this context.

The first examples of multilevel control volume methods were designed for the purpose of upscaling. Indeed, here the main focus was on deriving transmissibilities (which can be interpreted as the product appearing in the boundary integral of the left-hand side of Eq. 2.6 below) [11]. As such, coarse basis functions were never explicitly derived, although iterative two-level methods have nevertheless been implemented [5]. Interestingly, these iterative methods do not converge to the exact solution, but the resulting approximation is still considered of sufficient accuracy for applications.

Explicit calculations of coarse basis functions appeared in the setting of discretization methods with the introduction of the so-called multipoint flux approximation (see [1] for an overview). In the context of a multilevel approximation it appeared after the introduction of the multiscale finite element method [14], and was later applied in the control volume context [17]. The relationship between the multiscale approximations and domain decomposition solvers was then observed in [28].

The remainder of the paper is structured as follows: We make precise the mathematical context and model problem in Section 2. A hierarchy of multilevel control volume discretizations is formulated in Section 3. In Section 4 a method for post-processing an inexact solution such that the conservation structure is preserved for all grid levels is presented. Error bounds are considered in Section 5, Section 6 contains numerical examples and the paper is finalized by concluding remarks in Section 7.

2 Model equations and discretization

To make the presentation concrete, we consider the model equation for the fluid potential in porous media, which is the heterogeneous elliptic problem. Given a domain $\Omega \subset \mathbb{R}^d$, $d = \{2, 3\}$, with boundary $\partial\Omega = \Gamma_D \cup \Gamma_N$, a symmetric positive definite coefficient matrix \mathbf{a} , which is a measurable function of space, and right hand side f , we seek u which solves

$$-\nabla \cdot (\mathbf{a}\nabla u) = f \quad (2.1)$$

with the boundary conditions

$$u = u_D \quad \text{on } \Gamma_D \quad (2.2a)$$

$$-\mathbf{n} \cdot (\mathbf{a}\nabla u) = q_N \quad \text{on } \Gamma_N \quad (2.2b)$$

where u_D and q_N are prescribed Dirichlet and Neumann boundary conditions, respectively and \mathbf{n} is the outer normal

vector. Alternatively, the model problem can be stated on mixed form, to read: Find u and \mathbf{q} such that

$$\mathbf{q} = -\mathbf{a}\nabla u \tag{2.3a}$$

$$\nabla \cdot \mathbf{q} = f \tag{2.3b}$$

still subject to the boundary conditions (2.2a)–(2.2b). We will consider the mixed formulation in Section 4, but for now, we focus on the form (2.1)–(2.2a). The strong form of the equations as stated above requires relatively smooth coefficients and right-hand side f . However, in applications it is common to have no more regularity than \mathbf{a} being defined almost everywhere with uniformly bounded eigenvalues, and $f \in L_2$. Assuming Γ_D has nonzero measure, there exists a unique weak solution $u \in H^1(\Omega)$ satisfying (2.1)–(2.2b), and a unique weak solution pair $(u, \mathbf{q}) \in H^1(\Omega) \times H(\text{div})$ satisfying (2.2a)–(2.3b). Numerical approximations are typically developed to approximate these weak solutions.

Equations (2.1)–(2.3b) are ubiquitous in applications. In this work, we are motivated by flow in geological porous media, but the results presented herein will be of interest for a wider range of applications including heat flow in composite materials and fluid flow in biological tissues. In our setting, u represents a fluid potential, \mathbf{a} represents the permeability of the medium, and f is a source or sink term of fluid and (2.3a) is the so-called Darcy’s law.

In practice, (2.1) forms part of a larger nonlinear system, whose stability may depend on certain conservation properties. In particular, often a hyperbolic conservation law transports mass on the flux \mathbf{q} , and it is then of importance that the conservation structure highlighted by the mixed formulation as stated in Eqs. 2.3a–2.3b is preserved by the discrete representation (see e.g., [24]). Such discretizations, both as applied to Eq. 2.1 and to hyperbolic laws, are known as control volume or finite volume discretizations and are of present of interest to us.

We proceed to derive the general form of a control volume discretization from Eq. 2.1. For simplicity define the space H_a^2 such that $u \in H_a^2$ if $\nabla \cdot (\mathbf{a}\nabla u) \in L^2$. Then, after multiplication by a test function and integration, we see that a solution of Eq. 2.1 with the assumption that $u \in H_a^2$ also solves the weak variational problem: Find $u \in U \subset H_a^2$ such that

$$a(u, v) = f(v) \quad \forall v \in V \subseteq L^2 \tag{2.4}$$

Here, the bilinear form a is defined by

$$a(u, v) = - \int_{\Omega} \nabla \cdot (\mathbf{a}\nabla u) v \, dx \quad \text{and} \quad f(v) = \int_{\Omega} f v \, dx$$

Note that the spaces U and V need not in general coincide, but need to be compatible, as is ensured by the formulation

(2.4) and the requirement $u \in H_a^2$. The formulation is equivalent to the mixed variational formulation of the control volume methods [38].

To construct finite-dimensional approximations of these function spaces, and thus a numerical discretization, we partition the domain into a finite-dimensional grid with d cells, which we will denote ω_i , with boundary $\partial\omega_i$. We refer to the boundary between two cells as a face, independent of the spatial dimension, and write $\gamma_k = \partial\omega_{ij}$ to indicate that face k divides cells i and j . Throughout the manuscript, index i and j will denote generic cells, while k denotes a face. When sums or unions are taken of cells and faces, the summation variable will be r and s , respectively.

The control volume methods are identified as the methods obtained when the trial functions lie in the subspace $V_h \subset V$ consisting of piecewise constant functions of the cells ω_i . Various control volume methods can then be defined by their choice of space U_h . These are to a large extent nonconforming (e.g. $U_h \not\subset U$), which is acceptable as long as U_h lies in the dual of V_h with respect to the bilinear form a . The simplest choice is known as the control volume finite element method, wherein U_h is piece-wise constant on a triangulation dual to the cells ω_i [45],[4],[42], while more advanced constructions are developed in recent work ([1],[10]).

We thus assume that Eq. 2.4 has been discretized by a control volume method such that the resulting finite-dimensional problem can be identified with the solution of a discrete variational problem: Find the solution vector $\mathbf{u}_h \in R^n$ such that

$$a(\boldsymbol{\psi}_h \cdot \mathbf{u}_h, \boldsymbol{\phi}_h \cdot \mathbf{e}_j) = f(\boldsymbol{\phi}_h \cdot \mathbf{e}_j) \quad \forall j = \{1, \dots, n\} \tag{2.5}$$

Here, $\boldsymbol{\psi}_h$ is a vector of linearly independent basis functions $\boldsymbol{\psi}_h = [\boldsymbol{\psi}_{h,1}, \dots, \boldsymbol{\psi}_{h,n}]$ spanning the discrete space U_h . These basis functions are classically referred to as the shape functions of the solution. Similarly $\boldsymbol{\phi}_h$ is a vector of linearly independent basis functions spanning V_h referred to as the test functions. Finally, \mathbf{e}_j is the j th unit vector of dimension n , and in accordance with the definition of control volume methods, we will always choose $\boldsymbol{\phi}_{h,j} = \boldsymbol{\phi}_h \cdot \mathbf{e}_j$ as the indicator function for cell j .

The name control volume stems from the local conservation structure of the method, as for V_h chosen as above, the variational problem (2.5) can after integration by parts be written as follows: Find $u_h \in U_h$ such that

$$\int_{\partial\omega_j} (-\mathbf{a}\nabla u_h) \cdot \mathbf{n} \, dx = \int_{\omega_j} f \, dx \quad \forall j = \{1, \dots, n\} \tag{2.6}$$

We note that as Eq. 2.6 represents the physical conservation principle, it is common to state it directly, and consequently derive the differential form given in Eq. 2.1. We have chosen the opposite approach to highlight the structural similarities with the Galerkin finite element method and, for the coarse

problems introduced in Section 3, the Galerkin formulation of multigrid methods. We will also write the systems (2.5) and (2.6) in terms of the discrete bilinear form: Find the discrete solution vector \mathbf{u}_h such that

$$a_h(\mathbf{u}_h, \mathbf{e}_j) = f_h(\mathbf{e}_j) \quad \forall j = \{1, \dots, n\} \tag{2.7}$$

or equivalently, solve the linear system

$$\mathbf{A}\mathbf{u}_h = \mathbf{b} \tag{2.8}$$

The different control volume methods are usually discussed in terms of the surface integrals appearing on the left-hand side of Eq. 2.6 (see e.g. [45],[1],[10]). In this spirit, a generic control volume discretization of the left-hand side of Eq. 2.6 reads

$$\begin{aligned} \int_{\partial\omega_i} (-\mathbf{a}\nabla u_h) \cdot \mathbf{n} \, dx &= \sum_{\gamma_k \in \partial\omega_i} \int_{\gamma_k} (-\mathbf{a}\nabla u_h) \cdot \mathbf{n} \, dx \approx \sum_{\gamma_k \in \partial\omega_i} \sum_{r \in S_F(\gamma_k)} t_{k,r} u_h^r \\ &= \sum_{\gamma_k \in \partial\omega_i} q_{h,k} \end{aligned} \tag{2.9}$$

Here, we have introduced the transmissibilities $t_{k,r}$ of face k and u_h^r denotes the discrete potential in cell r , and

$$q_{h,k} = \sum_{r \in S_F(\gamma_k)} t_{k,r} u_h^r$$

is the discrete flux over γ_k . Finally $S_F(\gamma_k)$ (subscript F indicates face) represents the support of the flux discretization over γ_k . To simplify notation, we will also write

$$q_{h,k} = T_k u_h \tag{2.10}$$

where the transmissibility matrix T_k is nonzero only for cells in $S_F(\gamma_k)$. The only assumption we make on flux discretization is that its support is local, in that $S_F(\gamma_k)$ only contains cells sharing at least one vertex with γ_k , as will be the case for most control volume methods found in practice.

In practice, considerable care is taken in choosing an appropriate discretization method that preserves structures in the equations that are of importance for the application at hand. The control volume methods by construction preserve a conservation property and that make them the preferred choice for industrial simulations of flow in porous media (among the most popular are the Eclipse simulator [37] and GEM [7]).

3 Conservation-preserving coarse spaces

In this section, we formulate a multilevel hierarchy of discretizations of Eq. 2.5 and the flux discretization (2.10). We also discuss the basis functions used to describe the coarse spaces and the application of the resulting discretization hierarchy as linear solvers.

