



SPE 163649-MS

The Multiscale Finite-Volume Method on Stratigraphic Grids

Olav Møyner, Knut-Andreas Lie, SPE, SINTEF ICT.

Copyright 2013, Society of Petroleum Engineers

This paper was prepared for presentation at the SPE Reservoir Simulation Symposium held in The Woodlands, Texas USA, 1820 February 2013.

This paper was selected for presentation by an SPE program committee following review of information contained in an abstract submitted by the author(s). Contents of the paper have not been reviewed by the Society of Petroleum Engineers and are subject to correction by the author(s). The material does not necessarily reflect any position of the Society of Petroleum Engineers, its officers, or members. Electronic reproduction, distribution, or storage of any part of this paper without the written consent of the Society of Petroleum Engineers is prohibited. Permission to reproduce in print is restricted to an abstract of not more than 300 words; illustrations may not be copied. The abstract must contain conspicuous acknowledgment of SPE copyright.

Abstract

Finding a pressure solution for large and highly detailed reservoir models with fine-scale heterogeneities modeled on a meter scale is computationally demanding. One way of making such simulations less compute intensive is to employ multiscale methods that solve coarsened flow problems using a set of reusable basis functions to capture flow effects induced by local geological variations. One such method, the multiscale finite-volume (MsFV) method, is well studied for 2D Cartesian grids but has not been implemented for stratigraphic and unstructured grids with faults in 3D. We present an open-source implementation of the MsFV method in 3D along with a coarse partitioning algorithm that can handle stratigraphic grids with faults and wells. The resulting solver is an alternative to traditional upscaling methods, but can also be used for accelerating fine-scale simulations. To achieve better precision, the implementation can use the MsFV method as a preconditioner for Arnoldi iterations using GMRES, or as a preconditioner in combination with a standard inexpensive smoother.

We conduct a series of numerical experiments in which approximate solutions computed by the new MsFV solver are compared with fine-scale solutions computed by a standard two-point scheme for grids with realistic permeabilities and geometries. On one hand, the results show that the MsFV method can produce accurate approximations for geological models with pinch-outs, faults, and non-neighboring connections, but on the other hand, they also show that the method can fail quite spectacularly for highly anisotropic systems in a way that cannot efficiently be mitigated by iterative approaches.

The MsFV method is, in our opinion, not yet sufficiently robust to be applied as a black-box solver for models with industry-standard complexity. However, extending the method to realistic grids is an important step on the way towards fast and accurate multiscale solution of large-scale reservoir models. In particular, our open-source implementation provides an efficient framework suitable for further experimentation with partitioning algorithms and MsFV variants.

Introduction

Multiscale methods have been proposed as a way of bridging the gap in resolution between geological models (cell sizes: centimeters to decimeters in the vertical direction, meters to tens of meters in horizontal direction) and dynamic simulation models (cell sizes: meters to tens of meters in vertical direction, tens of meters to hundred of meters in horizontal direction). As an alternative to traditional upscaling techniques, multiscale methods (Efendiev and Hou 2009) can resolve fine-scale quantities with reduced computational complexity (Kippe et al. 2008) on highly detailed reservoir models by creating basis functions that relate localized flow problems on the fine scale (geological model) to a global flow problem on a much coarser scale (dynamic model). Herein, we consider the multiscale finite-volume (MsFV) method (Jenny et al. 2003) in which the basis functions in the MsFV method are computed using a dual-grid formulation with unitary pressure values at each vertex of the coarse blocks. The method has been extended to provide qualitatively correct approximations to wide variety of problems in subsurface flow, including density-driven flow (Lunati and Jenny 2008; Künze and Lunati 2012), compressible flow (Lunati and Jenny 2006; Hajibeygi and Jenny 2009; Zhou and Tchelepi 2008), three-phase and compositional flow (Lee et al. 2008; Hajibeygi and Jenny 2013), well modeling (Wolfsteiner et al. 2006; Jenny and Lunati 2009), and fractured systems (Hajibeygi et al. 2011b; Sandve et al. 2013).

To overcome accuracy problems observed for strongly heterogeneous cases, and generally be able to systematically reduce fine-scale residuals, an iterative method was introduced by Hajibeygi et al. (2008) and further developed in (Hajibeygi and Jenny 2011). The key idea is to use the MsFV coarse-scale operator to reduce low-frequency errors and a combination of locally computed correction functions and inexpensive smoothers to reduce high-frequency errors. Likewise, the method has been formulated in an algebraic manner (Zhou and Tchelepi 2008; Lunati et al. 2011) and applied as a preconditioner (Lunati et al. 2011; Zhou and Tchelepi 2012; Wang et al. 2012). The close relationship between MsFV and domain-decomposition methods is discussed in more detail in (Nordbotten and Bjørstad 2008; Sandvin et al. 2011; Nordbotten et al. 2012; Sandvin et al. 2013). A particular advantage of the MsFV method, compared with multigrid and domain-decomposition methods, is that the multiscale

method can reconstruct a conservative flux field at any iteration stage, which is crucial if the flux field is used to solve transport problems.

Application to industry-standard stratigraphic and unstructured grids is a remaining challenge for the MsFV method. The modeling approaches used in industry today are predominantly structured. However, highly irregular and complex grids having unstructured connections and degenerate cell geometries arise naturally when representing structural frameworks like faults, joints, and deformation bands, and/or stratigraphic architectural characteristics like channels, lobes, clinofolds, and shale shale/mud drapes. Similarly, unstructured connections are induced when local grid refinement, structured or unstructured, is used to improve the modeling of (deviated) wells. With the exception of a highly idealized model of faults represented on matching 2D structured grids (Hajibeygi et al. 2011a), the MsFV method has so far mainly been formulated and applied to grids with a regular topology. Extending the MsFV formulation from uniform Cartesian to industry-standard grids with complex geometries/topologies and high aspect/anisotropy ratios is therefore a key step on the road to widespread adoption in industry. In particular, it is desirable to develop automated coarsening strategies that perform well for complex models and preferably adapts to geological features and well paths to ensure optimal accuracy for a chosen level of coarsening. This, in turn, requires robustness and a high degree of flexibility with respect to the size and shape of the coarse blocks. Extending the original geometric formulation (Jenny et al. 2003) of the MsFV method to grids with unstructured connections has proved difficult. However, the more recent operator formulation (Lunati et al. 2011) casts the method in an algebraic form and offers much of the flexibility needed to develop a MsFV method for grids with unstructured connections and irregular geometries. Using the operator formulation, the key challenges are to develop: (i) automated and robust methods for constructing primal and dual coarse grids, and (ii) appropriate formulations of a conservative coarse-scale operator and localization conditions for the basis functions that are robust to large aspect ratios and strong variations in media properties at the faces of the coarse blocks. Herein, we start by discussing how to formulate the MsFV method for general unstructured grids and then present a set of coarsening methods that can create the required primal/dual grids for a reasonable class of stratigraphic grids. To this end, we first create a primal coarse grid using techniques developed for the multiscale mixed finite-element method, for which each coarse block can consist of an (almost) arbitrary simply-connected set of polyhedral cells (see e.g., (Aarnes et al. 2006, 2008; Natvig et al. 2011; Alpak et al. 2012; Lie et al. 2012)). Then, the accompanying dual grid is created based on geometrical considerations. We show that the resulting method can produce solutions that are in good agreement with fine-scale reference solutions for realistic stratigraphic grids with pinch-outs, faults, and non-neighboring connections. However, as discussed recently by several authors (Nordbotten et al. 2012; Zhou and Tchelepi 2012; Sandvin et al. 2013; Wang et al. 2012), the method may introduce (surprisingly) large errors for highly heterogeneous and anisotropic problems on Cartesian grids that cannot be eliminated in an efficient manner by the standard iterative approaches. We demonstrate that this behavior may be further aggravated for realistic 3D stratigraphic grids containing cells with large differences in face areas, e.g., because of pinch-outs or across slip faults. We also present a simple method for modifying the reduced-problem boundary condition used for localization of basis functions, which will mitigate the problem in some cases.

The MsMV solvers discussed herein have been implemented as a separate module in the open-source Matlab Reservoir Simulation Toolbox (MRST 2013) and can be freely downloaded and used within the GPL 3.0 software license. (The MsFV module was recently used by Sandve et al. (2013) to study fractured geothermal reservoirs modeled on 2D triangular grids; the resulting solvers are also freely available as a module in MRST.) To support reproducible research, scripts and data necessary to run numerical experiments reported in the paper will be published on the MRST website.

The Multiscale Finite-Volume Method

To keep the presentation of new ideas as simple as possible, we will in the following only consider incompressible flow and neglect the effects of capillarity and gravitational forces; these effects can easily be included as discussed in the literature. Combining mass conservation, $\nabla \cdot \vec{v} = q$, with Darcy's law, $\vec{v} = -\mathbf{K}\lambda_t \nabla p$, we obtain a Poisson-type, elliptic pressure equation

$$-\nabla \cdot (\mathbf{K}\lambda_t \cdot \nabla p) = q \quad (1)$$

for the fluid pressure p and the total velocity v . Here, q denotes fluid sources, \mathbf{K} is the permeability tensor, and $\lambda_t = \sum_{\alpha} \lambda_{\alpha}$ is the total mobility, where the mobility λ_{α} of phase α equals the ratio of the relative permeability k_{α} and viscosity μ_{α} of phase α .

In the geometric formulation, the MsFV method can be informally described as first computing a reduced flow problem along each face of the dual coarse blocks to obtain localized flow problems inside each dual coarse block. Once these problems have been solved and assembled into a basis, the interactions between the basis functions over the primal coarse grid gives a coarse system with *multipoint connections*. The solution of the coarse pressure system can be extended to the fine grid using the subresolution of the pressure basis functions. To properly account for the fine-scale effects of wells, compressibility, and gravity, it is common to construct so-called correction functions (Lunati and Jenny 2008) that take care of the particulate part of the solution cannot be represented by the homogeneous basis functions. Conservative fine-scale fluxes can be constructed by solving another set of local problems with Neumann boundary conditions inside each primal coarse block.

The algebraic form of the MsFV method uses vector partitions, permutations, and manipulation and elimination of matrix blocks in the underlying fine-scale discretization to derive basis functions, correction functions, and the coarse-scale operator. This formulation will be our starting point, and to make the paper more self-contained, we will in the following briefly review its details. Later, however, we will use concepts from the geometric description of the method whenever we feel that this makes

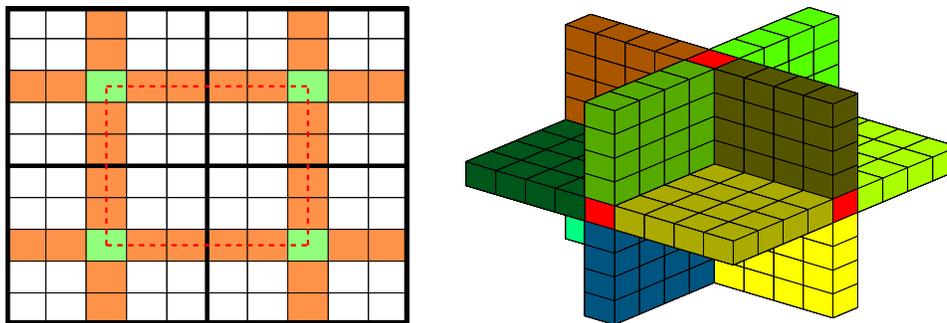


Fig. 1— The left plot shows a primal coarse grid (thick black lines) along with a single block of the corresponding dual grid (dashed red lines). Node cells are shown in green, edge cells in orange, and interior cells in white; there are no face cells in 2D. The right plot shows 'edge' cells in red and 'face' cells in different colors for a 3D Cartesian grid.

the discussion more transparent to the reader. Unless stated explicitly otherwise, our setup of the method follows previous approaches, including parts of the method that are not reviewed in this section (e.g., correction functions, iterations, smoothing, etc). This means that for Cartesian grids fine/coarse grids, our multiscale method will produce the same results as the method described e.g., in (Lunati et al. 2011). Finally, we encourage the reader to consult (Lunati et al. 2011; Zhou and Tchelepi 2012; Møyner 2012) for more details about the restriction and prolongation operators used in the operator formulation, or the published literature the geometric formulation of the MsFV method for more details about the reduced-problem boundary condition used for localization of basis functions, the correction functions, reconstruction of conservative fine-scale fluxes, smoothing steps, and iterations.

Fine-scale discretization. A multiscale method will neither have better accuracy nor contain less numerical artifacts than the method used to discretize the flow equations on the fine scale. In industry, the predominant discretization approach is to use a two-point flux approximation (TPFA) method (i.e., the finite-volume counterpart of a seven-point finite-difference stencil on Cartesian grids in 3D), even though it is well known that this method is only consistent for K-orthogonal grids and will therefore not converge in the general case. On the other hand, this method is guaranteed to be monotone, unlike consistent (linear) discretization methods like the multipoint flux approximation (MPFA) schemes, mimetic finite differences, and mixed finite elements.

