Grid Adaption for Upscaling and Multiscale Methods

J.R. Natvig (SINTEF) K.-A. Lie* (SINTEF) S. Krogstad (SINTEF)
Y. Yang (ExxonMobil Upstream Research Company)
X.H. Wu (ExxonMobil Upstream Research Company)

September 13, 2012

Abstract

A Dirichlet–Neumann representation method (DNR) was recently proposed for upscaling and simulating flow in reservoirs. The DNR method expresses coarse fluxes as linear functions of multiple discrete pressure values along the boundary and at the center of each coarse block. The number of coarse fluxes and pressure values at the boundary can be adjusted to improve the accuracy of simulation results, and in particular to resolve important fine-scale details. Improvement over existing approaches is substantial especially for reservoirs that contain high permeability streaks or channels. Recent multiscale methods provide an efficient means to obtain fine-scale fluxes or pressures at the cost of solving a coarsened problem. However, these methods can also be utilized as upscaling methods that are flexible with respect to geometry and topology of the coarsened grid. In this work, we compare the multiscale mixed finite element (MsMFE) method and the DNR approach for accurate upscaling. Both methods can be expressed in mixed form, with local stiffness matrices obtained as inner products of numerically computed basis functions with fine-scale sub-resolution. These basis functions are determined by solving local flow problems with piecewise linear Dirichlet boundary conditions for the DNR method and piecewise constant Neumann conditions for MsMFE. Adding discrete pressure points in the DNR method corresponds to subdividing coarse faces and hence increasing the number of basis functions in the MsMFE method. The methods show similar accuracy for 2D Cartesian cases, but the MsMFE method is more straightforward to formulate in 3D and implement for general grids.

Introduction

Being able to understand and predict flow and transport processes is decisive to enhance the recovery from hydrocarbon reservoirs. Porous rocks are typically highly heterogeneous and exhibit a multiscale behavior in the sense that small-scale flow paths determine the overall displacement of fluids in a reservoir. Describing all pertinent flow processes with a single model is impossible and flow modeling is therefore divided into separate steps according to physical scales: from rock models on the micro scale, via facies models and geological models, to simulation models on the macro scale. Upscaling is inevitable to transfer parameters and effective properties up in the model hierarchy.

Our primary interest herein is the upscaling from geological models to simulation models. To accurately model heterogeneous rock formations, geo-cellular grid models have very complex geometries and topologies and may contain millions of cells. Even for models with a few tens or hundred thousand cells, a typical forward simulation of a subsurface flow system will require hours of computer time. Upscaling is therefore often a necessary step to reduce model sizes and reduce the turnaround time for workflows used to assess different model assumptions, explore parameter space, and quantify the large uncertainty that is usually associated with reservoir characterization. When geological models are coarsened to obtain simulation results faster, the cells in the resulting simulation model can be quite large and may contain localized geological features. The upscaled models are obviously easier to simulate, but are also only approximations of the original model. Errors introduced in the upscaling process may be small when the flow field is relatively smooth, but could be large otherwise, in particular for models that contain channels and high-permeability streaks. A critical technical challenge in upscaling is to be able to coarsen geo-cellular models to reduce simulation times while maintaining a high degree of accuracy for the simulation results.

To this end, several methods have been proposed. In the Dirichlet–Neumann representation method [10], DNR for short, expressions are derived for flow rates as linear functions of the pressure value at the center and multiple discrete pressure values along the of each coarse block. The number of pressure values at the boundary is flexible and may be chosen to provide an adequate representation of pressure profiles and flow distribution throughout a dynamical simulation. In the multiscale mixed finite-element method [2], MsMFE for short, one constructs a set of special basis functions by which the effects of the fine-scale heterogeneity can be incorporated into the discretized coarse-scale flow problem in a way that is consistent with the local fine-scale properties of the differential operators. The basis functions are computed by solving localized flow problems driven by source terms. The DNR and MsMFE methods are similar in the sense that both allow straightforward reconstruction of fine-scale flow solutions and can hence be used as part of a multiscale computational procedure. On the other hand, the methods can also be seen as complementary: whereas the basis functions in the MsMFE method are localized and determined by specifying Neumann boundary conditions on fluxes, the local flow solutions are determined by specifying Dirichlet boundary conditions for the pressure in the DNR method.

The purpose of the paper is two-fold. First, we extend the DNR method from Cartesian to fully unstructured grids. To this end, we borrow ideas from the MsMFE methods, write the DNR method on mixed form, and notice that adding pressure points corresponds to subdividing coarse faces and hence increasing the number of basis functions in the MsMFE methods. Second, we compare the accuracy and robustness of the two methods and investigate how subdivision of interfaces between blocks in the coarsened model can be utilized to reduce the error induced by the artificial boundary conditions used to localize the computation of basis functions and Dirichlet–Neumann maps.