3.1 A hierarchy of discretizations

We are concerned with multilevel preconditioners that take the same form as control volume discretization methods. Let us first consider the two-level setting, as the multilevel method follows by recursion. We associate the index $l = 0$ with the discretization level, and $l = 1$ with the coarse level; we will throughout the paper use index l and m to indicate grid levels. The subscript h_l denotes a discrete quantity on level l , and we sometimes use subscript h for h_0 to be consistent with standard notation for discretization methods, cf Section 2. Consider a coarsening procedure that leads to a coarse solution vector \mathbf{u}_{h_1} defined as the solution of

$$a_h(\boldsymbol{\psi}_{h_1} \cdot \mathbf{u}_{h_1}, \boldsymbol{\phi}_{h_1} \cdot \mathbf{e}_j) = f_h(\boldsymbol{\phi}_{h_1} \cdot \mathbf{e}_j) \quad \forall j = \{1, \dots, n_1\} \tag{3.1}$$

Here, $\boldsymbol{\psi}_{h_1}$ is an interpolation matrix, composed of individual interpolation vectors $\boldsymbol{\psi}_{h_1,l}$ $l = 1, \dots, n_1$, such that $\boldsymbol{\psi}_{h_1} = [\boldsymbol{\psi}_{h_1,1}, \dots, \boldsymbol{\psi}_{h_1,n_1}]^T$, where n_1 is the number of cells at level 1. We refer to the individual interpolation vectors as *discrete basis functions*. In analogy to control volume methods, these are not a priori uniquely defined, and our construction of discrete basis functions will be specified in Section 3.2. Similarly, $\boldsymbol{\phi}_{h_1} = [\boldsymbol{\phi}_{h_1,1}, \dots, \boldsymbol{\phi}_{h_1,n_1}]^T$ is a compression matrix, composed of individual compression vectors referred to as the *discrete test functions*. To preserve the control volume structure, the discrete test functions on level 1 are required to be piece-wise constants, which we define through a sum of test functions on level 0. The piece-wise constants on level 1 define a partitioning into coarse cells, as discussed in the next section, and thus the conservation property holds in analogy to Eq. 2.5.

Equation 3.1 can itself be interpreted as a control volume discretization of Eq. 2.1. To see this, we write the coarse discrete system in analogy to Eq. 2.6 as follows: Find \mathbf{u}_{h_1} such that

$$a_{h_1}(\mathbf{u}_{h_1}, \mathbf{e}_j) = f_{h_1}(\mathbf{e}_j) \quad \forall j = \{1, \dots, n_1\} \tag{3.2}$$

Note that, in particular, this equation can be related to the continuous problem by substitution of the definition of the discretization method as follows: Find \mathbf{u}_{h_1} such that

$$a_h([\boldsymbol{\psi}_{h_1}] \cdot \mathbf{u}_{h_1}, [\boldsymbol{\phi}_{h_1}] \cdot \mathbf{e}_j) = f([\boldsymbol{\phi}_{h_1}] \cdot \mathbf{e}_j) \quad \forall j = \{1, \dots, n_1\} \tag{3.3}$$

The expression in square brackets defines the composite basis functions of the discretization method induced by the composition of the fine-scale discretization and the algebraic coarsening.

The multilevel setting now follows by recursion. We index the levels $l = 0, \dots, N$, where $l = 0$ denotes the discretization, and $l \geq 1$ are coarse levels of the preconditioner. The number of degrees of freedom on each level is denoted n_l . Then, ψ_{h_0} is a vector of length $n_0 = n$, while ψ_{h_l} $l \geq 1$ are n_{l-1} by n_l matrices defining a hierarchy of coarse approximations. Furthermore, let the vector Ψ_{h_l} be defined as the induced basis function $\Psi_{h_l} = \psi_{h_0} \cdot \dots \cdot \psi_{h_l}$. Let similar notation hold for the test functions with $\Phi_{h_l} = \phi_{h_0} \cdot \dots \cdot \phi_{h_l}$. Then \mathbf{u}_{h_l} satisfies, for $l = 0, \dots, N$, the discrete system of equations given in terms of the continuous operator as

$$a(\Psi_{h_l} \cdot \mathbf{u}_{h_l}, \Phi_{h_l} \cdot \mathbf{e}_j) = f(\Phi_{h_l} \cdot \mathbf{e}_j) \quad \forall j = 1, \dots, n_l \tag{3.4}$$

We denote the continuous approximation associated with discretization level l as $u_{h_l} \equiv \Psi_{h_l} \cdot \mathbf{u}_{h_l}$.

3.2 Coarse basis functions and flux discretizations

We identify two natural special cases of coarse basis functions. For nested grids, where simple geometric structures (triangles, squares, simplexes, etc.) can be identified for all cells on all levels, it is possible to use the lowest-order finite element shape functions on all levels. This leads to the so-called control volume finite element (CVFE) discretization on level 0 ([4],[42]), and the coarse space problems can be similarly identified as CVFE discretizations. Although the CVFE discretization has not become popular for applications in porous media due to poor compatibility with discontinuous coefficients, it is relatively easy to analyze, and can be considered as a model discretization.

The more natural coarse spaces, from the perspective of an algebraic construction and the relationship to multiscale methods and upscaling, is to create coarse spaces that are in a certain sense a PDE-harmonic coarsening of the finer scale ([29]). For reference, a short description of the construction is provided here with emphasis on details of importance for the multilevel reconstruction presented in the next section. Further information on the construction of the coarse spaces can be found in ([14],[17] [9],[36]).

3.2.1 Coarse primal and dual grids

Given a grid on level $l-1$ (this can be the finest level, or any coarsening thereof) with cells ω_i^{l-1} , we first define a coarse grid on level l . Define index sets C_i^l such that a coarse cell ω_i^l is defined as

$$\omega_i^l = \bigcup_{r \in C_i^l} \omega_r^{l-1}$$

The coarse cells ω_i^l should be simply connected unions of fine cells ω_j^{l-1} , for illustrations see Fig. 1, and also Section 6.2.2. To define a post-processing method that is applicable to general grids, mappings between cells and faces on different levels are needed. We refer to the set of subcells of ω_i^l on level $m < l$ by $C^m(\omega_i^l)$. Similarly, the face between two coarse cells $\gamma_k^l = \partial\omega_{ij}^l$ consists of faces on level $l-1$, and we define E_k^l such that

$$\gamma_k^l = \bigcup_{s \in E_k^l} \gamma_s^{l-1}$$

The set of sub-faces of γ_k^l on level $m < l$ are denoted $G^m(\gamma_k^l)$ Finally, we define

$$H^m(\omega_i^l) = \bigcup_{\omega_r^m \in C^m(\omega_i^l)} \partial\omega_r^m$$

as the faces of all those cells on level m that are subcells of cell ω_i^l .

As noted earlier, the coarse test functions are simply defined as piecewise constants on the cells ω_i^l . However, to compute basis functions ψ_{h_l} , a dual coarse grid is needed. To that end, within each coarse cell ω_i^l a fine-scale cell denoted $\omega_{V_i}^{l-1}$ is defined as the vertex of the dual grid that is associated with ω_i^l ; for simplicity we will take this as the centermost cell in ω_i^l . Denote by V^{l-1} the set of all vertex cells on level $l-1$. The coarse basis function of cell ω_i^l will be associated with this vertex cell.

Next, for each primal face $\gamma_k^l = \partial\omega_{ij}^l$ a dual coarse edge is formed by a continuous path of fine scale cells ω_j^{l-1} between the respective vertices of the dual grids. The only assumptions made on the dual edge are that it should only




Fig. 1 A Cartesian grid with fine and coarse (thick lines) cells; a finer grid is also indicated. The dual coarse grid is indicated by grey cells. The basis function centered in the black cell has support in all four surrounding dual coarse cells, and thus contribute to the flux expressions of all coarse edges shown in the figure. The coarse discretization is thus a 9-point stencil

intersect with a single primal face, and that a fine-scale cell ω_j^{l-1} can at most belong to one dual edge. For a 2D problem, the dual vertices and edges will together partition the domain in the sense that fine-scale cells that are not part of the dual vertices or edges form disjoint subdomains; see Fig. 1. If $\Omega \subset R^3$, the dual grid is completed by defining dual faces that connect the dual edges so that each dual face intersects a single primal edge.

3.2.2 Coarse basis functions

We next describe the computation of basis functions that resolve heterogeneities in the coefficient \mathbf{a} . We focus on the basis function of ω_i^l , and for convenience, we drop a subscript and use the notation $\psi_i = \psi_{h_l,i}$ in this section. The value of the basis function is set to unity in the dual vertex within the coarse cell l , and to zero in all other vertex cells, that is

$$\psi_i(\omega_j^{l-1}) = \delta_{ij} \quad \omega_j^{l-1} \in V^{l-1} \tag{3.5}$$

Equation 3.5 serves as boundary conditions for the confinement of Eq. 2.5 to each dual edge that has $\omega_{V_i}^{l-1}$ as a vertex. For this problem flow that is normal to the edge is neglected that is, a 1D equation that reads

$$-\nabla \cdot (\mathbf{a}_{1D} \nabla u) = 0 \tag{3.6}$$

where the one-dimensional, tangential component of the coefficient field, \mathbf{a}_{1D} is discretized and solved on each dual edge. This construction can be achieved algebraically based on the discretization on level $l - 1$ ([35]). On dual edges that do not have ω_j^{l-1} as a vertex, the basis function is set to zero.

Next, the computed values of the basis function provide boundary conditions for local problems in 2D, these are either on interior sub-domains (if $\Omega \subset R^2$) or on dual faces ($\Omega \subset R^d, d > 2$). In the former case the restriction of Eq. 2.5 to the local domain is solved while for the latter only the component of \mathbf{a} tangential to the 2D surface is considered, i.e. we solve (3.6) with \mathbf{a}_{1D} replaced by \mathbf{a}_{2D} . In both cases, the previously computed values of the basis functions on the dual edges serve as boundary conditions for the 2D problem. If $\Omega \subset R^3$, the computation is finalized by considering the interior subdomains. Note that the resulting basis function has support in all dual coarse cells that overlap with ω_i^l , see Fig. 1. The support of the basis function of ω_i^l , expressed in terms of cells on level $m < l$ is denoted $S_C^m(\omega_i^l)$.

The computation of basis functions are computed for all coarse cells ω_i^l completes the definition of the coarse discretization (3.4) on level l .