To develop the multiscale formulation, we assume that the computational domain is represented on the fine scale by a grid Ω_h consisting of cells c_i . Then, the corresponding fine-scale discretization of Eq. 1 results in a linear system on the form

$$A_h \mathbf{p}_h = \mathbf{q}_h, \quad (2)$$

where the matrix A_h corresponds to a set of mass balance equations for the cell-center pressures \mathbf{p}_h with the source terms being collected in \mathbf{q}_h . Although the general form Eq. 2 includes consistent discretizations like MPFA schemes, we will in the following only use the standard TPFA method in all numerical examples.

Categorization of cells and coarse grids. In addition to the fine grid Ω_h , the MsFV method requires two different coarse grids to be defined. The *primal coarse grid*, denoted as Ω_H , consists of coarse blocks B_j that are composed of fine cells $c_i \in \Omega_h$ such that $\Omega = \cup_{j \in [1, N]} B_j$, with every fine cell belonging to one and only one block B_j . The coarse grid does not need to be defined explicitly to implement the MsFV method, but a categorization of each c_i to some B_j is required to get conservative flux.

The *dual coarse grid* $\tilde{\Omega}_H$ has corners corresponding to the block centers of Ω_H . Because Ω_H is defined by fine cells, its edges run alongside the fine-cell edges. The dual grid $\tilde{\Omega}_H$, however, has edges running through fine cells so that we can define a set of fine cells as edge cells. This concept is central to the operator formulation (Lunati et al. 2011; Lunati and Lee 2009), wherein all fine cells c_i are categorized according to their positions in $\tilde{\Omega}_H$ using a so-called wire-basket ordering (Smith et al. 1996). The categories can be explained as follows (see Figure 1 for two simple illustrations):

- **Nodes** (n): Cells containing the centers of Ω_H , i.e., points where three coarse faces of \tilde{B}_j intersect for 3D Cartesian grids.
- **Edges** (e): Cells on the edge of the faces of $\tilde{\Omega}_H$, i.e., cells common to two faces of one block \tilde{B}_j for 3D Cartesian grids.
- **Faces** (f): Cells on a faces of the dual coarse grid $\tilde{\Omega}_H$, but not on the edges of the face. For 3D Cartesian grids, this would be equivalent with being on a single face for some \tilde{B}_j . For 2D domains this category is omitted.
- **Inner** (i): Cells not on the faces of $\tilde{\Omega}_H$.

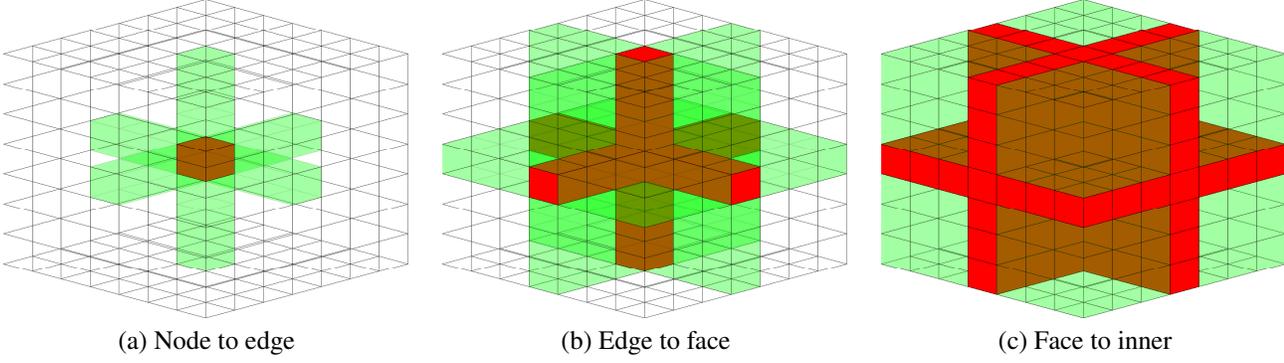


Fig. 2— Illustration of the three extrapolation stages of the MsFV formulation. In each plot, the previous step is shown in red, with the current in green.

Permuting the system matrix. Once all cells have been categorized, we construct a permutation matrix P that reorders Eq. 2 so that the unknowns are grouped by categorization of their corresponding cells:

$$P\mathbf{p}_h = \mathbf{p} = \begin{bmatrix} \mathbf{p}_i \\ \mathbf{p}_f \\ \mathbf{p}_e \\ \mathbf{p}_n \end{bmatrix}, \quad PA_hP^T = A = \begin{bmatrix} A_{ii} & A_{if} & 0 & 0 \\ A_{fi} & A_{ff} & A_{fe} & 0 \\ 0 & A_{ef} & A_{ee} & A_{en} \\ 0 & 0 & A_{ne} & A_{nn} \end{bmatrix}. \quad (3)$$

Here, matrix block A_{kl} represents the influence from cells of category l to the mass balance of cells of category k . To develop the approximate multiscale method, the equations for the interior cells are kept, while the equations for the edge and face cells are modified to reflect the reduced-boundary conditions used for localization. These conditions are applied by setting A_{fi} and A_{ef} to zero and removing the corresponding part from the diagonal blocks (see, e.g., (Lunati et al. 2011; Møyner 2012; Wang et al. 2012) for more details),

$$(M_{kk})_{rr} = (A_{kk})_{rr} + \sum_s (A_{ks})_{rs}. \quad (4)$$

Once the fine-scale system is reduced to upper block-triangular form by removing connections so that edges only depend on the values of nodes, the faces only on the edges, and so on, we obtain the multiscale system matrix,

$$\begin{bmatrix} A_{ii} & A_{if} & 0 & 0 \\ 0 & M_{ff} & A_{fe} & 0 \\ 0 & 0 & M_{ee} & A_{en} \\ 0 & 0 & 0 & M_{nn} \end{bmatrix} \begin{bmatrix} \mathbf{p}_i \\ \mathbf{p}_f \\ \mathbf{p}_e \\ \mathbf{p}_n \end{bmatrix} = \begin{bmatrix} \mathbf{q}_i \\ \mathbf{q}_f \\ \mathbf{q}_e \\ \mathbf{q}_n \end{bmatrix}, \quad (5)$$

which can be solved using backward substitution. (The localization assumption is the only source of error in the multiscale approximation.) If we for a while assume that the values \mathbf{p}_n in cells of category 'node' are known, it follows immediately from Eq. 5 that the unknown pressure values in the 'edge' cells are given by $M_{ee}\mathbf{p}_e + A_{en}\mathbf{p}_n = \mathbf{q}_e$, and hence can be computed by first extrapolating values from 'node' cells $A_{en}\mathbf{p}_n$, and then solving for the rest of the category, $M_{ee}^{-1}(\mathbf{q}_e - A_{en}\mathbf{p}_n)$. Notice that this corresponds to the solution of reduced problems with unitary pressures specified at node points used to provide localization in the geometric description of the method. The face and inner unknowns can be computed similarly by first extrapolating values from the previous category and then inverting the diagonal matrix block. This procedure is illustrated in Figure 2. Altogether, the solution has the form

$$\mathbf{p} = B\mathbf{p}_n + C\mathbf{q}, \quad (6)$$

where the matrices B and C are defined as follows:

$$B = \begin{bmatrix} A_{ii}^{-1}A_{if}M_{ff}^{-1}A_{fe}M_{ee}^{-1}A_{en} \\ M_{ff}^{-1}A_{fe}M_{ee}^{-1}A_{en} \\ M_{ee}^{-1}A_{en} \\ I \end{bmatrix}, \quad C = \begin{bmatrix} A_{ii}^{-1} & -A_{ii}^{-1}A_{if}M_{ff}^{-1} & A_{ii}^{-1}A_{if}M_{ff}^{-1}A_{fe}M_{ee}^{-1} & 0 \\ 0 & M_{ff}^{-1} & -M_{ff}^{-1}A_{fe}M_{ee}^{-1} & 0 \\ 0 & 0 & M_{ee}^{-1} & 0 \\ 0 & 0 & 0 & 0 \end{bmatrix}. \quad (7)$$

Here, we see that the matrix B consists of the basis functions that extend the coarse pressure solution \mathbf{p}_n to the fine grid, whereas C consists of the correction functions that handle fine-scale behavior of fluid sources (wells).

Coarse-scale system. When backwards substituting the system, we need to find values for the coarse node pressure \mathbf{p}_n . Because the cells representing nodes in the coarse grid are not direct neighbors in the fine grid, we cannot simply find M_{nn} in the same manner as for example M_{ee} . Instead, we find a volume-conservative coarse-scale system $M_{nn}\mathbf{p}_n = \mathbf{q}_n$ defined over the primal coarse grid, by summing up the mass balance equations for all the cells that belong to each block. To this end, we introduce a summation matrix χ that has one row χ_j for each primal coarse block B_j , defined such that element $\chi_{ij} = 1$ if cell c_i belongs to B_j and $\chi_{ij} = 0$ otherwise. Applying this summation operator to the reordered system and inserting Eq. 6, we obtain

$$\chi\mathbf{q} = \chi A(B\mathbf{p}_n + C\mathbf{q}) \iff \chi AB\mathbf{p}_n = \chi(I - AC)\mathbf{q}. \quad (8)$$

This allows us to construct a coarse pressure system which can be combined with the basis functions from Eq. 6 to construct a fine-scale pressure solution. For the nature of the operators used in these systems, we again refer to (Lunati et al. 2011; Lunati and Lee 2009). While additional steps are required to create iterative variants and construct conservative flux, these steps are the same for all variants of the method and have been covered thoroughly in earlier papers on the MsFV method.

Generating Coarse Grids

As explained above, the MsFV method requires two different coarse grids: a primal grid on which one formulates the global coarse-scale equations and a dual grid on which one computes the basis functions. Previous publications on the MsFV method have only considered regular Cartesian grids, for which a regular primary coarse grid and its shifted dual can be trivially obtained by introducing the tensor product of a load-balanced linear distribution along each axial direction. The same approach can be applied to partition the indices of stratigraphic grids that have a uniform Cartesian topology. However, most real-life models will have unstructured connections, degenerate cell geometries, and inactive cells, which all make the process of defining a reasonable pair of dual-primal coarse grids difficult. In the rest of this section, we will present methods that can be used to generate basic dual/primal coarse grids for a reasonably large class of stratigraphic grids.

Primal grid. The MsFV method was implemented using the Matlab Reservoir Simulation Toolbox (MRST 2013), in which all grids are assumed to consist of a set of non-overlapping polyhedral cells with matching planar faces. Grids with non-matching faces, arising for instance at a fault in the stratigraphic format, are converted into matching grids by splitting non-matching faces into a set of matching (sub)faces. All grids are represented in a general format in which cells, faces, and vertices as well as the connections between cells and faces are stored explicitly (Lie et al. 2012).

Starting from a fine-scale grid (structured, stratigraphic, or fully unstructured), the toolbox offers many different methods for generating primal coarse grids, ranging from simple methods that impose a load-balanced partition in index space to advanced methods for generating non-uniform coarse grids that can be constrained to geological features like high/low permeability structures, facies, stone types, or saturation regions, or adapted to flow indicators of various types. These coarsening methods have been developed for upscaling of flow and transport and for use with the multiscale mixed finite-element method (Aarnes et al. 2006; Kippe et al. 2008; Natvig et al. 2011; Hauge 2010; Alpak et al. 2012; Natvig et al. 2012). Extensive numerical experiments show that whereas a regular partition may be sufficient to get good accuracy in many cases, there are also cases for which improved robustness and increased accuracy is observed if the coarse grid is adapted to wells, heterogeneity, or flow features. However, since most of the advanced methods require some user experience, we will only utilize simple partition methods herein.

Dual grid. The dual grid will be represented implicitly as a partition that categorizes each fine cell into one of the four categories 'inner', 'face', 'edge', or 'node'. The key challenge is to ensure that node cells are sufficiently connected through edge cells, that edge cells are sufficiently connected through face cells, and that inner cells associated with one dual coarse block are not directly connected to inner cells of another dual block. If these conditions are not fulfilled, it will generally not be possible to compute basis functions and extrapolate pressure values between the different cell categories.

For Cartesian grids, a dual grid is trivially obtained by shifting the primal grid in each of the axial directions, a strategy that can easily be generalized to stratigraphic grids without non-neighborly connections. Using a logical partitioning scheme in index space has the advantage of being easy to implement and requires no lookup of actual node positions. Figure 1 shows such a dual partitioning for a simple 2D Cartesian grid. In geometrical terms, the algorithm can be explained as follows: for each pair of primal coarse blocks, draw a line between the centroids of the corresponding node cells, i.e., the dashed red line connecting two green cells in the figure. If a cell is intersected by one of these lines, it will be categorized as an edge cell (red cells in the figure).