Grid and discretization

We consider the following single-phase flow problem

$$\vec{v} + \mathbf{K}\nabla p = 0, \quad \nabla \cdot \vec{v} = q, \tag{1}$$

for $\vec{x} \in \Omega$. Here, $\mathbf{K}$ is the permeability tensor, $\vec{v}$ is the fluid velocity, p is the fluid pressure, and q represents sources and sinks. In addition, (1) needs to be augmented with boundary conditions on $\partial\Omega$.

To discretize (1), we partition Ω into a set $\{\Omega_i\}$ of N_Ω non-overlapping polyhedral grid cells, each of which have n_i planar polyhedral faces $\{\Gamma_{i_k}\}$. We set $N = \sum_i n_i$. The set of polyhedral cells form a conformal unstructured grid in the sense that each of the N_Γ faces that is not part of the outer boundary is shared by two grid cells. The parameters of (1) are constant on each grid cell. The discrete representation of (1) on a single grid cell can be written on the form

$$\mathbf{v}_i = \mathbf{T}_i(\mathbf{e}p_i - \pi_i), \tag{2}$$

where $\mathbf{v}_i$ denotes the vector of outward fluxes on Ω_i, p_i the pressure at the cell center, $\mathbf{e}$ the $n_i \times 1$ vector of ones, π_i the vector of pressures at face centers, and the matrix $\mathbf{T}_i$ is the *transmissibility matrix* for cell Ω_i. Different choices of $\mathbf{T}_i$ results in different methods like the standard two-point flux approximation, the mixed finite-element method, the related mimetic finite difference schemes, and (in a certain sense) some multipoint flux approximations.

Given all transmissibility matrices $\{\mathbf{T}_i\}$, the global system in mixed-hybrid form becomes

$$\begin{bmatrix} \mathbf{T}^{-1} & \mathbf{C} & \mathbf{D} \\ \mathbf{C}^\mathsf{T} & \mathbf{0} & \mathbf{0} \\ \mathbf{D}^\mathsf{T} & \mathbf{0} & \mathbf{0} \end{bmatrix} \begin{bmatrix} \mathbf{v} \\ -\mathbf{p} \\ \pi \end{bmatrix} = \begin{bmatrix} \mathbf{0} \\ \mathbf{q} \\ \mathbf{0} \end{bmatrix}, \tag{3}$$

where $\mathbf{v}$ is the stacked vector of cell fluxes, $\mathbf{p}$ are the cell pressures, and π is the vector of pressures for each interface. In the system matrix, $\mathbf{T}$ is the $N \times N$ block-diagonal matrix with blocks $\mathbf{T}_i$; $\mathbf{C}$ is the $N \times N_\Omega$ block-diagonal matrix in which each block is the $n_i \times 1$ vector $\mathbf{e}$ of ones; and $\mathbf{D}$ is the $N \times N_\Gamma$ matrix in which each row has a single unit entry that identifies the interface corresponding to the entry in $\mathbf{v}$.

In the following, we will operate on a coarse grid in which each grid block is formed by amalgamating cells from an underlying fine grid. The interface between two coarse blocks may or may not have been subdivided; either way, each coarse face consists of a (connected) set of cell faces from the underlying grid. In other words, we keep the notation and let the coarse grid consist of N_Ω coarse grid blocks $\{\Omega_i\}$ and N_Γ coarse interfaces $\{\Gamma_j\}$, defined so that each block Ω_i is a simply connected set of (fine) grid cells and each coarse interface Γ_j is a connected set of fine-grid faces forming (part of) the interface between two coarse blocks or (part of) the outer boundary of a single coarse block. Each coarse grid block, Ω_i has a boundary that is the union of n_i coarse interfaces.

The Dirichlet–Neumann Method

The Dirichlet–Neumann representation method for computing flows in reservoirs is presented in [10] for Cartesian grids in two spatial dimensions. In this section, we briefly review the method and present a slight reformulation that simplifies the extension to unstructured grids in two and three spatial dimensions.

For a coarse block Ω_i with boundary $\partial\Omega_i$ formed by n_i coarse interfaces, we can compute a discrete Dirichlet–Neumann map $\mathbf{M}_i = (\mathbf{m}_1, \ldots, \mathbf{m}_{n_i}) \in \mathbb{R}^{n_i \times n_i}$ such that

$$\mathbf{v}_i = \mathbf{M}_i \pi_i. \tag{4}$$

Each column in $\mathbf{M}_i$ is a vector of interface fluxes computed from shape functions for flux ψ_{i_k} defined on the fine grid restricted to Ω_i. The n_i shape functions are solutions of

$$\vec{\psi}_{i_k} + \mathbf{K}\nabla\phi_{i_k} = 0, \quad \nabla \cdot \vec{\psi}_{i_k} = 0, \tag{5}$$

in Ω_i with Dirichlet boundary conditions given on $\partial\Omega_i$. We will come back to the specification of boundary conditions later.