3.2.3 Coarse flux expressions

For the post-processing that will be introduced in Section 4, and for the goal of constructing an integrated framework

for linear solvers, discretizations and up- and downscaling, we also require coarse equivalents of the flux expression (2.9). To be consistent with (3.3), the flux discretizations are formulated in terms of the computed basis functions. We first introduce a downscaling of the flux, and analogous to Eq. 2.10, we write the discrete flux of face $\gamma_k^0, q_{h_0,k}$, in terms of potentials on level l as

$$q_{h_0,k} = T_k^0 (\psi_{h_1} \cdot \dots \cdot \psi_{h_l}) \mathbf{u}_{h_l} = T_k^l \mathbf{u}_{h_l} \tag{3.7}$$

Here $(\psi_{h_1} \cdot \dots \cdot \psi_{h_l}) \mathbf{u}_{h_l}$ gives a mapping of potentials from level l to the discretization level, and we have defined the transmissibility matrix T_k^l to simplify the expression. Fluxes are extensive quantities, thus the flux over $\gamma_k^l > 0$ in terms of potentials on level $l \geq m$ is given as

$$q_{h_m,k} = \sum_{\gamma_s \in E_k^m} q_{h_0,s} = \sum_{\gamma_s \in E_k^m} T_s^l \mathbf{u}_{h_l}$$

As discussed in Section 2, the support $S_F(\gamma_k)$ is limited to cells that share at least one vertex with γ_k^0 . We similarly define $S_F^l(\gamma_k^m)$ as the support of the flux discretization of γ_k^m in terms of potentials on level $l > m$, that is, cells on level l which give a nonzero contribution to (3.8). We will refer to faces with common support as belonging to the same *interaction regions*, in accordance with discretization method terminology [1]. Due to the localization of the basis functions, $S_F^l(\gamma_k^m)$ will have a similar size as the fine-scale discretization. For instance, for a uniform coarsening of a Cartesian grid the coarse support will in general contain 6 and 18 points in 2D and 3D, respectively.

3.3 Application as linear solvers

The hierarchy of coarse spaces described in this section can provide approximations to a linear system of equations by computing coarse-scale corrections to any solution, and this naturally leads to a framework for iterative linear solvers. Although much effort has been put into constructing coarse spaces that represent the essential features of the fine-scale problem, e.g., [36], any coarse approximation will be of low rank, and there will unavoidably be cases where the coarse solution space is orthogonal to significant components of the solutions, thus any coarse solution is unsatisfactory even in the setting of an inexact solver. This realization has led to a refocus from pure multiscale methods, which are confined to a single iteration, to iterative solvers. Given an initial guess u^0 , an iterative scheme for (2.8) defines a sequence of solutions $\{u^1, u^2, \dots\}$ by

$$u^{n+1} = u^n + B (b - Au^n) = u^n + Br^n$$

where $B \approx A^{-1}$ in our case represents the coarse scale approximations, $r^n = b - Au^n$ is the residual, and we have for simplicity dropped the bold face notation.

A crucial component of a coarse scale correction strategy is the treatment of the residual. Two options are natural: A treatment that is consistent with the Schur complement formulation of the multilevel method, or a classical smoothing of the residual. The first method, which we will refer to as B_{SC} is inferior as a smoother, and should therefore not be applied to reduce the size of the residual. However, B_{SC} by construction preserves the conservation structure of the discrete problem. Therefore a single application by B_{SC} to the residual of any approximate solution u^n , which need not be constructed from the multilevel hierarchy, ensures that a conservative flux field can be associated with the resulting u^{n+1} by the postprocessing technique introduced in the next section. The full derivation of the Schur complement formulation is in the interest of space omitted herein, the reader is referred to previous literature for details [28]. For the second method, an appropriate smoother S_{hl}^v such as Jacobi or Gauss-Seidel, is applied on all grid levels [12]. The projection and restriction operators for the residual are defined by the coarse interpolation matrices ψ_{hi} and their transpose, respectively. The approximation B_{SM} can either be applied as a stand-alone solver, or accelerated by a Krylov subspace method.

When the approach discussed in this section is implemented in a two-level setting with aggressive coarsening, it leads to sub-structuring methods and coarse-scale control volume finite element methods [28]. In the multilevel setting, this has not previously been discussed, primarily because of the difficulties associated with reconstructing conservative fluxes from an inexact potential in a multilevel setting. In the next section, we give a resolution to this issue.

4 Multilevel conservative reconstruction

Our goal is to formulate solvers that can be terminated at any time in the iterative process, and still retain the conservation structure of the equations. That is, we are seeking solvers that produce approximation pairs $(u_h^*, \mathbf{q}_h^*) \approx (u_h, \mathbf{q}_h)$ such that (both in the discrete and continuous sense)

$$\nabla \cdot \mathbf{q}_h^* = f, \quad \mathbf{q}_h^* \approx -\mathbf{a} \nabla u_h^*$$

This requires post-processing the approximate solution obtained by the multilevel method, and will be efficient if the post-processing is local in the hierarchical sense. It is important to note that the approximated potential and flux cannot be consistent, in the sense that $\mathbf{q}_h^* = -\mathbf{a} \nabla u_h^*$, as by uniqueness this would represent the exact solution.

A property of control volume discretizations which follows from the flux $q_{h,k}$ in (2.9) being defined uniquely for each cell face is that the discrete set of equations given in (2.6) is equivalent to the compatibility condition

for Neumann boundary value problems. Therefore, from the solution u_h^* , discrete fluxes \mathbf{q}_h^* are given by Eq. 2.10, and these discrete face fluxes can be interpreted as the Neumann data for a local problem within each cell. This boundary-to-interior map provides a definition of a continuous approximate flux \mathbf{q}^* , which satisfies the conservation structure. This construction is used in practice when the elliptic equation (2.1) is coupled to hyperbolic transport equations [16], but also has theoretical relevance which is exploited in error analysis, as we will see in the following section. The key property to be observed is that it is the conservative structure of the method that allows for a completely local post-processing of the discrete solution onto a continuous field.

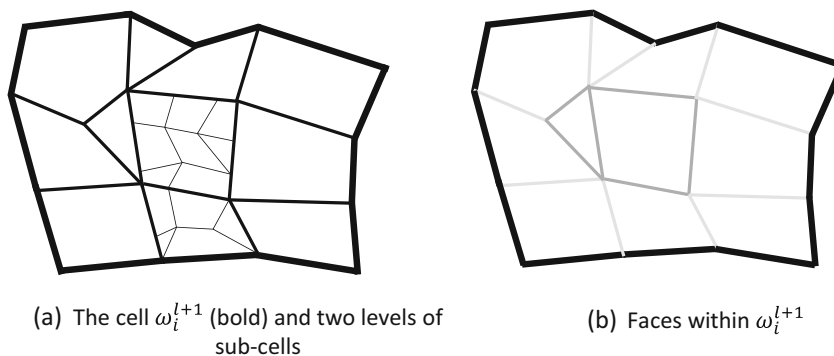
The possibility of post-processing control volume fluxes is also the key aspect that makes the control volume multi-scale methods attractive [17]. Thus, in a two-scale setting, a coarse control volume method can be defined as given in Section 3, and if the coarse linear system is solved exactly, the projection of coarse-scale fluxes to the fine scale satisfies the compatibility condition locally within each coarse cell. Thus the solution can be post-processed to provide conservative fluxes on the finer scale. However, a key ingredient is the access to Neumann boundary discretizations, which for control volume methods typically requires special treatment. These are directly available at the finest level from the control volume discretization, and as such pose no problems for two-level methods, but are not directly available in the coarse discretization. Therefore, multilevel implementations have so far mostly considered one level of reconstruction [35], although a multilevel method for structured grids was recently introduced in [23]. In the following, we will show how the multilevel post-processing can be performed for general grids, and in a way that is consistent with the underlying discretization.

4.1 Multilevel topology

In this section, we detail the hierarchical multilevel post-processing. It is applicable to general unstructured grids, as may arise from algebraic coarsening strategies. To describe the flux reconstruction on level l we assume that conservative fluxes for all faces γ_k^m , ($m > l$) are known together with their partitioning into subfluxes on level l . Note that this calls for an exact solver to applied at the coarsest level N . We will show how the post-processing can be performed locally within a single cell on level $l + 1$, ω_i^{l+1} , illustrated in Fig. 2a. To this end we formulate a linear system describing conservation for subcells of ω_i^{l+1} on level l and with flux boundary conditions projected from level $l + 1$.

Since we only consider a single cell on level $l + 1$, we simplify notation by letting $C^l = C^l(\omega_i^{l+1})$ for the subcells

Fig. 2 Left: A multilevel grid with a coarse cell (bold) with two finer levels of grids indicated. Right: The splitting of the faces into the boundary faces γ_B (black), faces that are adjacent to the boundary γ_V (dark grey) and the truly interior faces γ_I (light grey). a The cell ω_i^{l+1} (bold) and two levels of sub-cells. b Faces within ω_i^{l+1}



of ω_i^{l+1} on level l , and write $H^l = H^l(\omega_i^{l+1})$ for the union of all faces on level l that are either inside or at the boundary of ω_i^{l+1} . We will introduce a splitting of H^l into faces that are respectively on the boundary of ω_i^{l+1} , in the vicinity (in terms of the flux discretization) to the boundary and in the interior, see Fig. 2b. We denote the three sets γ_B, γ_V and γ_I respectively, and define them as

$$\begin{aligned} \gamma_B &= \{\gamma_k \in H^l: \text{if } \gamma_k = \partial\omega_{ij}, \text{ then } \omega_i \notin C^l \text{ or } \omega_j \notin C^l\} \\ \gamma_V &= \{\gamma_k \in H^l: \text{if } \gamma_k = \partial\omega_{ij}, \text{ then } \{\omega_i, \omega_j\} \in C^l \text{ but } S_F^l(\gamma_k) \not\subset C^l\} \\ \gamma_I &= \{\gamma_k \in H^l: \text{if } \gamma_k = \partial\omega_{ij}, \text{ then } \{\omega_i, \omega_j\} \in C^l \text{ and } S_F^l(\gamma_k) \subset C^l\} \end{aligned}$$

4.2 Formulation of a local linear system

The postprocessing is performed by solving the local linear system that governs conservation for the cells in C^l with local sources and the prescribed boundary conditions. The flux expressions for the faces in γ_I are straightforwardly assembled into the target local linear system. However, the assignment of boundary conditions to faces that are in the interior of the global domain, and thus have no boundary discretization available, requires some care. The flux expressions for the faces in γ_B are replaced by the known boundary conditions, while faces in γ_V pose greater difficulties since the flux stencils have support outside C^l and are thus incompatible with a local post-processing method. The flux expressions for faces in γ_V must thus be modified so that dependencies on exterior cells are replaced by information contained in the known fluxes over γ_B . On level this can easily be done, since on the finest level the discretization of flux boundary conditions is available. On coarser levels this is however not the case and the elimination of the exterior cells becomes more involved. Indeed, this elimination represents the key challenge to the formulation of fully multilevel multiscale finite volume methods.