The same approach can be applied to more general geometries/topologies. First, a node cell is defined for each primal block by choosing the cell whose centroid lies closest to the mean centroid of all cells inside the block. Then, for each pair of neighboring blocks we draw a line between the centroids of the node cells and define the corresponding edge as the union of all cells whose centroid lie within a prescribed distance to the line. This may give excessive edge cells, which in turn makes the multiscale approximation less prone to errors, as the reduced system is less sensitive to single cells with high variations in permeability. However, the linear systems corresponding to the edges will be larger, and for this reason we will use the system topology to strip away superfluous edge cells. First, we observe that inner cells from different dual coarse blocks should appear in different cycles of the directed graph derived from A_{ii} . (Cells belonging to different cycles are not connected and can hence be computed independent of each other when inverting the matrix). To remove excessive cells, we loop through the edge cells in a certain order

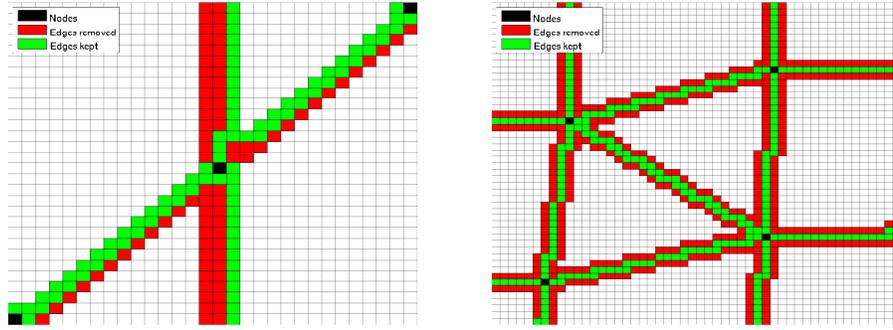


Fig. 3— The thickness of edges in the dual grid can be adjusted by a culling algorithm, here shown for a small and somewhat larger subset of a 2D of grid. The initial edges are intentionally made extra thick for illustration purposes.

and for each cell visited change its category to 'inner' if this operation does not reduce the number of cycles in A_{ii} . By using the system graph based on the discretizations to define connections, this can be done quickly while ensuring that the final linear equation sets remain solvable. For an example of this, see Figure 3.

Defining a dual grid for 3D grids with complex geometries is nontrivial; Møyner (2012) discusses several approaches that work in special cases. Herein, we present a more general approach that starts by defining sets of *common neighbors* and letting the partition algorithm work on these sets to produce the final grid. A set of *common neighbors* will be defined as three coarse blocks chosen so that every coarse-block is neighbor to at least one of the others and neighbors the other either directly or via some block not in the set. This allows us to cover both Cartesian-like neighborhood structures as well as unstructured grids, as can be seen in Figure 4. A triangle is then constructed between the centroids of the nodes of these neighbors as shown in the rightmost part of Figure 4. The left plot in Figure 5 shows the three planes defined for the three triplets of common neighbors that can be defined from the upper corner cell.

Once triangles have been generated for all triplets of common neighbors, any cells intersected by each plane will be added to the 'face' category. Any cells that are on the edge of two triangles of a triple, or on three or more triangles of all triplets, are on the path between two direct neighbors and will be categorized as on the edges, see Figure 5. Note that if the coarse block has more neighbors than its node cell has faces, edges will overlap because some edge cells are associated with more than one pair of nodes. This does not reduce the accuracy of the approximation, but reduces the number of cycles in M_{ee} , i.e., reduces the number of cell groups that can be computed in parallel when inverting M_{ee} .

While not always stated when discussing the MsFV method, having dual edges that touch the boundary is important to be able to impose boundary conditions other than no-flow. The observant reader may have noticed that our algorithm only categorizes cells within sub-domain bounded by the node cells, see Figure 6. There are two alternative methods to extend the dual partition to the whole domain: Either add virtual coarse blocks that do not contain any fine cells outside of the domain, or simply use the coarse-face centroids on the boundary to construct additional triangles. The rightmost plot in Figure 6 shows how this enables the rest of the dual grid to be constructed.

Numerical Experiments

In this section, we apply the MsFV method to a set of models with increasing degree of topological and geometrical complexity. The first example is a simple 3D box model, the second is a corner-point model with a single fault and highly heterogeneous petrophysical properties, whereas the third is a synthetic model with a curved fault represented using 2.5D prismatic or PEBI grids. The last two examples uses the corner-point geometries and permeabilities from real-life models: a model giving a detailed description of the bedding structure on the core scale, and a simulation model of a Norwegian Sea reservoir. We keep the flow physics as simple as possible in all examples since the main goal of this paper is to demonstrate that the capability of the MsFV method to handle realistic grids. When the test problems include wells, these are handled by the inclusion of correction functions (Lunati and Lee 2009).

In all examples, we will compare the approximate multiscale solutions to the fine-scale reference solutions using scaled L_∞ and L_2 norms,

$$\|\mathbf{u}\|_\infty = \frac{\|\mathbf{u}_{fs} - \mathbf{u}_{ms}\|_\infty}{\|\mathbf{u}_{fs}\|_\infty}, \quad \|\mathbf{u}\|_2 = \frac{\|\mathbf{u}_{fs} - \mathbf{u}_{ms}\|_2}{\|\mathbf{u}_{fs}\|_2} \quad (9)$$

where \mathbf{u} is either the vector of cell pressures or fine-scale fluxes.

SPE10 box. To demonstrate the accuracy that can be expected on 3D models with realistic aspect ratios and heterogeneity and hence set the stage for the discussion of the MsFV method applied to industry-standard stratigraphic grids, we start by considering a $60 \times 220 \times 35$ box model with a fine-scale cell size of $20 \times 10 \times 2$ feet. The box model is populated with three different sets of

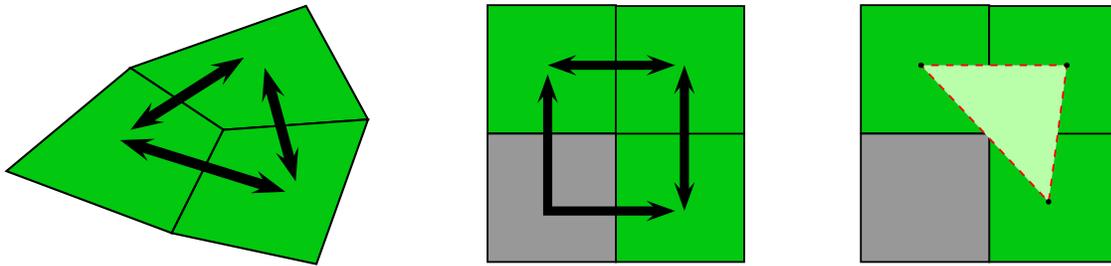


Fig. 4— Two neighbor sets: the left plot shows an unstructured grid in which all blocks are neighbors of each other, while the middle plot shows a Cartesian subset in which two of the neighbors are connected via a fourth coarse block shown in gray. The right plot shows how a triangle (red) is constructed from the centers of the neighbors.

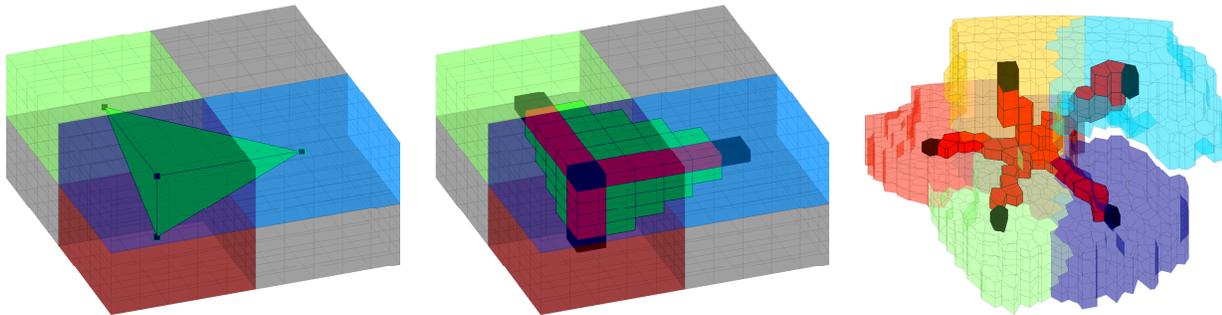


Fig. 5— Creating a dual partition: from a block and three neighbors, three planes are generated (left). Once the planes have categorized the cells, the resulting nodes (black), edges (red) and faces (green) are determined (middle). The left plot shows edge cells emanating from a single coarse cell in a subset of a 2.5D perpendicular bisector (PEBI) grid.

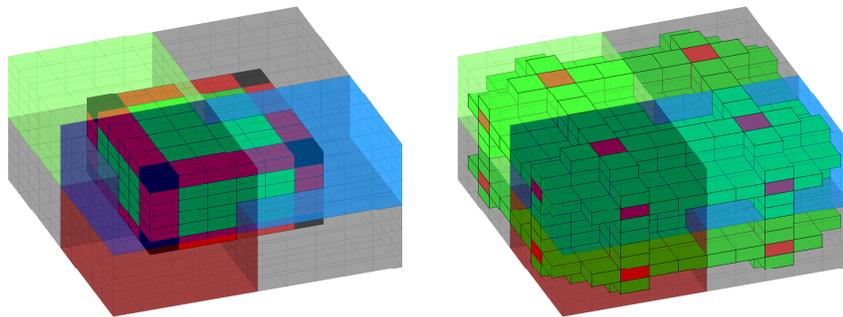


Fig. 6— Using triangular planes between common neighbors only defines the dual grid in a subset of the domain, bounded by the node cells (left). Adding virtual coarse blocks to the outside of the domain lets the dual grid reach the boundary faces (right).

permeabilities: isotropic, homogeneous permeability equal 100 mD and two different subsamples of Model 2 from the 10th SPE Comparative Solution Project (Christie and Blunt 2001), which is a geo-cellular model with highly heterogeneous petrophysical properties that was originally created as a challenging benchmark case for upscaling methods. Later, subsets of this model have become popular tests for multiscale and other numerical methods. In particular, various formulations of the MsFV method have been tested extensively using the SPE10 data set; 2D layers of the full model set have been used in some form in almost all previous papers discussing the method. Herein, we use permeability sampled from the upper 35 layers of the shallow-marine Tarbert formation, and permeability sampled from the bottom 35 layers of the fluvial Upper Ness formation. Flow is driven from east to west by imposing Dirichlet boundary conditions of 500 bar and 250 bar along the east and west faces of the model. The fine grid is partitioned into coarse blocks of $5 \times 5 \times 5$ fine cells each, giving a total of 3696 degrees of freedom for the coarse system.

Table 1 reports the discrepancies between the fine-scale and MsFV pressure and flux solutions measured in the relative L_2 and L_∞ norms, Eq. 9. With homogeneous permeability, the analytical solution describes a linear pressure drop from east to west. In this case, the localization assumption is exact and the multiscale method reproduces the fine-scale solution to within prescribed tolerances in linear algebra. The permeabilities in the Tarbert formation vary several orders of magnitude, but the spatial variation is relatively smooth and Figure 7 shows that the method resolves the qualitative behavior of the pressure solution

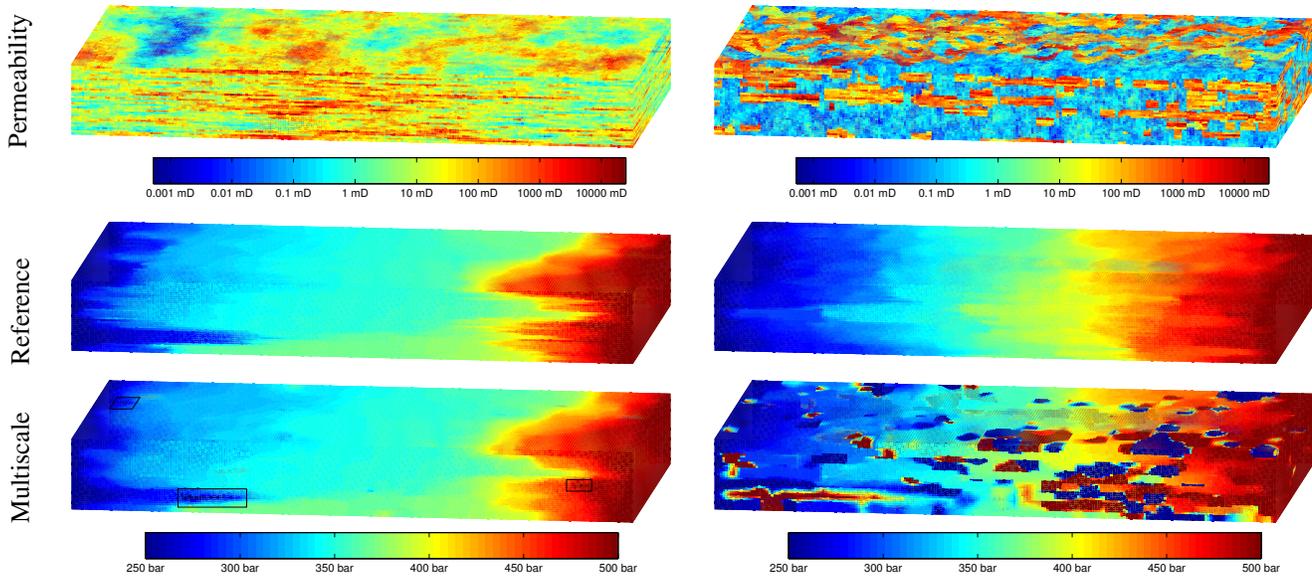


Fig. 7— Pressure computed for the box model with permeability sampled from the Tarbert (left) and the Upper Ness (right) formations of the SPE10 data set. The vertical aspect ratio has been scaled by a factor 4 in these plots to make viewing easier.