To arrive at a local system we make two key observations: Firstly the modification of the flux expression for

faces in γ_V is naturally performed in the interaction region framework, since this was the essential tool for computing the original flux expression. Consider a single interaction region as illustrated in Fig. 3, and assume it contains cells that are not in C^l . Denote the cells in the region interior, C_I , and exterior, C_E , depending on whether they are members of C^l or not. The faces in the region, denoted γ_R , can be split into three categories:

$$\begin{aligned} \gamma_{R,I} &= \{\gamma_k \in \gamma_R: \text{if } \gamma_k = \partial\omega_{ij}, \text{ then } \{\omega_i, \omega_j\} \in C_I\} \\ \gamma_{R,B} &= \{\gamma_k \in \gamma_R: \text{if } \gamma_k = \partial\omega_{ij}, \text{ then } \omega_i \in C_I, \omega_j \in C_E\} \\ \gamma_{R,E} &= \{\gamma_k \in \gamma_R: \text{if } \gamma_k = \partial\omega_{ij}, \text{ then } \{\omega_i, \omega_j\} \in C_E\} \end{aligned}$$

The flux expression for $\gamma_{R,I}$ must be modified so that the dependencies on potentials in the external cells C_E are replaced by the known fluxes over faces in $\gamma_{R,B}$. This replacement gives a twofold modification to the flux computation: The flux expression is altered, and the prescribed




Fig. 3 An interaction region with cells grouped into exterior (\mathbf{u}_E) and interior (\mathbf{u}_I) cells. The prescribed boundary fluxes (\mathbf{q}_B , thick black arrows) are split into their sub-fluxes on a finer level (illustrated by thin black arrows). The uppermost face is in γ_V , and the sub-fluxes (\mathbf{q}_I grey arrows) over this induced by the boundary conditions are computed in the elimination of the exterior cells \mathbf{u}_E

boundary conditions set up an induced flux over the faces in $\gamma_{R,I}$, as illustrated by grey arrows in Fig. 3. This induced flux comes in addition to the flux driven by the potential in the cells C_I as computed by the local linear system. The challenge is to compute both these modifications for general grid configurations in a way that is consistent with the multilevel discretization.

As is explained below, if there are as many cells in C_E to be eliminated as there are constraints in the form of known boundary fluxes, the elimination is carried out by solving a linear system. However, as illustrated in Fig. 3, there might be more exterior cells than there are constraints, and more equations are needed to close the system. To provide the extra equations necessary, we utilize the following aspect of the multilevel approach: The boundary conditions are available not only on level l ; but as discussed in Section 3.2.3 the sub-fluxes can also be computed with higher resolution at level $0, \dots, l - 1$. Moreover, if $\gamma_k^l \in \gamma_{R,I}$, then $E^m(\gamma_k^l) \in \partial\omega_i^{l-1}$ for some i . That is, the sub-faces of γ_k^l are on the boundary of all post-processing domains on the finer levels, and when the post-processing is performed on these finer levels, the boundary conditions will be assumed known. In other words, the modified transmissibilities and the induced fluxes must be computed on all grid levels, and the natural way to do this is to carry out the elimination on the finest level, and then sum the expressions to get the modifications on coarser levels, reminiscent of the computations in Section 3.2.3. In addition to rendering modified transmissibilities and induced fluxes, this strategy will both provide the necessary equations to eliminate the exterior unknowns.

4.3 Elimination in an interaction region

Again, consider the single interaction region as illustrated in Fig. 3. Let u_I and u_E be the potential in cells in C_I and C_E , respectively. Moreover, let n_b be the number of boundary faces on level l in the interaction region and let n_e be the number of cells in the interaction region that are exterior to the local problem (2 and 3 in Fig. 3, respectively). Motivated by the above discussion we consider flux expressions on level l via the mapping T_k^l , and for notational convenience we drop the superscript identifying the grid level. We further split the transmissibility matrix into the contribution from cells in C_I and C_E and combine T_k for all faces in $\gamma_{R,B}$ and $\gamma_{R,I}$ to write $T_{B,I}$ and $T_{B,E}$, and similarly $T_{I,I}$ and $T_{I,E}$. The exterior faces in $\gamma_{R,E}$ are not part of elimination. Furthermore, the face fluxes on level 0 within the region are divided into known boundary fluxes, q_B and the interior fluxes q_I that should be computed. With this notation, (3.7) combined for all cells in $\gamma_{R,B}$ and $\gamma_{R,I}$ can be written as

$$\begin{pmatrix} T_{B,I} & T_{B,E} \\ T_{I,I} & T_{I,E} \end{pmatrix} \begin{pmatrix} u_I \\ u_E \end{pmatrix} = \begin{pmatrix} q_B \\ q_I \end{pmatrix} \tag{4.1}$$

When q_B is known, u_E can be eliminated from the first equation in Eq. 4.1. However, (4.1) describes conservation on level l while our primary aim is to produce conservative fluxes on level l , and the linear system should be reformulated accordingly. To that end, flux expressions for sub-faces belonging to the same face on level l are added. This gives n_b equations to eliminate u_E . If $n_b = n_e$, the system is closed, and no more equations are needed to carry out the elimination. We note that this case includes two-dimensional Cartesian grids, thus Cartesian grids with uniform coarsening are simpler from the perspective of hierarchical flux reconstruction than general grids. Furthermore, even though u_E in this case can be eliminated without the involvement of sub-fluxes, fluxes on level 0 must still be considered for conservation on the finer levels.

The requirement of conservation on level l produces a sufficient number of equations for some grids. However if $n_b < n_e$, as is the case for unstructured grids as well as for certain interaction regions on 3D Cartesian grids, more equations are needed to close the system. To define these, we make the choice of using higher moments of flux conservation over the edge. Let \mathbf{x}_k^f be the coordinate of the fine-scale face k , and let \mathbf{x}_0 be the midpoint of the interaction region. Moments of conservation are computed as weighted sums of the rows in Eq. 4.1 corresponding to boundary faces, with a row weighting of $(\mathbf{x}_k^f - \mathbf{x}_0)^\alpha$ computed according to the coordinate of the corresponding sub-face, and $\alpha = 1, 2, \dots$ incremented by one for each sum taken. A sufficient number of discrete moment integrals are computed to close the system with n_e equations in total.

The resulting linear system, containing both conservation on level l , and possibly also higher order moment integrals, reads

$$\hat{T}_{B,I}u_I + \hat{T}_{B,E}u_E = q_B$$

There are now as many equations for the boundary fluxes as there are exterior cells. The equation is solved for u_E and inserted into the flux expression for the interior faces on level l , which becomes

$$q_I = \left(T_{I,I} - T_{I,E} \hat{T}_{B,E}^{-1} \hat{T}_{B,I} \right) u_I + T_{I,E} \hat{T}_{B,E}^{-1} q_B \tag{4.2}$$

We recognize the term $T_{I,E} \hat{T}_{B,E}^{-1} \hat{T}_{B,I}$ as the modification introduced to the transmissibilities for faces in $\gamma_{R,I}$, and $T_{I,E} \hat{T}_{B,E}^{-1} q_B$ as the flux induced by the Neumann boundary conditions. Both these modifications are currently formulated for faces on level 0, but they can be added to be valid for level l , or any level in between.

We note that a crucial assumption of the methodology is that $\hat{T}_{B,E}$ is invertible. We have no proof that this will always be the case; however during extensive numerical tests this has always been the case in practice.

The elimination of exterior unknowns is carried out consistently with the conservative fluxes over the boundaries, and if necessary, consistent with moment integrals of these fluxes. The viability of the elimination hinges on the number of exterior cells being less than the number of sub-faces on the finest grid. This imposes a lower bound on the coarsening ratios admissible. A practical example of this is that for a multilevel 3D Cartesian grid hierarchy, the coarsening ratio from level 0 to 1 must be at least 5 for the post-processing to work.

4.4 Assembly of the linear system

The above elimination is carried out for all elements in γ_V so that the resulting flux expressions are dependent only on cells in $C^l(\omega_i^{l+1})$. These can be combined with the unmodified flux expressions for γ_I to form a local linear system for ω_i^{l+1} on the form

$$D_i^{l+1} F_{\gamma_V \cup \gamma_I} u_C = f_C + D_i^{l+1} q_N + D_i^{l+1} q_{\text{ind}}$$

Here, D_i^{l+1} is the local divergence operator, $F_{\gamma_V \cup \gamma_I}$ are the flux expressions for all interior faces, f_C are the sources within D_i^{l+1} , q_N are the imposed Neumann boundary conditions, and q_{ind} are the induced fluxes caused by the boundary conditions. The linear system is solved for u_C , and fluxes over the faces as well as over all sub-faces are computed from (4.2). These fluxes are conservative, and may serve as boundary conditions for level $l - 1$. This construction thus allows for a local, hierarchical post-processing of the inexact solution to yield a conservative flux field. The local and hierarchical aspect is important, since it ensures that the post-processing has a computational cost which scales no worse than the cost of the preconditioner itself. Furthermore, the local aspect of the reconstruction also guarantees that there is no regional bias in the post-processed fluxes, as might result from algorithms based on a sweep over the domain.

5 Error bounds

Iterative linear solvers commonly use reduction of the residual error as a stopping criterion. This is reasonable when an “essentially exact” solution is sought. Our objective is different, in that we seek approximate solutions that may potentially have significant deviation from the exact solution to the linear system of equations. In this context, it is necessary to have access to alternative error controls.

5.1 A posteriori error bound on fluxes

To analyze the error in the approximate fluxes (or equivalently, the H^1 semi-norm of the approximate potentials), we

turn to the framework of guaranteed error bounds [32]. In particular, we recall the abstract a posteriori error bound for problems with Dirichlet boundaries $\Gamma_D = \partial U$ [18]:

$$\|\mathbf{a}^{-1/2} (\mathbf{q} - \mathbf{q}_h^*)\| \leq \inf_{s \in H_1^1} \|\mathbf{a}^{-1/2} (\mathbf{q}_h^* + \mathbf{a} \nabla s)\| + \sup_{\substack{\phi \in H_1^1 \\ \|\phi\|=1}} (\nabla \cdot (\mathbf{q} - \mathbf{q}_h^*) \cdot \phi) \tag{5.1}$$

Here the norms are the usual L_2 norms, and the space H^1 refers to the space of H_1 functions satisfying the boundary data. This inequality bounds a weighted error measure for the flux. To proceed, we note that by construction the multilevel post-processed flux has exactly the same divergence as the discrete flux computed from the exact solution of the discrete problem (2.5), and the last term of Eq. 5.1 evaluates to 0. Thus, it remains only to approximate the first right-hand-side term. This term is a measure of the distance between the post-processed flux and a solution to the flux law, for some potential s . An accurate error bound can be obtained by solving the minimization problem over all s , leading to a dual problem. However, we will be satisfied with a coarser estimate at the benefit of a locally computable error expression. Any choice of s will lead to a bound, and since we expect $\mathbf{q}_h^* \approx -\mathbf{a} \nabla u_h^*$, we expect u_h^* to be a candidate for s which is close to the optimal choice. From the triangle inequality

$$\|\mathbf{a}^{-1/2} (\mathbf{q} - \mathbf{q}_h^*)\| \leq \inf_{s \in H_1^1} \|\mathbf{a}^{-1/2} (\mathbf{q}_h^* + \mathbf{a} \nabla s)\| \leq \|\mathbf{a}^{-1/2} (\mathbf{q}_h^* + \mathbf{a} \nabla u_h^*)\| + \inf_{s \in H_1^1} \|\mathbf{a}^{1/2} \nabla (u_h^* - s)\|$$

The inequality represents a weighted H^1 seminorm for the error potential u_h^* . Following [18], we identify the two terms on the righthand side as the error from the solver and the discretization, respectively, and thus write

$$\|\mathbf{a}^{-1/2} (\mathbf{q} - \mathbf{q}_h^*)\| \leq E_s + E_h \tag{5.2}$$

The error from the linear solver E_s is immediately computable, since both the approximate solution u_h^* and an approximate flux \mathbf{q}_h^* are available. The discretization error can again be bounded from above by choosing any s in H^1 . The simplest choice is the Oswald interpolation I_O based of u_h^* on element-wise N -linear polynomials, defined by the average potential on corners [18]. This leads to

$$E_h \leq \|\mathbf{a}^{1/2} \nabla (u_h^* - I_O(u_h^*))\|$$

The two components of the error thus allow us to consider when the linear solver is more accurate than the error bound for the underlying discretization, after which it is debatable whether to invest additional computational effort to improve the solution to the linear system.

We note that the error can also be estimated directly from the post-processed flux approximation without splitting the

evaluation into two components. This has the advantage of avoiding the triangle inequality, providing a slightly sharper estimate at the cost of losing the insight into the structure of the error.

5.2 An a posteriori estimate by backward error analysis

An alternative perspective, of particular relevance for applications, is gained by recalling that the coefficient \mathbf{a} is itself often given to only a very low accuracy [15]. Thus, considering the approximation pair (u_h^*, \mathbf{q}_h^*) , it is possible to ask if there exists a (symmetric) tensor field \mathbf{a}^* such that

$$\mathbf{q}_h^* = -\mathbf{a}^* \nabla u_h^*$$

in the discrete sense. In general such a tensor field exists almost everywhere, but is not unique. Furthermore, the family of tensor fields \mathbf{a}^* is easily computable. The pair u_h^* and \mathbf{q}_h^* can then be interpreted as the exact solution to an approximate problem (a similar idea is discussed in [28]). This view is similar to the approach of seeking modified equations to which a numerical scheme is exact, which is a common analysis tool for hyperbolic conservation laws (see e.g. [24]). From the perspective of applications, the quality of the solution can then be judged by the distance of \mathbf{a} from the set of tensor fields \mathbf{a}^* , relative to the uncertainty in \mathbf{a} itself [32]. We may also consider the quality of the field \mathbf{a}^* itself. The physical problem will always have a symmetric positive definite coefficient field \mathbf{a} , and this property can be used as a measure of the quality of \mathbf{a}^* . Since there will only exist positive definite \mathbf{a}^* if $\mathbf{q}_h^* \cdot \nabla u_h^* < 0$ everywhere, we define γ_{SE} as the set of faces with sign errors in the flux, that is

$$\gamma_k \in \gamma_{SE} \quad \text{if} \quad \mathbf{q}_{h,k}^* \cdot \nabla u_{h,k}^* > 0$$

We then measure the relative size of the associated with unphysical values of \mathbf{a}^* by

$$E_t = \left(\sum_{\gamma_k \in \gamma_{SE}} |\mathbf{q}_{h,k}^*|^2 \right)^{1/2} / \|\mathbf{q}_h^*\|$$

5.3 A priori error estimates

We support the concept of using the multilevel hierarchy as a solver by highlighting the relationships to classical multilevel methods. This allows us to discuss the methodology in terms of a priori error estimates, at least for simple model problems.

Coarse spaces as defined in Section 3 can be combined to a multilevel solver, following the approach of multigrid methods [44]. However, standard analysis for these types of problems is not straightforward, since the discrete method defined by Eq. 2.6 is not symmetric, and thus does not

introduce a natural energy norm for the problem [46]. Furthermore, a monotone coarse problem with local support (at any level) can in general not be guaranteed to exist [27],[29]. While the subspace correction framework for multilevel methods has been extended in certain cases to finite element discretizations of parabolic equations, which leads to non-symmetric discrete equations (see e.g., [33],[20]), we have not attempted to adapt the subspace correction framework to the current setting.

The convergence of the multilevel method as a solver does not require symmetry; however, the multilevel approximations need to be convergent in the sense of discretizations [12]. We will therefore assume that the hierarchy of coarse problems preserves an approximation property of the form

$$\|u_{h_l} - u_{h_{l+1}}\|_X \leq c h_{l+1}^\beta \|u_{h_l}\|_Y$$

for certain norms given by spaces X and Y , and where c is a generic constant and $\beta > 0$.

Following the classical work of Hackbusch [12], it is then well known that a convergent multilevel method can be constructed if the hierarchy of coarse spaces is combined with appropriate smoothers $S_{h_l}^v$ (e.g., Jacobi or Gauss-Seidel), satisfying

$$\|S_{h_l}^v u_{h_l} - u_{h_l}\|_Y \leq \frac{c}{v^\gamma} h_l^{-\beta} \|u_{h_l}\|_X$$

This result holds without any further assumptions on the discrete system, such as symmetry or monotonicity; however, it is weak in the sense that the constants c are not specified, and that the number of smoothing steps v may be large.

In general, convergence of control volume methods is nontrivial, and only conditional results on certain geometries are known for specific choices of basis functions (see e.g. [22]). Thus, the Hackbusch argument for convergence of the multilevel methods can only be verified for special cases. An exception is multilevel CVFE discretizations. The control volume finite elements were introduced in Section 3 and are obtained for discretization of triangular elements by choosing piece-wise linear basis functions. Furthermore, the algebraic approach herein coincides with the CVFE method for constant coefficient fields \mathbf{a} . It is known that the approximation properties for sufficiently smooth problems are of the form [4]

$$|u_{h_l} - u_{h_{l+1}}|_1 \leq c h_{l+1} \|u_{h_l}\|_2$$

Multilevel convergence is thus established by choosing any standard smoother. In the numerical examples below, we

will verify the convergence of the multilevel solver using a Gauss-Seidel smoother.

6 Numerical validation

In this section, we investigate the performance of the multilevel control volume method. The purpose of these examples is twofold: Firstly, we illustrate the utility of obtaining conservative fluxes even for inexact solutions. Secondly, we stress the importance of working in a convergent framework, wherein a high-accuracy solution can be obtained if this is deemed necessary. To illustrate the utility of the presented framework it is of interest to study both two-level methods with aggressive coarsening and multilevel methods with minimal coarsening ratios. We apply a Gauss-Seidel (GS) smoother in all simulations. The number of both pre- and post-smoothing steps in GS is set equal to the coarsening ratio in each dimension unless otherwise stated.

Most test cases involve Cartesian grids on the unit cube in two and three spatial dimensions; a hexagon-dominated grid on the unit square is considered in Section 6.2.2. In all examples the coefficient \mathbf{a} is isotropic and aligned with the coordinate axes. As noted in Sections 2 and 3, the coarse hierarchy makes no assumptions on the particular form of the fine-scale discretization scheme, except it should be a control volume scheme. In this work we discretize (2.1) on the finest grid using the MPFA O(0)-method [1],[8], which for aligned coefficient \mathbf{a} reduces to a two-point flux stencil for the Cartesian grids. The boundary conditions are no-flow for the unstructured grid, while periodic boundaries are considered otherwise.

The numerical tests are divided into two parts. Section 6.1 consider multilevel methods and establish their convergence both for media with heterogeneities that are well represented on the coarse grid and for random heterogeneities. In Section 6.2 two-level methods and the acceleration of these by GMRES are investigated.

To assess the performance of the methods, both the L_2 relative error of the potential as compared to the fine-scale solution, as well as the post-processed flux error given by E_S in Eq. 5.2 is considered. For the flux error to be meaningful, a SC correction and flux post-processing is performed after the final GS iterations, and the resulting potential and flux errors are reported

For benchmarking, the linear systems are also solved with smoothed aggregation algebraic multigrid (hereafter AMG) solver [41] as implemented in [3] and the AMG residual is reported. We have also tried the classical AMG method [40] with similar results as the ones reported here. Note that since an approximated solution from AMG does not preserve the conservation structure, it is not meaningful to consider flux errors for the AMG solution. It would

have been possible to balance the nonconservative AMG fluxes as discussed in Section 3.3, but this approach is not pursued herein. Giving a fair comparison of the performance of different solvers is as always a challenge. When a multilevel control volume (hereafter CV) solver is used, the number of iterations is the natural measure of cost. In these cases the solver setups are made as equal as possible, thus in Section 6.1 AMG is applied as a stand-alone solver, with the number of pre- and post-smoothing steps equal for the two solvers. When CV methods with only two levels are considered in Section 6.2, the comparison is not that straightforward, as the relative cost of an iteration may not be comparable. Nevertheless we choose to use the number of iterations as a proxy for computational cost. In Section 6.2, the number of pre- and post-smoothing steps for AMG is set to 3, and the iterations are accelerated by GMRES.