Table 1— Discrepancies between fine-scale and MsFV solutions computed on $60 \times 220 \times 35$ box model with three different permeability distributions. The coarse grid has $12 \times 44 \times 7$ blocks.

Model	Error in pressure		Error in flux	
	L_2	L_∞	L_2	L_∞
Homogeneous	1.74858e-12	2.00261e-12	1.55120e-11	1.34408e-10
Tarbert	2.43395e-02	1.84410e+00	4.91018e-01	2.27574e+00
Upper Ness	1.59527e+00	1.02909e+02	2.81082e+01	1.56371e+02

very well. However, while the coarse-scale operator produces physically reasonable values for large parts of the domain, it is not monotone everywhere and violates the range of the boundary conditions in several blocks (e.g., the three small regions marked in the figure). These oscillations are the reason why the L_∞ error in Table 1, is larger than unity. The highly heterogeneous Upper Ness formation is very challenging for the MsFV method as the permeability may exhibit orders of magnitude variations in permeability at small scales. Large areas with high errors have previously been observed for strongly heterogeneous problems in 2D (Kippe et al. 2008; Hajibeygi et al. 2008; Sandvin et al. 2011) and the right column in Figure 7 shows that these problems are aggravated in 3D.

To explain the cause of the large errors, we consider a simplified 2D version of Upper Ness consisting of a set of high-permeable channels on a low-permeable background. Flow is driven from left to right by a unit pressure drop. Figure 8 shows how large contrasts along the edge segments give poor localization, which again makes the coarse-scale operator non-monotone, giving non-physical pressure values far outside the range prescribed by the boundary conditions. Lack of monotonicity for the MsFV coarse-scale operator has been studied previously by Hesse et al. (2008). Here, we note that whereas the general features of the pressure field may be reproduced with near-unity aspect ratios, the effect is greatly amplified with gradually larger aspect ratios and contrasts as seen in the rightmost subplot of Figure 8. For simple grids, it should be possible to automatically generate dual partitions that account for heterogeneities to diminish the non-monotonicity of the coarse-scale operator. How to do the same for realistic stratigraphic models with high complexity is, at best, a topic for future research.

It has previously been suggested to use MsFV as a preconditioner for GMRES to reduce the errors in both pressure and flux reduced towards machine precision (Lunati et al. 2011). Similarly, it has been suggested to use a small number of iterations with an inexpensive smoother to remove oscillatory errors near edges of dual blocks that cannot be removed by the coarse-scale operator and correction functions (Hajibeygi and Jenny 2011). The changes in the pressure distribution are then added to the system as source terms, and a new MsFV solution is computed with the new source terms to ensure conservative flow. In this approach, conservative fine-scale fluxes can be constructed after any global MsFV step, which means that unlike in a multigrid or domain-decomposition method, it is possible to abort the iteration process before the fine-scale residual has converged and still obtain a conservative flux approximation. Unfortunately, several authors have recently observed that such iterative strategies will not work well for strongly heterogeneous cases like Upper Ness (Nordbotten et al. 2012; Zhou and Tchelepi 2012; Sandvin

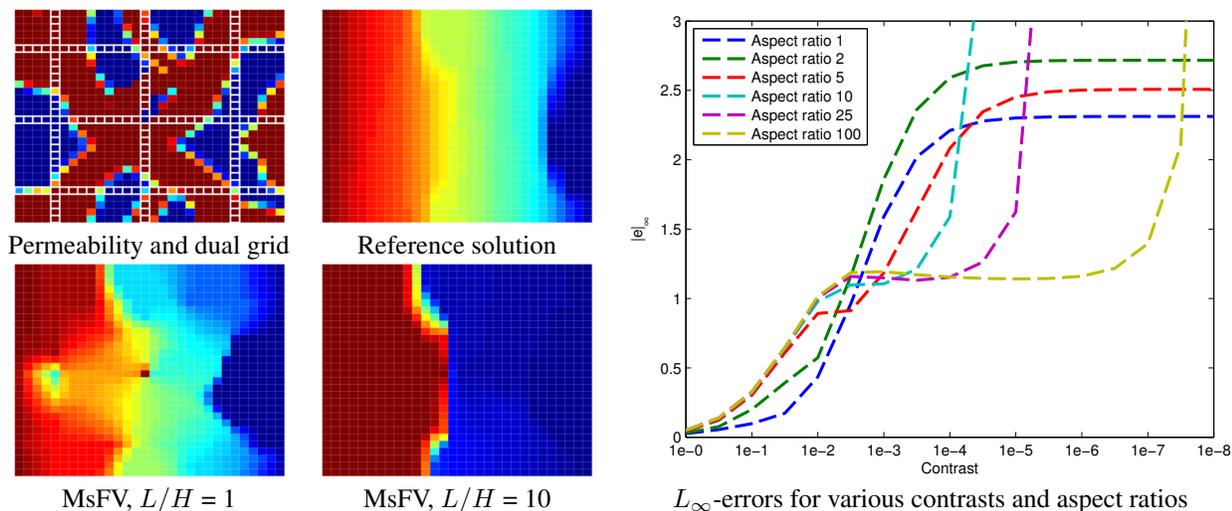


Fig. 8— Dual centers in low-permeable regions can pose problems for the MsFV method, especially for grids with large aspect ratios. The test case is a high-permeable channel (red color=1 darcy) on a low-permeable background (blue color=0.001 mD) inside a box with dimensions $L \times H$. Flow is driven by a unit pressure drop from left to right.

et al. 2013; Wang et al. 2012). The MsFV approximation shown in Figure 7 contains large regions where the pressure solution is orders-of-magnitude wrong because of non-monotonicity in the coarse-scale operator. Correction functions are based on the similar localizations and cannot remove local errors on the interfaces of the dual blocks. Using MsFV plus correction functions as a preconditioner in an iterative approach will therefore not be very efficient (using a constant pressure may, in fact, be a better first-step preconditioner). The iterations will converge, but the solution quality is very poor until a large number of iterations have been completed. In regions with large contrasts in permeability, the solution will be patchy and discontinuous until the late stages of the iteration process. This means that there will be large high-frequency errors at the block boundaries that must be reduced by GMRES or the inexpensive smoother, which neither are efficient without an effective preconditioner.

Altogether, this example summarizes observations we have made for a large number of Cartesian cases in 2D and 3D: for problems that are not strongly heterogeneous, the MsFV method produces qualitatively correct solutions and iterative approaches can be used to converge the fine-scale residual with reasonable efficiency. On the other hand, the method may fail quite spectacularly on certain strongly heterogeneous cases, causing iterations to converge painstakingly slow. Similar observations are discussed more thoroughly by Nordbotten et al. (2012); Zhou and Tchelepi (2012); Sandvin et al. (2013); Wang et al. (2012), who all clearly demonstrate that MsFV may be unsuited as a preconditioner for strongly heterogeneous problems.

Slip-faults. In this example, we will consider slip-faults, which are often seen in realistic models. For simplicity, we will use a synthetic setup consisting of a rectangular domain that contains a single fault across which there has been a significant displacement. To investigate the effect of non-conforming grids across the fault face, we first consider a $10 \times 40 \times 50$ Cartesian model with a vertical fault in the middle of the domain, three different permeability realizations (homogeneous, lognormal, and lognormal layered), and a unit pressure drop from left to right; see Figure 9. To provide different degrees of non-conformity, we start from a conforming grid and gradually increase the heave until we have added the height of a single fine cell. For each displacement, a new dual grid is generated based on a fixed $1 \times 4 \times 1$ primal partition. As can be seen from the rightmost plot in Figure 9, the accuracy of the method is relatively insensitive to changes in the conformity of the grid. In principle, the overlap area between edge cells on opposite sides of the fault could have been only a fraction of a vertical cell face (normalized interface area close to zero), which potentially would have acted as a low-transmissibility barrier that induces oscillations, as discussed above. However, because edges in the dual grid are built based on the intersection of cells with a straight line, the algorithm will in such cases tend to create edges that are at least two cells thick, i.e., add more cells to the edges than what is strictly needed to create a topological connection between coarse nodes. As a result, we avoid edges where the only connection is over a small fraction of a single face, as can be seen in the normalized interface areas, which vary between 0.4 and 1.5. The error increases somewhat with aspect ratio, but the most pronounced variations are caused by heterogeneities in the layered case. Figure 10 illustrates the main error mechanism for a case that in a certain sense is the reverse of Figure 8.

Next, we consider a sloping fault modeled using a corner-point grid in which the pillars are vertical at the east and west sides of the domain and gradually become more skew to align with the fault in the middle of the domain. Figure 11 shows a realization of the structural model that has been partitioned into a primal and a dual coarse grid. To partition across the fault, coarse blocks divided by the fault are split in two to ensure that the primal blocks are relatively homogeneous in petrophysical properties, which is one of the assumptions behind a coarse pressure system in the first place. Because of the displacement, almost all the coarse blocks next to the fault are connected to two blocks on the opposite side, thereby giving a primal coarse-grid with non-neighboring

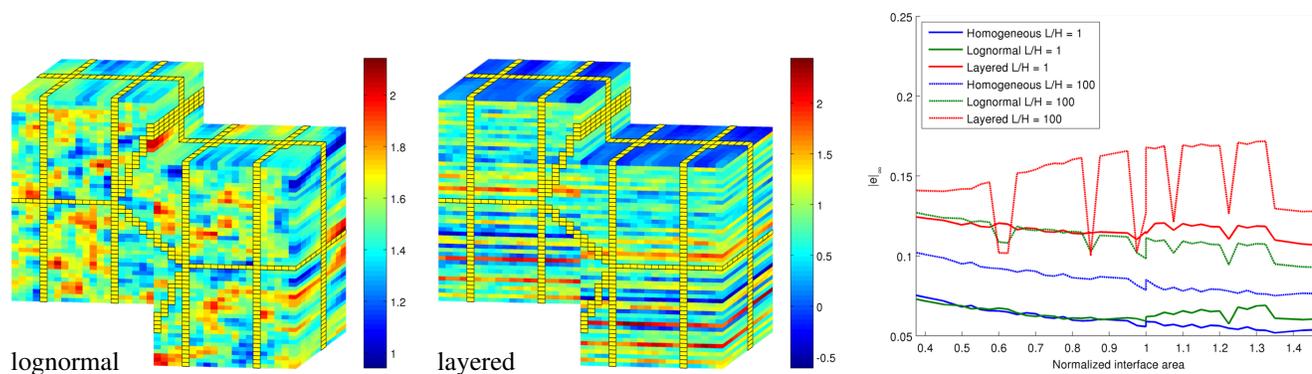


Fig. 9— The left and middle plots show the logarithm of two permeability realizations for a model with a single vertical fault and grid cells with a unit aspect ratio, $L/H = 1$; the faces of the dual grid corresponding to a primal $1 \times 4 \times 1$ grid are shown in yellow. Starting with this matching grid, the fault heave is increased a fraction of a cell height until the cells on opposite sides of the fault match again. The right plot shows the corresponding L_∞ pressure errors as a function of the normalized area of the intersection of the fault plane and the 'edge' cells; a unit area corresponds to an interface equal to a single face of a fine cell.