6.1 Multilevel methods

We first consider a multilevel solver with minimal coarsening ratios applied to get a multigrid style method. The iterations are implemented according to [12] and both V - and W -cycles are applied. This setting is ideal for a residual-smoothing strategy, and we therefore report results obtained by iterations with the GS method followed by a final SC flux equilibration. The right hand side is given by

$$f = \sin(3\pi x) \cos(3\pi y) \cos(3\pi z)$$

and we set $z = 0$ for 2D grids.

6.1.1 2D: Homogeneous and structured heterogeneous cases

The first test case involves a homogeneous medium in 2D with 3^6 cells in each dimension and five levels of coarsening. The results obtained are shown in Fig. 4. For the CV solver, the V -cycle iterations exhibit a smooth convergence, with asymptotic rates of 0.065 and 0.075 for the potential and flux, respectively. The AMG solver converges somewhat slower. Considering the flux errors and the error estimates, we note that E_S gives a fairly good estimate of the actual flux error, and that after four iterations the estimated flux error is smaller than the discretization error, thus the iterations could have been terminated according to the discussion in Section 5.1, and a conservative flux field could have been reconstructed. This is a crucial difference to the AMG solver, or indeed any standard linear solvers, wherein no special care is taken to preserve the conservation structure. As a result, an approximated solution from a standard linear solver will yield a conservation error that, due to

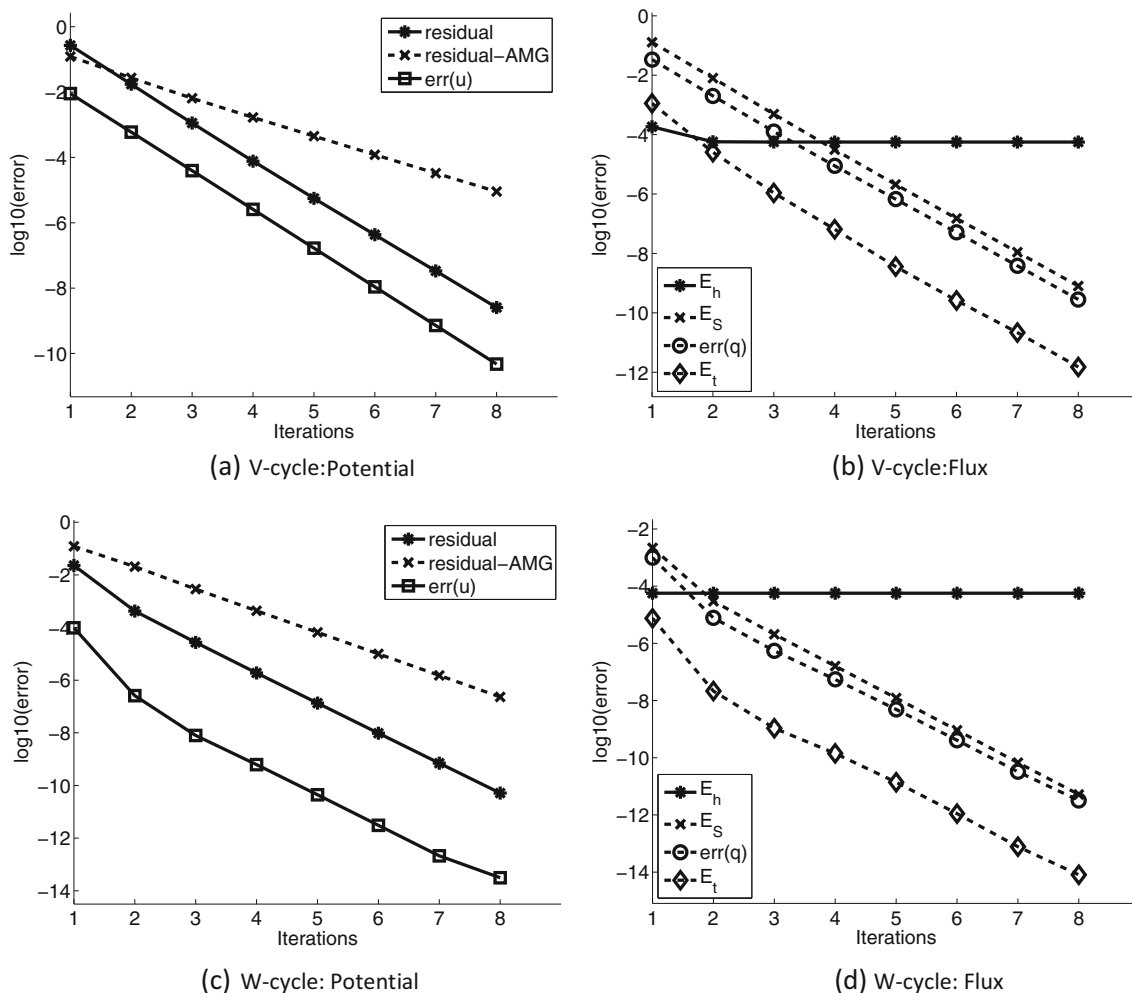


Fig. 4 Convergence of potential and flux errors, and estimators for flux and discretization errors on a homogeneous 2D problem with both V- and W-cycles. a V-cycle

Eqs. 2.1–2.3b is on the order of the linear solver residual. We further note that the error measure E_t also decreases rapidly.

The *W*-cycle converges much faster for the first iterations, with an error reduction of almost three orders of magnitude from the first to the second iteration. According to the error estimates, the iterations could have been terminated at this point. The convergence rate then slows, but the *W*-cycle as expected consistently produces more accurate solutions than the *V*-cycle. The error estimates for the flux are consistently within a factor of 5 of the real error.

Next, we consider media with structured heterogeneities in the form of low-permeable inclusions. The coefficient \mathbf{a} is defined to have the value $10^{-\kappa}$ on the interior of the coarse dual cells and unity on the boundaries of the dual grid, as illustrated in Fig. 5a. The upscaled medium will thus be homogenous and isotropic. The domain is again two-dimensional with 3^6 cells in each dimension and 5 levels of coarsening. Similar to the homogeneous case, the

error reduction factors are quickly stabilized, and only the mean reduction factors for the CV solver with varying values of κ are reported; see Table 1. The convergence rates



Fig. 5 Conceptual drawings of media with structured (left) and random (right) heterogeneities. White marks fine-scale cells with high permeabilities. Black lines indicate coarse cells. For the structured case, the low permeable area coincides with the interior of coarse dual cells

Table 1 Asymptotic convergence rates obtained by V-cycles for a 2D medium with a structured heterogeneity with varying degree of contrast

	$\kappa = 5$	$\kappa = 4$	$\kappa = 3$	$\kappa = 2$	$\kappa = 1$	$\kappa = 0$
Red e(u)	0.066	0.066	0.066	0.066	0.066	0.065
Red e(q)	0.071	0.071	0.071	0.071	0.071	0.075

are observed to be stable with respect to the coefficient contrast both for the flux and the potential.

The above tests establish the convergence of the multi-level method as a solver. Furthermore, the convergence is robust with respect to structured low-permeable inclusions. The case of $\kappa > 0$ is not considered, as the structured coarse spaces utilized herein are inappropriate for that coefficient field, and the method consequently fails. This situation can be remedied by altering the coarse partitioning. The use of adaptive coarse partitioning, including algebraic aggregation, is currently under development for the methods presented herein, and will not be considered in this communication.

6.1.2 3D: Random heterogeneities

To further investigate the error estimates we consider a 3D grid with random heterogeneities as illustrated in Fig. 5b, and a coefficient contrast of 10. The 3D Cartesian grid is one case where the post-processing sets limits on the admissible coarsening ratios; if a coarsening ratio of 3 is applied in each direction, there will be interaction regions where there are fewer sub-interfaces than there are exterior unknowns (4 and 6, respectively, to be precise). Thus, in accordance with the discussion in Section 4.3, we consider a grid with 45 unknowns in each direction, and apply a coarsening of 5, followed by a coarsening of 3. The number of Gauss-Seidel smoothing steps is set to 3 in this simulation, and we only

report the result from a V-cycle; the W-cycle shows similar behavior.

The results obtained are presented in Fig. 6. We observe that the convergence is slower than for the 2D examples. Not unexpectedly, the AMG solver, which can better account for heterogeneities, shows a faster convergence. However, taking the error estimates into account, the CV solver could have been terminated after seven iterations, at which point the AMG residual still is of considerable magnitude. Again, the relative size of fluxes associated with unphysical values of \mathbf{a}^* decreases rapidly, as measured by E_t .

6.2 Two-level methods

Next, we consider the CV solver with a single level of coarsening rendering a domain decomposition-style method, which in principle is identical to the algebraic formulation of the multiscale finite volume method [28]. When the CV method is applied as a stand-alone solver for this case, the convergence is slow, and we therefore accelerate the iterations by GMRES. We consider two test cases: A Cartesian grid with unstructured, binary coefficient field, and an unstructured grid with anisotropic coefficient field.

6.2.1 Heterogeneous medium

We consider a grid with 21^2 cells in each direction in 2D and with a coarsening ratio of 21. The medium is isotropic, with

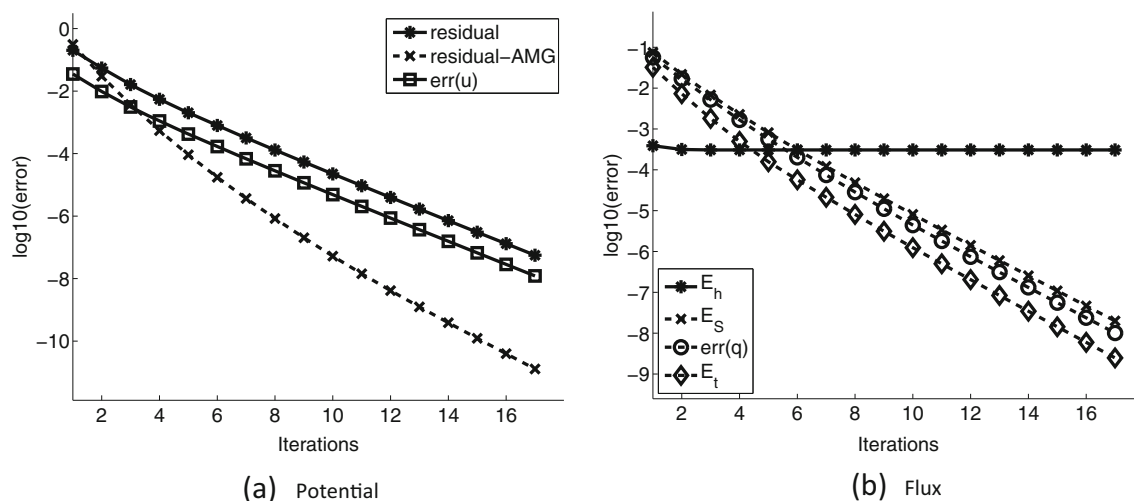


Fig. 6 Convergence of potential and flux errors, and estimators for flux and discretization errors on a heterogeneous 3D problem with a V-cycle

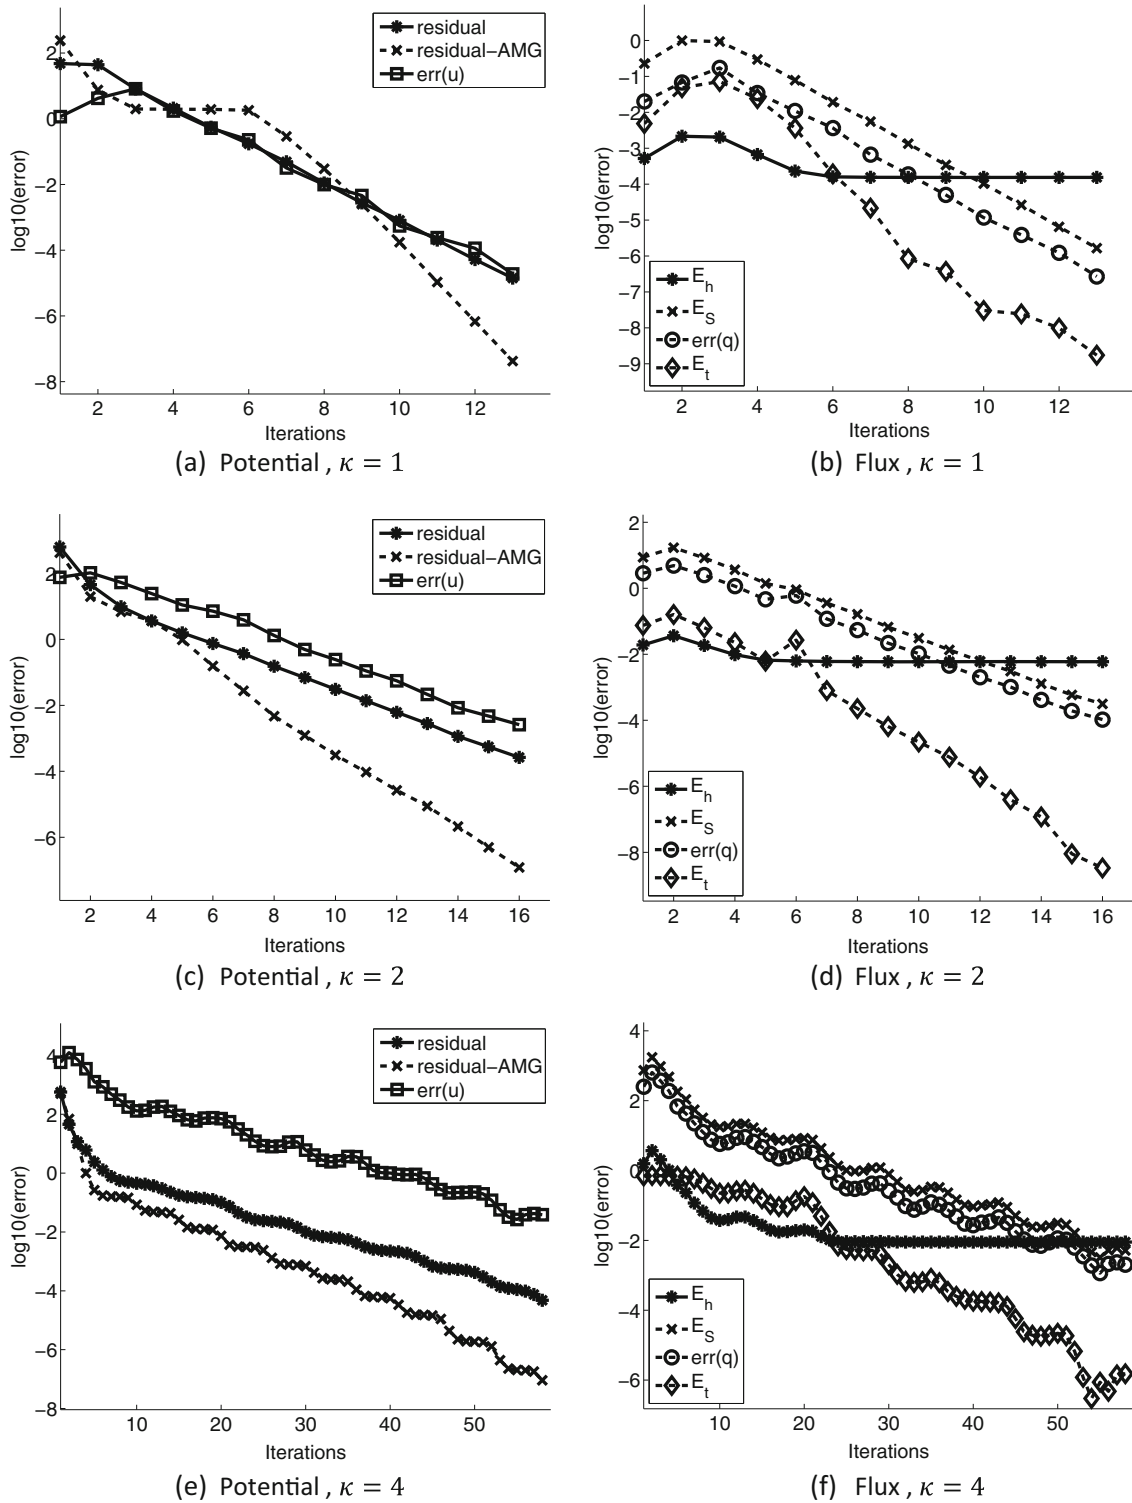


Fig. 7 Convergence of potential and flux errors, and estimators for flux and discretization errors on a heterogeneous 2D problem with a Cartesian grid




Fig. 8 Illustration of the unstructured grid used in section 6.2.2. Thick black lines show edges of the primal grid. The dual grid is formed by vertex cells (dark grey), edge cells (light grey) and interior cells (white)

a random binary coefficient field as illustrated in Fig. 5b, with values of 1 and $10^{-\kappa}$, $\kappa = \{1, 2, 4\}$ The righthand side is given by

$$f = \sin(3\pi x) \cos(3\pi y)$$

Figure 7 shows the convergence for the various errors and error estimates. We note that in this case, the multiple levels of coarsening make AMG a more efficient preconditioner with a quicker convergence of the residual. However, the error estimates indicate that an early termination of the CV iterations is admissible, making the CV solver competitive for $\kappa = \{1, 2\}$ When $\kappa = 4$, the situation is less clear, as both methods converge slowly in this case, and $E_S < E_h$ only when AMG is close to convergence.

6.2.2 Unstructured grids

The previous test cases have established the viability of the inexact solver on media where the coefficient \mathbf{a} was either homogeneous or had no spatial correlation. In this final example, we consider a setup that is closer to the challenges that are met in practical subsurface simulations. The coefficient \mathbf{a} is homogeneous, but with an anisotropy ratio of 10, where the eigenvectors are not aligned with the grid. The fine-scale grid consists of 201^2 cells that are the dual cells of a regular triangulation, forming predominantly hexagons, but also pentagons and squares along the boundary, see Fig. 8 for an illustration. Note that in this case, the fine-scale MPFA discretization does not reduce to a two-point flux expression. Figure 8 also illustrates the coarse primal and dual grids; the full coarse grid has 14^2 cells. We note that on the coarse scale cells that are not adjacent to the boundary (the middle cell in Fig. 8) will have 7-point stencils. A point source is placed in the lower left corner of the grid, and an equally sized sink is located in the upper right corner. No-flow boundary conditions are assigned.

The convergence history for the iterations is shown in Fig. 9. Again, we see that the CV and AMG preconditioners are comparable in terms of the residual reduction rate. The estimated discretization error is in this case much larger than the flux error, and thus the multiscale solution may in this case be deemed sufficiently accurate. With regards to the error estimates, we note that the results of [18], upon which our estimates (5.1)–(5.2) are based, are only valid for problems with Dirichlet boundary conditions, and thus do not provide guaranteed upper bounds for the present case. Nevertheless, we still consider the estimates as useful error indicators for controlling the iterative solver.

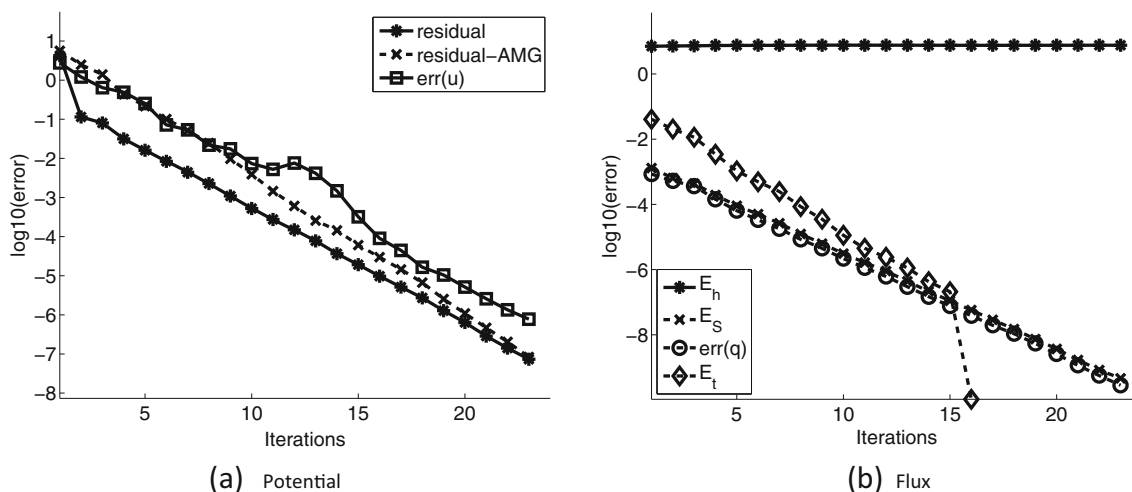


Fig. 9 Convergence of potential and flux errors, and estimators for flux and discretization errors on a heterogeneous 2D problem with an unstructured grid

Finally, we note that after 16 iterations with the CV solver, there are no fluxes with the wrong sign, as measured by E_t .