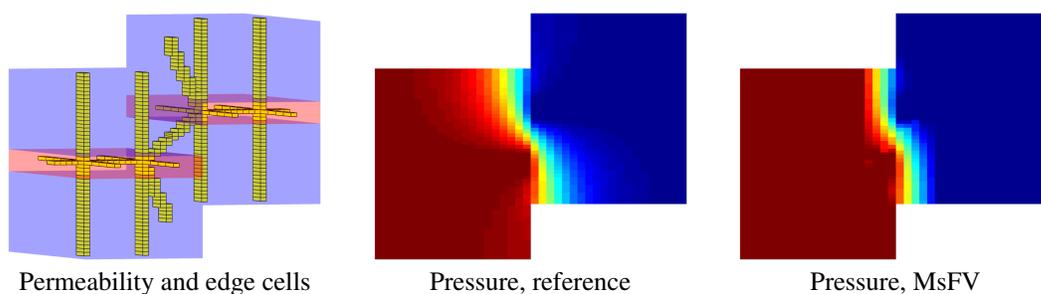


Fig. 10— Errors caused by permeability contrasts at non-matching fault blocks. With a ratio of 10^5 between the high-permeable (red) and low-permeable layers (blue), the localization imposed along the sloping edge gives a low transmissibility between the coarse blocks on opposite sides of the fault in the multiscale coarse system.

connections. As a result, four of the primal blocks shown in the side-view plot will have five edges emanating from their node cell instead of the usual four edges seen for blocks having only logically Cartesian neighbors.

Table 2 reports the discrepancies between the fine-scale and MsFV pressure and flux solutions measured in the relative L_2 and L_∞ norms for a setup using the homogeneous permeability and the two $60 \times 220 \times 35$ subsamples from SPE10 used in the previous example on a $12 \times 44 \times 7$ coarse grid. For the homogeneous permeability, the error is significantly higher than in Table 1 since the analytical solution is no longer a simple linear pressure drop. However, the fact that the error is low indicates that the coarse grids should be applicable for more complex permeability fields. For the Tarbert formation, the MsFV method is able to reproduce the qualitative behavior of the fine-scale solution, but the errors are one order of magnitude larger than for the box model. The displacement at the fault introduces discontinuities in the permeability which have an unfavorable effect on the reduced problems computed along edges. This effect will also be amplified by the skewed cell geometries and the reduced face areas of non-matching grid blocks at the fault. For the Upper Ness formation, the strong pressure oscillations observed in the previous example are amplified even further, rendering the resulting approximation more or less useless.

To improve upon this situation, let us revisit the iterative method using MsFV as a preconditioner to GMRES. As explained above, it is the localization assumption that causes oscillations. In particular, the first step of the process that extrapolates from

Table 2— Discrepancies between the fine-scale and the MsFV solution computed for the single-fault model with three different permeability distributions realized on a $60 \times 220 \times 35$ fine grid that is partitioned into a $12 \times 44 \times 7$ coarse grid.

Model Permeability	Error in pressure		Error in flux	
	L_2	L_∞	L_2	L_∞
Homogeneous	1.35204e-03	3.01027e-03	4.42839e-02	3.89119e-01
Tarbert	1.68418e-01	2.26869e+01	1.68660e+00	3.93173e+00
Upper Ness	6.14998e+01	3.35768e+03	3.36793e+02	2.04909e+03

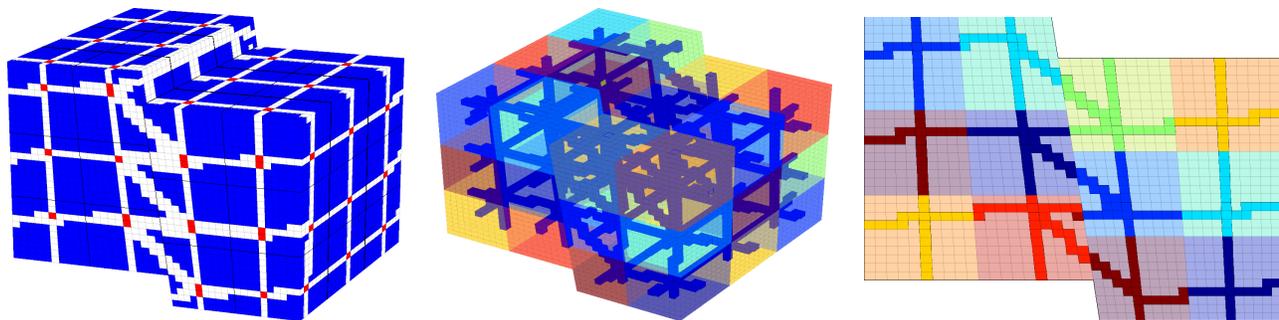


Fig. 11— Primal and dual coarse grids generated for a $40 \times 30 \times 31$ realization of the single-fault model. The left plot shows the 'inner' cells (blue), the 'face' cells (white), the 'edge' cells (red), and the primal coarse grid (thick line) for a subset of the whole model. The middle plot shows the different blocks in the primal $4 \times 3 \times 3$ coarse grid colored in different transparent colors to show the 'edge' cells inside, colored in blue. In the right plot, the 'edge' cells have been given a different color for each primal block.

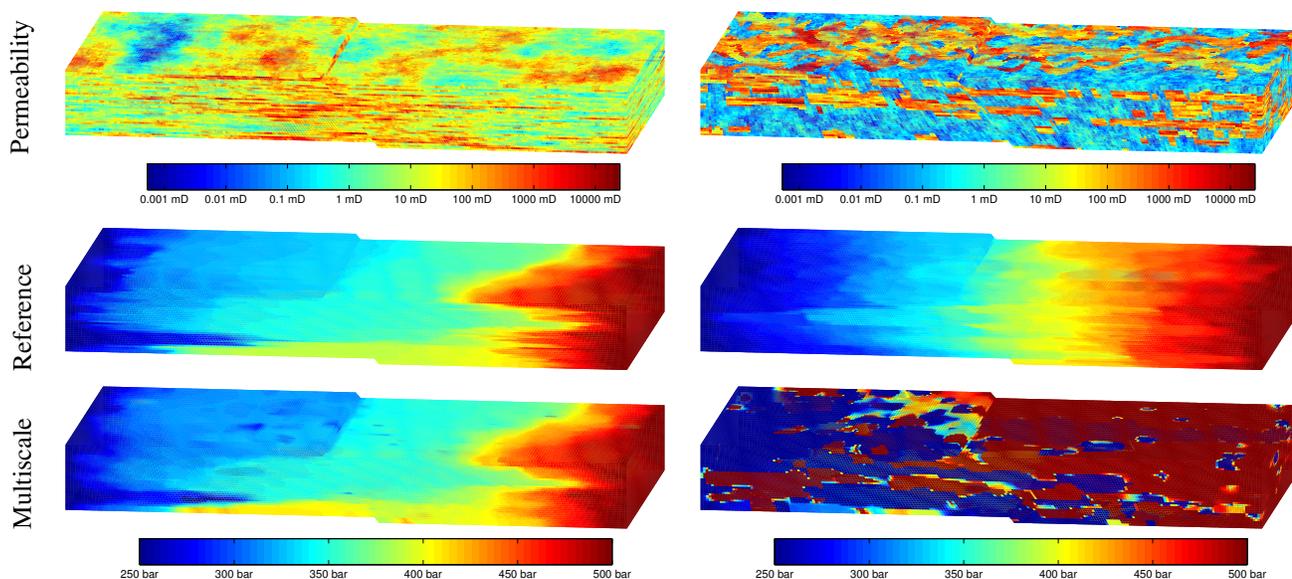


Fig. 12— Pressure computed for the single-fault model with permeability sampled from the Tarbert (left) and the Upper Ness (right) formations of the SPE10 data set. The vertical aspect ratio has been scaled by a factor 4 in these plots to make viewing easier.

nodes to edge cells is very sensitive to large variations in permeability. To mitigate these problems, we propose to replace the permeability values used to compute the localization of basis functions by the distance weighted harmonic average over all cells in a neighborhood around each cell. That is, if \mathbf{x}_i denotes the centroid and K_i the permeability of a node, edge, or face cell c_i , and $\mathcal{N}(i)$ defines the neighborhood used to compute the average, we replace K_i by a new value

$$\tilde{K}_i = \frac{1}{|\mathcal{N}(i)| \sum_{j \in \mathcal{N}(i)} w_j / K_j}, \quad w_j = 1 - \frac{\|\mathbf{x}_j - \mathbf{x}_i\|_2}{\sum_{j \in \mathcal{N}(i)} \|\mathbf{x}_j - \mathbf{x}_i\|_2}. \quad (10)$$

This retains most of the permeability values used to compute the localization in relatively homogeneous areas, but smooths the permeability in areas of high contrast. For the edge and node cells, this is done using cells that are within a logical distance of five cells. For the face cells we use a logical distance of three cells, whereas the inner cells are left untouched. The transmissibilities corresponding to the smoothed permeability $\tilde{\mathbf{K}}$ are then used for the multiscale solutions in a preconditioned GMRES solve. We retain the original system matrix for calculation of residuals to ensure that we are still solving the correct fine-scale system. This will obviously converge slower as the system nears machine precision, but will give much better solutions for the first iterations when solving highly challenging petrophysical cases.

To investigate the effect of the new strategy, we consider a somewhat smaller model with a $30 \times 110 \times 15$ permeability realization sampled from the 15 lowest layers of Upper Ness. Approximate solutions obtained after 100 GMRES iterations for the original and the improved method are plotted in Figure 13, and the associated errors are reported in Table 3. For completeness,

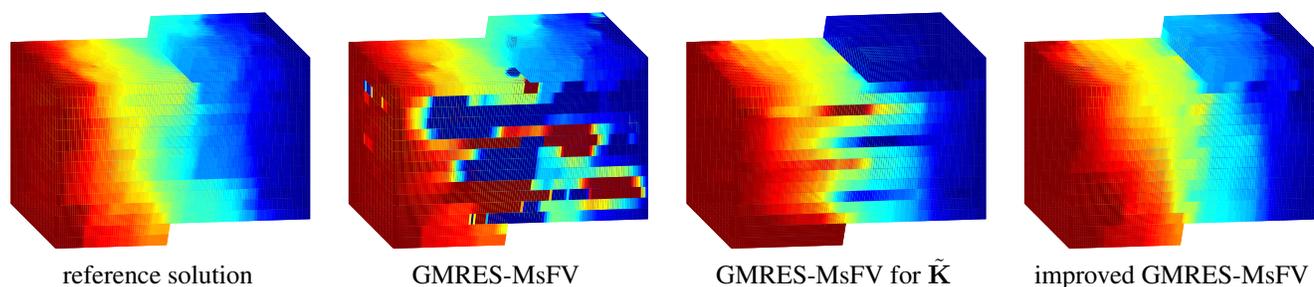


Fig. 13— Approximate pressure solutions computed by the iterative MsFV method using 100 GMRES iterations and a $6 \times 22 \times 15$ coarse grid on the faulted $30 \times 110 \times 15$ Upper Ness model.

Table 3— Discrepancy between the fine-scale solution and an approximate solution computed with the iterative MsFV method using 100 GMRES iterations and a $6 \times 22 \times 15$ coarse grid on the faulted $30 \times 110 \times 15$ Upper Ness model.

Iterative scheme		Error in pressure		Error in flux	
Basis functions	System matrix	L_2	L_∞	L_2	L_∞
Original	Original	2.60704e+00	7.19786e+01	1.39525e+00	2.91762e+00
Smoothed	Smoothed	9.26928e-02	6.14530e-01	9.97036e-01	9.99043e-01
Smoothed	Original	7.72174e-02	4.43286e-01	9.96027e-01	9.98863e-01

we have also included a simulation in which $\tilde{\mathbf{K}}$ is used to generate the system matrix. Although this obviously will converge to an incorrect solution, the result is a much better approximation than what is obtained by GMRES-MsFV using the original permeability. In fact, the L_2 error of the pressure is reduced by a factor 30 and the L_∞ by a factor 100. Using the smoothed permeability only for computing basis functions gives even better results. Comparing with Table 2, we also observe that since GMRES minimizes residuals in a fine-scale conservation equation, the error reduction is larger in flux than in pressure.

Because it is easy to detect where the error is large after an initial solve, a possible preconditioner employing the smoothing strategy proposed above could smooth the permeability only in certain areas for the first iterations. This would result in a much smoother residual, and hence the rest of the iterations could use the original transmissibilities, accelerating convergence in high contrast areas.

A curved fault. Reservoir simulations on realistic grids provide other cell geometries and more complex topologies than seen for the simple fault in the previous example. For instance, cells can often be general polyhedral rather than hexahedral. Likewise, there may not be an inherent cell ordering that can be easily exploited to generate primal coarse blocks. To investigate the capabilities of the MsFV solver on unstructured grids, we will first look at another synthetic case with a vertical fault that describes an S-curve in the horizontal plane. To represent the fault, we will use two different fine grids. The first grid was created by fitting a triangular grid to the S-curve in the horizontal plane and then extruding it to a prismatic grid in 3D. The second grid was created by first constructing the Voronoi diagram for the areal triangulation and then extruding the resulting polygonal grid to a 2.5D polyhedral PEBI grid. The latter grid will not represent the fault exactly.