We end the result section with a comment on parameter heterogeneities. The results presented for the two-level method show that the present formulation of the CV preconditioner is at least as good as an AMG preconditioner for homogeneous problems, also when the medium is anisotropic. However, the performance of the CV preconditioner deteriorates faster than with AMG when the parameter field becomes heterogeneous. If the GS smoother in the CV preconditioner is replaced by a more efficient smoother such as block-ILU, the two-level CV approach has been shown to be comparable to AMG also for highly heterogeneous parameter fields [47], [43]. Ultimately, the presented examples emphasize that the method presented herein forms a compromise between a dedicated solver and a multiscale method. While the convergence performance will not in all cases be competitive with dedicated solvers such as AMG, the approach forms a viable compromise between upscaling methods such as multiscale methods and fine-scale solvers.

7 Concluding remarks

In this paper, we have formulated a fully multilevel control volume discretization for elliptic equations and considered its application as an inexact linear solver. The design philosophy behind the linear solver is that key properties that are preserved in the discretization of the continuous problem, should also be preserved by the linear solver. In the context of control volume methods for the scalar elliptic equation, this leads us to define a method that provides a conservative flux field even if the associated potential is inexact. The key ingredient to achieve this is a hierarchical post-processing of the solution to get a conservative flux field. This post-processing has been considered previously for two-level methods, but the extension to multilevel methods for general grids is novel. Since the linear solver is designed with inexact solutions in mind, nuanced error control is desirable. To that end, we have derived guaranteed a posteriori error bounds for the flux field, as well as an estimate that is tailored for problems with significant uncertainty in the permeability field.

Our numerical tests of both two-level and multilevel methods show that the methods perform well in benchmarking with an algebraic multigrid solver. Furthermore, the sharpness of the error estimates is confirmed, and since the inexact solver need not converge to produce a conservative flux field, substantial computational savings can be made by terminating the iterations when the solver no longer constitute the main source of errors in the solution.

It is common practice to view linear solvers as independent from the other parts of the simulation tool, often on a plug-in basis, with lesser regard to the origins of the linear system or to desired properties of the produced solution. This paper illustrates how the linear solver on the contrary can in an advantageous way be designed as an integral part of the simulator. The current work thus forms part of a larger effort to obtain a holistic approach to spatial discretizations, linear solvers, upscaling, error control, and handling of uncertain parameters.

References

1. Aavatsmark, I.: An introduction to the multipoint flux approximations for quadrilateral grids. *Comput. Geosci.* **6**(3–4), 405–432 (2002)
2. Aziz, K., Settari, A.: *Petroleum reservoir simulation*. s.l. Chapman & Hill, London (1979)
3. Bell, W., Olson, L., Schroder, J.: PyAMG: Algebraic multigrid solvers in PYTHON v2.0. <http://www.pyamg.org> :Release 2.0 (2011)
4. Cai, Z., Mandel, J., McCormick, S.: The finite volume element method for diffusion equations on general triangulations. *SIAM J. Numer. Anal.* **18**(3), 392–402 (1991)
5. Chen, Y., Durlafsky, L.J., Gerritsen, M., Wen, X.H.: Acoupled local–global upscaling approach for simulating flow in highly heterogeneous formations. *Adv. Water Resour.* **26**, 1041–1060 (2003)
6. Chen, Z., Huan, G., Ma, Y.: *Computational methods for multiphase flows in porous media*. s.l. SIAM, Philadelphia (2006)
7. Computer Modelling Group, L.L.C.: GEM: generalized equation-of-state model compositional reservoir simulator, s.l.: s.n. (2012)
8. Edwards, M.G., Rogers, C.: Finite volume discretization with imposed flux continuity for the general tensor pressure equation **2**, 259–290 (1998)
9. Efendiev, Y., Hou, T.: *Multiscale finite element methods: theory and applications* s.l, Springer (2009)
10. Eymard, R., Herard, J.-M.: *Finite Volumes for complex applications V*. s.l., Wiley (2008)
11. Farmer, C.L.: Upscaling: a review. *Int. J. Numer. Methods Fluids* **40**(1–2), 63–78 (2002)
12. Hackbusch, W.: *Multi-grid methods and applications 2*. red. s.l. Springer (2003)
13. Hauge, V., Lie, K.-A., Natvig, J.: Flow-based coarsening for multiscale simulation of transport in porous media. *Comput. Geosci.* **16**(2), 391–408 (2012)
14. Hou, T., Wu, X.-H.: A multiscale finite element method for elliptic problems in composite materials and porous media. *J. Comput. Phys.* **134**(1), 169–189 (1997)
15. Howell, J., et al.: Sedimentological Parameterization of shallow-marine reservoirs. *Pet. Geosci.* **14**, 17–34 (2008)
16. Hægland, H., et al.: Improved streamlines and time-of-flight for streamline simulation on irregular grids. *Adv. Water Resour.* (2007)
17. Jenny, P., Lee, S., Tchelepi, H.: Multi-scale finite-volume method for elliptic problems in subsurface flow simulations. *J. Comput. Phys.* **187**(1), 47–67 (2003)
18. Jiránek, P., Strakos, Z., Vohralik, M.: A posteriori error estimates including algebraic error and stopping criteria for iterative solvers. *SIAM J. Sci. Comput.* **32**(3), 1567–1590 (2010)

19. Keilegavlen, E., Nordbotten, J., Aavatsmark, I.: Sufficient criteria are necessary for monotone control volume methods. *Appl. Math. Lett.* **22**(8), 1178–1180 (2009)
20. Kim, H., Xu, J., Zikatanov, L.: Uniformly convergent multigrid methods for convection-diffusion problems without any constraint on coarse grids. *Adv. Comput. Math.* **20**, 385–399 (2004)
21. Kippe, V., Aarnes, J.E., Lie, K.-A. *Comput. Geosci.* **12**(3), 377–398 (2008)
22. Klausen, R.A., Stephansen, A.F.: Convergence of multi-point flux approximation on general grids and media. *Inter. J. Numer. Anal. Model.* **9**(3), 584–606 (2012)
23. Künze, R., Lunati, I., Lee, S.: A multilevel multiscale finite-volume method. *J. Computatinoal Phys.* **255**, 502–520 (2013)
24. Leveque, R.: Numerical methods for conservation laws. s.l. Springer, Berlin (1992)
25. Lipnikov, K., Moulton, J., Svysatskiy, D.: A multilevel multiscale mimetic (M3) method for two-phase flows in porous media. *J. Comput. Phys.* **227**(14), 6727–6753 (2008)
26. Lunati, I., Tyagi, M., Lee, S.: An iterative multiscale finite volume algorithm converging to the exact solution. *J. Comput. Phys.* **230**, 1859–1864 (2011)
27. Nordbotten, J., Aavatsmark, I., Eigestad, G.: Monotonicity of control volume. methods. *Numer. Math.* **106**(2), 255–288 (2007)
28. Nordbotten, J., Bjørstad, P.: On the relationship between the multiscale finite-volume method and domain decomposition preconditioners. *Comput. Geosci.* **12**(3), 367–376 (2008)
29. Nordbotten, J., Keilegavlen, E., Sandvin, A.: I: R. Petrova, red. *Mass Conservative Domain Decomposition for Porous Media Flow*. s.l.:In Tech. (2012)
30. Nordbotten, J.M.: Adaptive variational multiscale methods for multiphase flow in porous media. *Multiscale Model. Simul.* **7**(3), 1455–1473 (2009)
31. Pasciak, J.E., Vassilevski, P.S.: Exact de Rham sequences of spaces defined on macro-elements in two and three spatial dimensions. *SIAM J. Sci. Comput.* **30**(5), 2427–2446 (2008)
32. Repin, S.: A posteriori error estimates for partial differential equations. s.l., Radon Series on Computational and Applied Mathematics (2008)
33. Rognes, Ø., Tai, X.-C.: Subspace correction methods for non-symmetric parabolic problems. *Linear Algebra Appl.* **332–334**, 205–235 (2001)
34. Russell, T.F., Wheeler, M.F.: Finite element and finite difference methods for continuous flows in porous media. I: R. E. Ewing, red. *Mathematics of Reservoir Simulation*. s.l. Soc. Ind. Appl. Math, 35–106 (1983)
35. Sandvin, A., Nordbotten, J., Aavatsmark, I.: A unified multi-level framework of upscaling and domain decomposition. s.l., Proc. of XVIII International Conference on Water Resources (2010)
36. Sandvin, A., Nordbotten, J.M., Aavatsmark, I.: Multiscale mass conservative domain decomposition for porous media flow. *Comput. Geosci.* **15**(3), 587–602 (2011)
37. Schlumberger: ECLIPSE technical description, s.l.: s.n. (2010)
38. Thomas, J., Trujillo, D.: Finite volume variational formulation. Application to domain decomposition methods. *Contemp. Math.* **157**, 127–132 (1994)
39. Toselli, A., Widlund, O.: *Domain Decomposition methods - algorithms and theory*. s.l. Springer, Berlin (2005)
40. Trottenberg, U., Schuller, A.: *Multigrid*. 1. red Orlando. Academic Press, Florida (2000)
41. Vaněk, P., Mandel, J., Brezina, M.: Algebraic multigrid by smoothed aggregation for second and fourth order elliptic problems. *Computing* **56**, 179–196 (1996)
42. Voller, V.R.: *Basic Control volume finite element methods for fluids and solids*. s.l., IISc Research Monographs Series (2009)
43. Wang, Y., Hajibeygi, H., Tchelepi, H.: Algebraic multiscale solver for flow in heterogeneous porous media. In: *Journal of Computational Physics*, Issue doi:[10.1016/j.jcp.2013.11.024](https://doi.org/10.1016/j.jcp.2013.11.024) (2013)
44. Wesseling, P.: *An introduction to multigrid methods*. s.l. Wiley, New York (1992)
45. Winslow, A.M.: Numerical solution of the quasilinear Poisson equation in a non-uniform triangular mesh. *J. Comput. Phys.* **1**, 149–172 (1966)
46. Xu, J.C.: Iterative methods by space decomposition and subspace correction. *SIAM Review* **34**, 581–613 (1992)
47. Zhou, H., Tchelepi, H.: Two-stage algebraic multiscale linear solver for highly heterogeneous reservoir models, June. *SPE J.* **17**(2), 523–539 (2012)