Coarse partitions with approximately 100 coarse blocks for both grids were created using Metis (Metis 2012) with two layers in the vertical direction for the PEBI grid and one for the prismatic grid. This gives fully unstructured primal coarse grids with complex connections, in particular for the PEBI grid. For the PEBI grid, we also generate a more regular partition, using cell centers to sample from a 10×10 partition of the grid's bounding box. All block intersecting the fault were split in all three partitions, and blocks with very few fine cells (< 50) were merged with their nearest neighbor. The resulting coarse grids exhibit no Cartesian features in either of the cases; this can be seen in Figure 14. As in the previous example, flow was driven from left to right by a pressure drop imposed as Dirichlet boundary conditions, here at 1000 and 250 bar respectively. Likewise, we use three different permeability realizations: (i) constant permeability equal 100 mD in all three spatial directions; (ii) permeability sampled from the Tarbert layers of SPE10; and (iii) permeability sampled from the Upper Ness layers of SPE10. Unlike in the previous example, the sampling from the SPE10 data does not account for displacement of layers across the fault, and the permeability fields will hence not be discontinuous at the fault. Table 4 reports pressure and flux errors of the initial MsFV solution for all three grids for the homogeneous permeability and the two PEBI grids for the two SPE10 subsamples.

With homogeneous permeability, the multiscale solutions are almost indistinguishable from the fine-scale solutions in the visual norm and plots are not included. However, comparing the structured and unstructured coarse grids for PEBI, we see that pressure error of the latter is more than ten times higher in the L_∞ norm. The reason is that the primal coarse grid generated by Metis is not 2.5D, or in other words, the lateral and upward faces do not always follow the layers and the vertical pillars of the fine-grid, respectively. Hence, the coarse-scale stencil will have larger grid-orientation effects induced by the reduced boundary

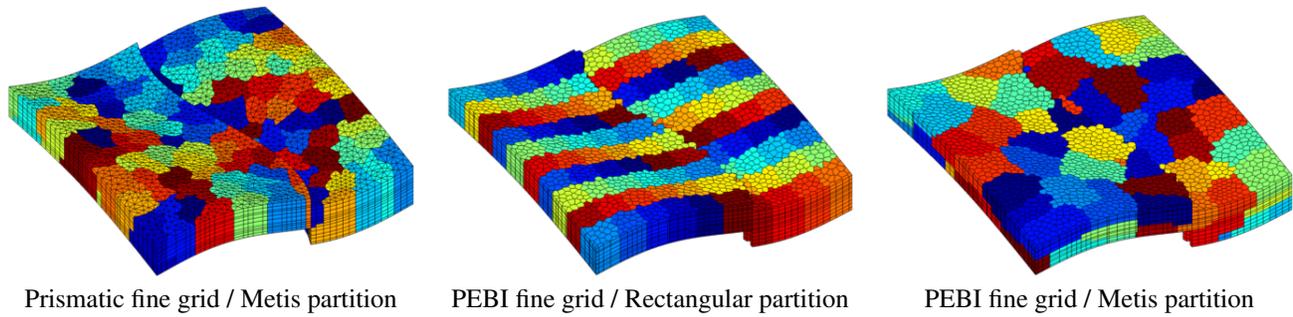


Fig. 14— A curved fault represented on a prismatic grid and on a 2.5D PEBI fine grid. The colors in the plots represent different primal coarse blocks; there are approximately 100 blocks in each grid.

Table 4— Discrepancies between the fine-scale and the MsFV solution computed for two different unstructured grids describing a curved fault for homogeneous permeability and anisotropic permeability subsampled from SPE10.

Model Perm	Cell type	Setup of simulation		Error in pressure		Error in flux	
		Partition	Smoother	L_2	L_∞	L_2	L_∞
100 mD	Prismatic	Metis	No	1.25392e-02	5.30042e-02	4.66819e-01	2.57628e+00
			5 bGS	1.19782e-02	3.09221e-02	1.35012e-01	8.74971e-01
100 mD	PEBI	Rectangular	No	1.20388e-02	4.07551e-02	4.36458e-01	7.13954e-01
			5 bGS	1.09899e-02	1.95476e-02	7.49334e-02	1.58961e-01
100 mD	PEBI	Metis	No	1.79198e-02	6.16854e-01	8.33699e-01	2.92127e+00
			5 bGS	2.65210e-02	7.23370e-02	3.68878e-01	1.63392e+00
Tarbert	PEBI	Rectangular	No	9.70071e-02	4.65787e-01	3.56573e+00	8.34111e+00
			5 bGS	6.32725e-02	1.83252e-01	8.05076e-01	1.70222e+00
Tarbert	PEBI	Metis	No	1.91794e-01	1.32520e+00	1.53620e+00	2.83060e+00
			5 bGS	7.22927e-02	6.99239e-01	1.31870e+00	3.68596e+00
Ness	PEBI	Rectangular	No	7.47695e-01	1.16558e+01	4.46944e+00	6.00056e+00
			5 bGS	1.76817e-01	3.16990e+00	4.31521e+00	1.94548e+01
Ness	PEBI	Metis	No	4.01894e+01	9.08244e+02	2.39994e+01	4.98474e+01
			5 bGS	1.07917e+02	1.80408e+03	3.21459e+03	1.17321e+04

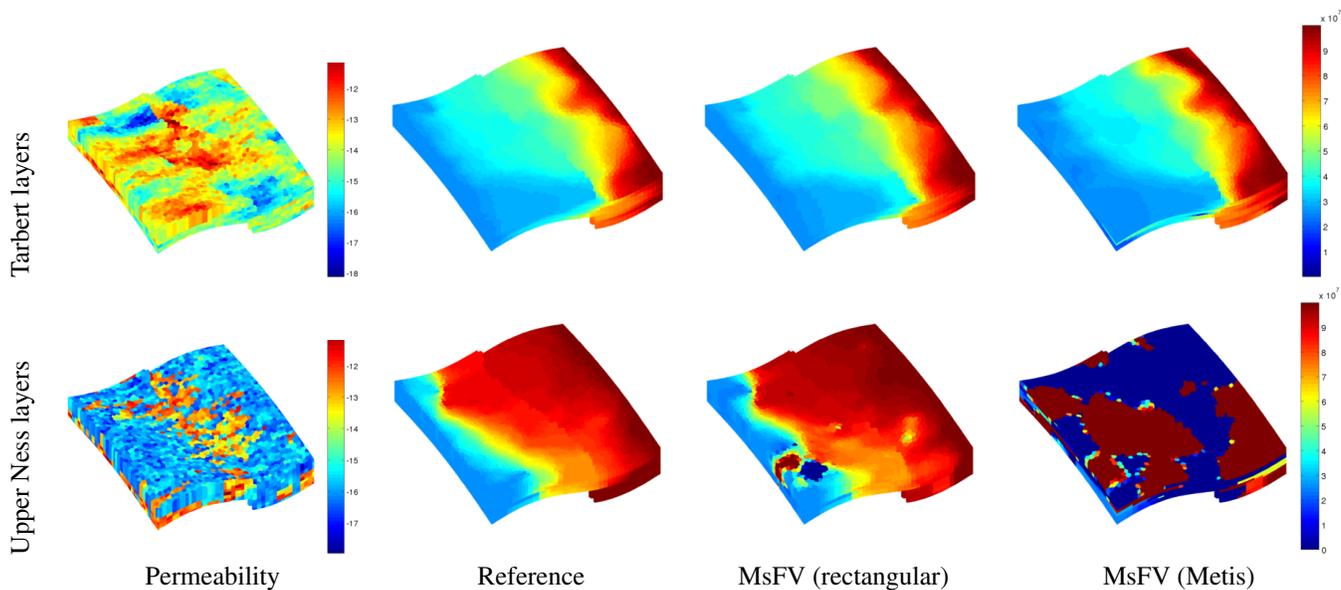


Fig. 15— A curved fault represented on a 2.5D PEBI grid with permeability sampled from SPE10. The last three columns show pressure in units Pascal computed by a TPFA method on the fine grid and by the MsFV method on two different coarse grids with approximately 100 primal blocks blocks.

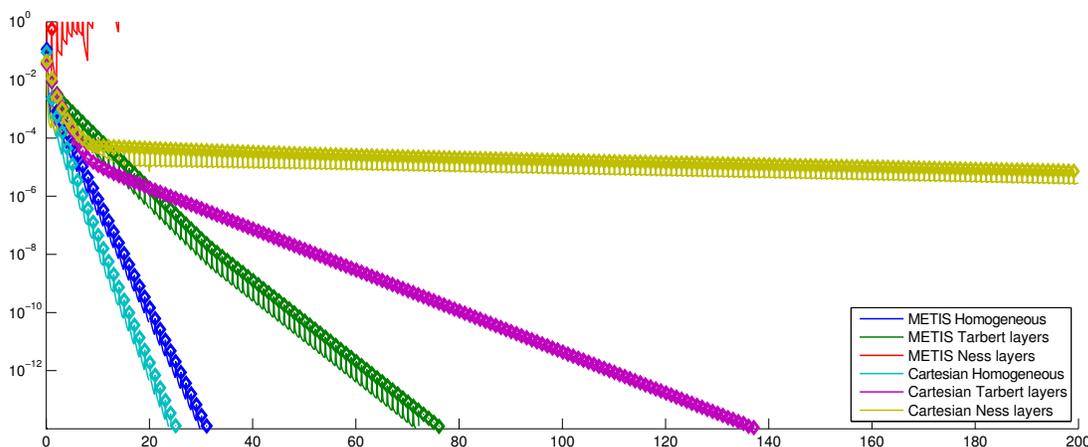


Fig. 16— Convergence history for the iterative MsFV method for the curved fault modeled on a PEBI grid. Ten smoothing steps with a block Gauss–Seidel method are used per iteration with blocks equal the primal coarse blocks. The solid lines show the residual after each smoothing steps, while the markers show the residual after the MsFV solve.

conditions used to localize the basis functions. As discussed above, we can use a smoother to remove errors induced at the faces of the dual coarse blocks and use the resulting difference in the solution as source terms for a new MsFV solution reusing the existing basis functions. Herein, we will use a simple block Gauss–Seidel (bGS) method which has a low computational cost. Table 4 also reports errors after five smoothing steps followed by a conservative reconstruction. The purpose of the smoother is to propagate high-frequency errors away from the faces of the dual blocks. As a result, we see that the L_∞ error decreased by almost a factor ten for the PEBI/Metis grids, while the L_2 pressure error *increases* slightly. As expected, the smoother has a smaller effect on the flux errors and on the pressure errors of the other two grids.

Table 4 also reports errors for the two PEBI grid populated with anisotropic permeabilities sampled from the Tarbert and Upper Ness layers of SPE10. Approximate solutions are shown in Figure 15. For Tarbert, the approximate solutions are still qualitatively correct, except for some small pressure oscillations seen on the Metis coarse grid. Notice also that the L_∞ flux error increases slightly, which may indicate that a few more smoothing steps are needed. For Upper Ness, the pressure field has large patches of oscillatory behavior, similar to what was observed and discussed above for the box and the single-fault models. We notice, in particular, that these problems are significantly amplified for the irregular coarse blocks generated by Metis. In fact, the errors introduced by the coarse-scale operator are so severe that the smoother diverges. In Figure 16, we investigate the convergence of the iterative approach. The figure reports the fine-scale residual for the two PEBI grids on all three permeability realizations using a setup with ten bGS smoothing steps per iteration. For homogeneous permeability, the iterative method converges within a few tens of iterations, as expected. For Tarbert, however, we observe an interesting behavior. Whereas the initial reduction of the residual is faster on the regular partition, the Metis grid has better asymptotic convergence because of the larger number of block connections used in the bGS iterations. In terms of efficiency, however, the cost of the extra couplings will to a certain degree balance out the improved convergence rate. For Upper Ness, the Metis grid diverges, as expected, whereas the convergence for the regular partition is painstakingly slow. The asymptotic convergence is dominated by the properties of the smoother, which in general will converge slow, with a rate that decays with increasing heterogeneity and aspect ratios. Nevertheless, as shown in the figure, the residual is efficiently reduced in the first 5 to 15 iterations when the initial MsFV solutions have significant small-scale errors near the faces of the dual coarse blocks induced by the localization assumptions.

A realistic bedding model Pinch-outs will, like faults, create unstructured non-neighboring connections and hence be one of the principal gridding challenges for real-life reservoir models. To exemplify, we consider a highly detailed, core-scale model of realistic bedding structures. Such models are used as input to derive directional permeability for a given lithofacies and identify net pay below the level of petrophysical log resolution. The bedding structure is realized on a $30 \times 30 \times 333$ corner-point grid in which (almost) all cells are affected in some degree by pinch: Although the model has around 300,000 cells initially, over two thirds will be inactive due to significant erosion, giving a fine grid of approximate dimensions $30 \times 30 \times 100$. The volumes of the cells, as well as the areas of the (vertical) faces, vary almost four orders of magnitudes. We partition the fine grid into $3 \times 3 \times 6$ coarse blocks and apply Dirichlet boundary conditions imposing a pressure drop from west to east.

The model has a large number of low-permeable shale layers pinched between the other high-permeable layers. As seen in the earlier examples, impermeable regions intersecting the dual grid may pose problems for the monotonicity of the coarse scale MsFV operator. We therefore consider three different permeability realizations: First a homogeneous, isotropic case to gauge the quality of the coarse grid; then a heterogeneous, anisotropic model in which the shale layers are replaced by high permeability; and finally the full model. As can be seen in Table 17, while the problems without shale give relatively good solution quality considering the upscaling factor, the shales introduce non-monotonicity and large errors. However, by using transmissibilities

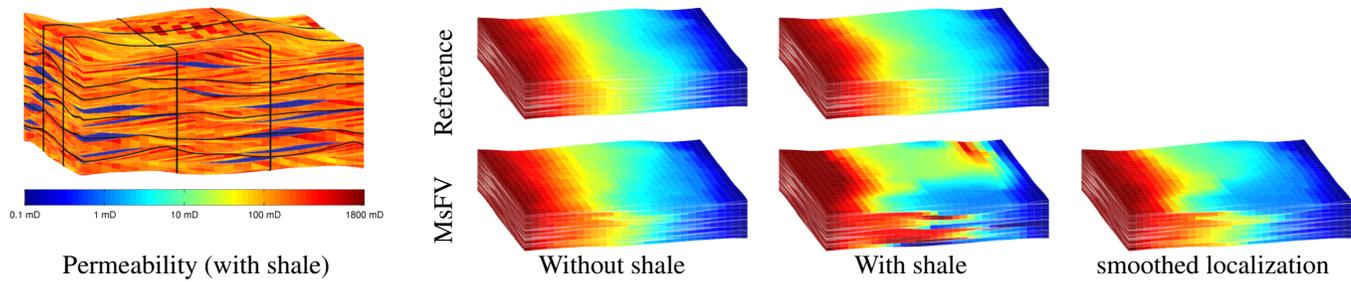


Fig. 17— A core-scale model of realistic bedding structure containing a high number of pinched layers and non-neighboring connections. The layers of shale (left) will drastically influence the solution quality for the MsFV operator, unless alternate boundary conditions are employed (right).

Table 5— Discrepancies between fine-scale and MsFV solutions for a geological model of realistic core-scale bedding structure.

Setup of simulation		Error in pressure		Error in flux	
Permeability	Localization	L_2	L_∞	L_2	L_∞
Homogeneous	Standard	5.68702e-02	6.47287e-01	1.19806e+00	6.67033e+00
Without shale	Standard	5.82791e-02	6.39894e-01	1.07475e+00	5.15355e+00
Full model	Standard	1.42394e+01	6.89412e+02	2.33271e+00	7.92730e+00
Full model	Smoothed	7.67519e-02	6.54025e-01	1.34098e+00	7.02256e+00

derived from smoothed permeabilities for the edge problems, as discussed above, the shales will no longer drastically impact the solution quality, which can be seen in the leftmost part of Figure 17.

A field model. In the last example, we will use the corner-point grid from the simulation model of the Norne reservoir from the Norwegian Sea along with a set of realistic petrophysical properties, as shown in Figure 18. Compared with the previous two examples, the simulation model has a significantly more complex geometry with erosions, pinch-outs, faults, and inactive cells. For instance, while 75% of the fine cells have six faces as in a Cartesian grid, the number of faces per fine cell varies all over from as little as 4 to as much as 19 because of several faults and pinched layers.

To obtain a coarse grid, we will first divide the model logically into coarse blocks as shown in the left part of Figure 19. To ensure that coarse blocks are split over faults as shown in the right plot of Figure 19, any connections across faults is temporarily removed, and the grid is post-processed to ensure that all cells within each coarse block are connected to all the other cells within the same block. If not, the block is split into new blocks that each consists of a connected set of cells.

Two producers operating at fixed bottom-hole pressure $p_{bhp} = 100$ bar are added to the reservoir along with a 500 bar pressure boundary condition at one of the outer boundaries at the opposite side of the wells, as can be seen along with permeability in Figure 18. This gives flow across the entire domain. The wells are added in using the correction terms as in Eq. 6. That is, one solves additional local problems in which the boundary conditions are set to zero at the nodes and source terms following a Peaceman well model are included. This allows the method to distinguish between multiple wells inside a single coarse block. Accurate resolution of wells is paramount to get correct flow patterns. Other possible approaches would be to include the wells directly in the coarse equations and instead rely on iterative techniques or coarse-grid refinement to capture fine-scale effects. As can be seen from Figure 20, the initial MsFV solution contains sub-resolution details near the wells.

We will first compute approximate solutions for uniform permeability equal to 100 mD to ensure that we can judge the quality of the coarse grids separately from the behavior of the method itself on the petrophysical parameters. For uniform permeability, Figure 20 shows that the multiscale solution strongly resembles the fine-scale solution, apart from a few small areas where the localization assumption reduces the accuracy of the method. This indicates that the dual grid has good quality in terms of giving a coarse pressure system and using the result to find fine-scale pressure.

The next step is to solve the same system with realistic permeability. The permeability is layered, heterogeneous, and anisotropic with six orders of magnitude difference between the minimum and maximum values. The resulting initial MsFV solution (Figure 20) seems to overestimate the pressure in the lower high-permeability layers. It should also be noted that the pressure is overestimated near one of the producers, which can be attributed to the fact that the MsFV method treats wells for the coarse system in an integral sense. As before, we can perform a series of inexpensive bGS iterations to produce additional source terms for a better multiscale solution: After ten smoothing steps the resulting multiscale solution is visually much closer to the reference and the flux error reported in Table 6 has been reduced by a factor ten.

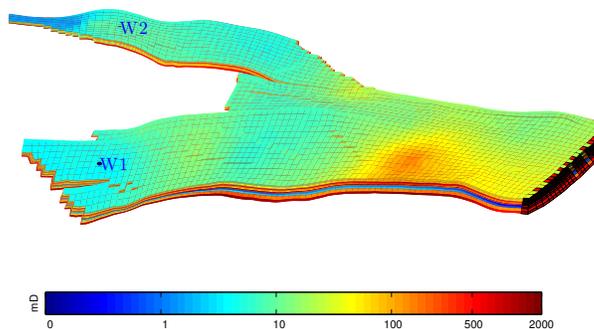


Fig. 18— Horizontal permeability and boundary conditions for the simulation model from the Norwegian Sea. Dirichlet boundary condition is imposed on one end of the reservoir (shown in red), while two producers (W1 and W2) are kept at fixed bottom-hole pressure at the other end of the model.

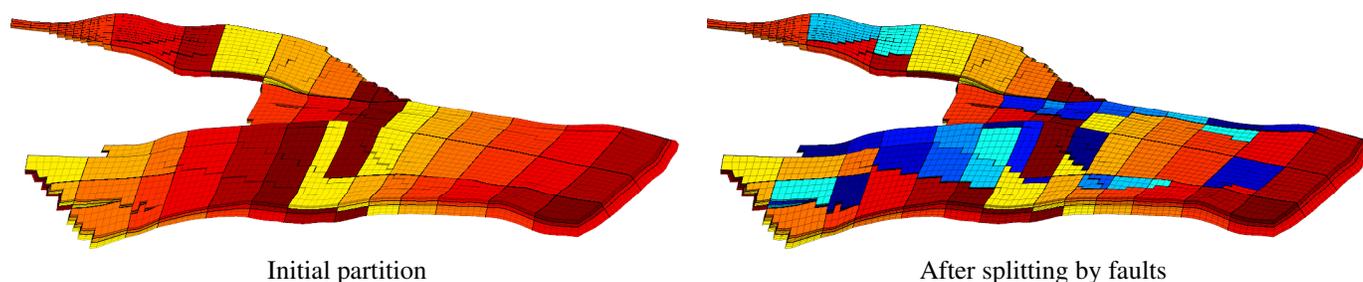


Fig. 19— The corner-point model is partitioned into blocks using the logical IJK cell numbering and then post-processed to create new coarse blocks along the faults.

Concluding Remarks

We have discussed how the algebraic (operator) formulation of the MsFV method can be applied to stratigraphic and unstructured grids provided a pair of primal and dual coarse grids is available. The main challenge is to generate primal-dual partitions that are consistent with the algebraic manipulations used to localize the multiscale method and produce a sensible coarse-scale system. To this end, we have presented an algorithm that creates a dual coarse grid from a given primal coarse grid. The algorithm is applicable to a large class of corner-point grids with faults, erosions, pinch-outs, inactive cells, etc., as well as more general grids with polyhedral cells. However, it may fail or be very time consuming if the primal partition is not a good starting point. To create a good primal coarse grid, one should ensure that the blocks are sufficiently regular and have faces large enough so that edges and interfaces of the dual grid can be constructed between nodes, i.e., the fine cells that encompass the centroids of the primal blocks. In addition, the quality of the primal grid can be increased by aligning coarse blocks with layers and pillars in the fine stratigraphic grid, splitting coarse blocks across faults, adapt to the permeability field to avoid large permeability contrasts within the primal blocks and hence reduce undesired effects of the localizations, etc.

A large number of numerical experiments, of which a few are reported herein, shows that the MsFV method, extended with our dual-partition algorithm, produces accurate approximations for stratigraphic grids with realistic heterogeneities and aspect ratios, if the heterogeneities are not too strong. Moreover, the MsFV solver can be applied as a preconditioner in iterative approaches to systematically drive the fine-scale residual toward zero. The main advantage of this approach is that one, unlike for multigrid and decomposition methods, can abort the iteration early and still obtain a conservative flux on the fine and coarse scale.

However, our numerical experiments also show that one can easily construct cases for which the current MsFV method produces solutions that deviate strongly from the fine-scale solution. This occurs in particular when the permeability field has strong

Table 6— Discrepancies between fine-scale and MsFV solutions on the real-field model.

Setup of simulation		Error in pressure		Error in flux	
Permeability	Smoothed	L_2	L_∞	L_2	L_∞
Homogeneous	No	6.80897e-02	7.86341e-01	2.65688e+01	5.01178e+01
Homogeneous	10 bGS	5.52130e-02	2.62088e-01	4.02353e+00	6.98880e+00
Realistic	No	8.87627e-02	6.03834e-01	5.29446e+00	7.91485e+00
Realistic	10 bGS	7.05802e-02	1.70915e-01	1.04733e+00	8.59376e-01

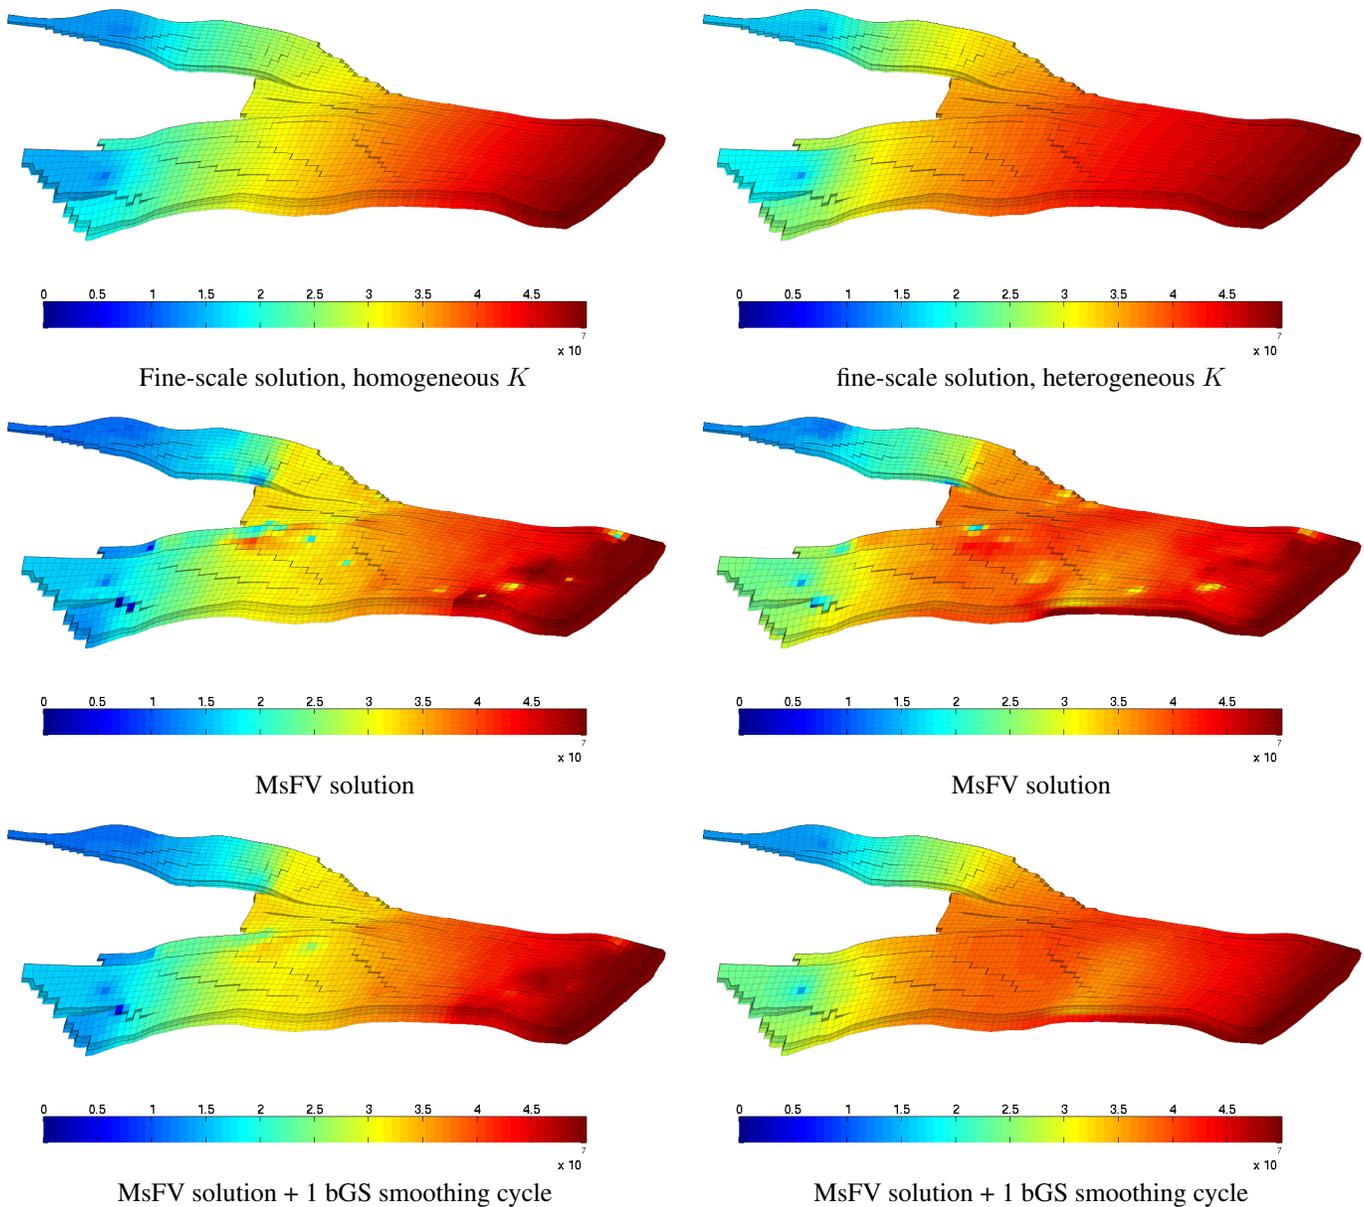


Fig. 20— Fine-grid and multiscale solutions computed on the field model with homogeneous permeability (left) and realistic permeability distribution (right).

low-high-low or high-low-high variations along the edges (and faces) of the dual coarse blocks, which effectively render the localization assumptions strongly erroneous. Such cases can be detected quite easily during grid partition, but are not simple to mitigate in an automatic manner by changing the coarse dual-primal partition. Further work should therefore focus on developing improved localizations and alternative coarse-scale operators that are better suited for high-contrast cases. Likewise, our experience with the MsFV on complex grids confirms lessons learned for the multiscale mixed-finite element method: Creating good coarse grids that both give high accuracy and reduced runtime requires a certain user proficiency. To make the method more robust and accessible to non-expert users, an important task for future research is therefore to develop automated or semi-automated methods for performing the internal grid coarsening and a set of guidelines for setting the necessary parameters, e.g., as discussed in (Aarnes et al. 2008; Hauge 2010).

Finally, the MsFV method has been implemented as a free and open-source module in MRST. The module is available online and the solvers discussed in the current paper can easily be used by other researchers for numerical experiments or continued research on the MsFV method.

Acknowledgments

The authors thank Bård Skaflestad, Halvor Møll Nilsen, Jostein R. Natvig, and Stein Krogstad for helpful discussions. The authors would also like to thank Statoil (operator of the Norne field) and its license partners ENI and Petoro for the release of the Norne data. Further, the authors acknowledge the Center for Integrated Operations at NTNU for cooperation and coordination of the Norne Cases.

References

- Aarnes, J. E., Krogstad, S., and Lie, K.-A. 2006. A hierarchical multiscale method for two-phase flow based upon mixed finite elements and nonuniform coarse grids. *Multiscale Model. Simul.*, 5(2):337–363 (electronic).
- Aarnes, J. E., Krogstad, S., and Lie, K.-A. 2008. Multiscale mixed/mimetic methods on corner-point grids. *Comput. Geosci.*, 12(3):297–315.
- Alpak, F. O., Pal, M., and Lie, K.-A. 2012. A multiscale method for modeling flow in stratigraphically complex reservoirs. *SPE J.*, 17(4):1056–1070.
- Christie, M. A. and Blunt, M. J. 2001. Tenth SPE comparative solution project: A comparison of upscaling techniques. *SPE Reservoir Eval. Eng.*, 4:308–317. Url: <http://www.spe.org/csp/>.
- Efendiev, Y. and Hou, T. Y. 2009. *Multiscale Finite Element Methods*, volume 4 of *Surveys and Tutorials in the Applied Mathematical Sciences*. Springer Verlag.
- Hajibeygi, H., Bonfigli, G., Hesse, M. A., and Jenny, P. 2008. Iterative multiscale finite-volume method. *J. Comput. Phys.*, 227(19):8604–8621.
- Hajibeygi, H., Deb, R., and Jenny, P. 2011a. Multiscale finite volume method for non-conformal coarse grids arising from faulted porous media. In *SPE Reservoir Simulation Symposium, The Woodlands, TX, USA, 21–23 February 2011*. SPE 142205-MS.
- Hajibeygi, H. and Jenny, H. A. 2013. Compositional multiscale finite-volume formulation. In *SPE Reservoir Simulation Symposium, The Woodlands, TX, USA, 18–20 February 2013*. SPE 163664-MS.
- Hajibeygi, H. and Jenny, P. 2009. Multiscale finite-volume method for parabolic problems arising from compressible multiphase flow in porous media. *J. Comput. Phys.*, 228(14):5129–5147.
- Hajibeygi, H. and Jenny, P. 2011. Adaptive iterative multiscale finite volume method. *J. Comput. Phys.*, 230(3):628–643.
- Hajibeygi, H., Karvounis, D., and Jenny, P. 2011b. A hierarchical fracture model for the iterative multiscale finite volume method. *J. Comput. Phys.*, 230(24):8729–8743.
- Hauge, V. L. 2010. *Multiscale Methods and Flow-based Gridding for Flow and Transport in Porous Media*. PhD thesis, Norwegian University of Science and Technology.
- Hesse, M. A., Mallison, B. T., and Tchelepi, H. A. 2008. Compact multiscale finite volume method for heterogeneous anisotropic elliptic equations. *Multiscale Model. Simul.*, 7(2):934–962 (electronic).
- Jenny, P., Lee, S. H., and Tchelepi, H. A. 2003. Multi-scale finite-volume method for elliptic problems in subsurface flow simulation. *J. Comput. Phys.*, 187:47–67.
- Jenny, P. and Lunati, I. 2009. Modeling complex wells with the multi-scale finite-volume method. *J. Comput. Phys.*, 228(3):687 – 702.
- Kippe, V., Aarnes, J. E., and Lie, K.-A. 2008. A comparison of multiscale methods for elliptic problems in porous media flow. *Comput. Geosci.*, 12(3):377–398.
- Künze, R. and Lunati, I. 2012. An adaptive multiscale method for density-driven instabilities. *J. Comput. Phys.*, 231(17):5557–5570.
- Lee, S. H., Wolfsteiner, C., and Tchelepi, H. 2008. Multiscale finite-volume formulation for multiphase flow in porous media: Black oil formulation of compressible, three phase flow with gravity. *Comput. Geosci.*, 12(3):351–366.
- Lie, K., Krogstad, S., Ligaarden, I., Natvig, J., Nilsen, H., and Skaflestad, B. 2012. Open-source MATLAB implementation of consistent discretisations on complex grids. *Comput. Geosci.*, 16:297–322.
- Lunati, I. and Jenny, P. 2006. Multiscale finite-volume method for compressible multiphase flow in porous media. *J. Comput. Phys.*, 216(2):616–636.
- Lunati, I. and Jenny, P. 2008. Multiscale finite-volume method for density-driven flow in porous media. *Comput. Geosci.*, 12(3):337–350.
- Lunati, I. and Lee, S. H. 2009. An operator formulation of the multiscale finite-volume method with correction function. *Multiscale modeling & simulation*, 8(1):96.
- Lunati, I., Tyagi, M., and Lee, S. H. 2011. An iterative multiscale finite volume algorithm converging to the exact solution. *J. Comput. Phys.*, 230(5):1849–1864.

- Metis 2012. Metis – serial graph partitioning and fill-reducing matrix ordering. url: <http://glaros.dtc.umn.edu/gkhome/views/metis>.
- Møyner, O. 2012. Multiscale finite-volume methods on unstructured grids. Master's thesis, Norwegian University of Science and Technology.
- MRST 2013. The MATLAB Reservoir Simulation Toolbox, version 2013a. <http://www.sintef.no/MRST/>.
- Natvig, J. R., Lie, K.-A., Krogstad, S., Yang, Y., and Wu, X.-H. 2012. Grid adaption for upscaling and multiscale methods. In *Proceedings of ECMOR XIII–13th European Conference on the Mathematics of Oil Recovery*, Biarritz, France. EAGE.
- Natvig, J. R., Skaflestad, B., Bratvedt, F., Bratvedt, K., Lie, K.-A., Laptev, V., and Khataniar, S. K. 2011. Multiscale mimetic solvers for efficient streamline simulation of fractured reservoirs. *SPE J.*, 16(4).
- Nordbotten, J. M. and Bjørstad, P. 2008. On the relationship between the multiscale finite-volume method and domain decomposition preconditioners. *Comput. Geosci.*, 12(3):367–376.
- Nordbotten, J. M., Keilegavlen, E., and Sandvin, A. 2012. Mass conservative domain decomposition for porous media flow. In Petrova, R., editor, *Finite Volume Method-Powerful Means of Engineering Design*. InTech Europe, Rijeka, Croatia.
- Sandve, T. H., Berre, I., Keilegavlen, E., and Nordbotten, J. M. 2013. Multiscale simulation of flow and heat transport in fractured geothermal reservoirs: inexact solvers and improved transport upscaling. In *Thirty-Eighth Workshop on Geothermal Reservoir Engineering Stanford University, February 11-13*, Stanford, California, USA.
- Sandvin, A., Keilegavlen, E., and Nordbotten, J. M. 2013. Auxiliary variables for 3d multiscale simulations in heterogeneous porous media. *J. Comput. Phys.*, 238(0):141–153.
- Sandvin, A., Nordbotten, J., and Aavatsmark, I. 2011. Multiscale mass conservative domain decomposition preconditioners for elliptic problems on irregular grids. *Comput. Geosci.*, 15:587–602.
- Smith, B. F., Bjørstad, P. E., and Gropp, W. D. 1996. *Domain decomposition*. Cambridge University Press, Cambridge. Parallel multilevel methods for elliptic partial differential equations.
- Wang, Y., Hajibeygi, H., and Tchelepi, H. A. 2012. Algebraic multiscale linear solver for heterogeneous elliptic problems. In *ECMOR XIII - 13th European Conference on the Mathematics of Oil Recovery*, Biarritz, France. EAGE.
- Wolfsteiner, C., Lee, S. H., and Tchelepi, H. A. 2006. Well modeling in the multiscale finite volume method for subsurface flow simulation. *Multiscale Model. Simul.*, 5(3):900–917 (electronic).
- Zhou, H. and Tchelepi, H. A. 2008. Operator-based multiscale method for compressible flow. *SPE J.*, 13(2):267–273.
- Zhou, H. and Tchelepi, H. A. 2012. Two-stage algebraic multiscale linear solver for highly heterogeneous reservoir models. *SPE J.*, 17(2):523–539.