Since each column of $\mathbf{M}_i$ in (4) is obtained from a solution of the homogeneous equation (5), $\mathbf{e}^\mathsf{T}\mathbf{M}_i$ is the zero vector and the matrix $\mathbf{M}_i$ has rank $n_i - 1$. To represent flow in coarse blocks containing source terms, (4) needs to be expanded by interface fluxes $\mathbf{m}_i$ from a non-homogeneous shape function. We compute the (approximate) effect of source terms in Ω_i by solving

$$\vec{\psi}_{i_0} + \mathbf{K}\nabla\phi_{i_0} = 0, \quad \nabla\cdot\vec{\psi}_{i_0} = \begin{cases} 1/|\Omega_i|, & \text{if } q = 0 \text{ in } \Omega_i \\ q/\int_{\Omega_i} q, & \text{otherwise,} \end{cases} \tag{6}$$

with homogeneous Dirichlet boundary conditions. The outward-directed fluxes on the interfaces of $\partial\Omega_i$ are then evaluated to form the vector $\mathbf{m}_i$. When source terms are included, we get the expression $\mathbf{v}_i = \mathbf{M}_i\pi_i + \mathbf{m}_i q_i$ for the out-fluxes of block Ω_i, where $q_i = \int_{\Omega_i} q$.

To approximate the solution of (1), we require continuity of the coarse flux across all coarse grid interfaces. Let $\mathbf{v}$ be the $N \times 1$ stacked vector of all outward-directed block interface fluxes $\mathbf{v}_i$, where $N = \sum_{i=1}^{N_\Omega} n_i$. Furthermore, if we collect the Dirichlet–Neumann maps of each coarse block in a $N \times N$ block-diagonal matrix $\mathbf{M}$, the non-homogeneous flux vectors in an $N \times N_\Omega$ block-diagonal matrix $\mathbf{m}$, and let π be the vector of all coarse grid interface pressures, we may write the linear system as

$$\begin{bmatrix} \mathbf{I} & -\mathbf{MD} \\ \mathbf{D}^\mathsf{T} & \mathbf{0} \end{bmatrix} \begin{bmatrix} \mathbf{v} \\ \pi \end{bmatrix} = \begin{bmatrix} \mathbf{mq} \\ \mathbf{0} \end{bmatrix}, \tag{7}$$

where $\mathbf{q}$ is the vector of block source terms. Each row in the $N \times N_\Gamma$ matrix $\mathbf{D}$ has a single non-zero entry equal one such that $\mathbf{D}\pi$ are the interface pressures in a block-wise ordering. This implies that $\mathbf{D}^\mathsf{T}\mathbf{v}$ is the vector of sums of the approximations to each interface flux. For internal interfaces, this sum should be zero to ensure flux continuity on the coarse grid, whereas for outer interfaces the flux should either be zero or equal any boundary conditions imposed on the flux. With a simple manipulation of (7), we get a linear system for the interface pressures

$$\mathbf{D}^\mathsf{T}\mathbf{MD}\pi = -\mathbf{D}^\mathsf{T}\mathbf{mq}. \tag{8}$$

Mixed-Hybrid Formulation. Equation (7) can be written in mixed-hybrid form (3) if we define a suitable interpretation of the coarse block pressure p_i and the transmissibility matrix $\mathbf{T}_i$ so that (2) is fulfilled for each block Ω_i. In particular, by multiplying (2) from left by $\mathbf{e}^\mathsf{T}$, we get $q_i = \mathbf{e}^\mathsf{T}\mathbf{v}_i = \mathbf{e}^\mathsf{T}\mathbf{T}_i(\mathbf{e}p_i - \pi_i)$ and hence

$$p_i = \frac{1}{\mathbf{e}^\mathsf{T}\mathbf{T}_i\mathbf{e}}\left(\mathbf{e}^\mathsf{T}\mathbf{T}_i\pi_i + q_i\right).$$

Substituting this expression back into (2), we obtain

$$\mathbf{v}_i = \underbrace{-\left(\mathbf{T}_i - \frac{1}{\mathbf{e}^\mathsf{T}\mathbf{T}_i\mathbf{e}}\mathbf{T}_i\mathbf{e}\mathbf{e}^\mathsf{T}\mathbf{T}_i\right)}_{\mathbf{M}_i}\pi_i + \underbrace{\frac{1}{\mathbf{e}^\mathsf{T}\mathbf{T}_i\mathbf{e}}\mathbf{T}_i\mathbf{e}}_{\mathbf{m}_i} q_i, \tag{9}$$

where we have indicated which parts must be equal to $\mathbf{M}_i$ and $\mathbf{m}_i$, respectively, for the method to be equivalent to the DNR method as defined above. By substitution and using the relations $\mathbf{M}_i\mathbf{e} = \mathbf{M}_i{}^\mathsf{T}\mathbf{e} = \mathbf{0}$ and $\mathbf{e}^\mathsf{T}\mathbf{m}_i = 1$, it can be verified that

$$\mathbf{T}_i = -\mathbf{M}_i + \frac{1}{\alpha_i}\mathbf{m}_i\mathbf{w}_i{}^\mathsf{T}$$

represents a family of transmissibility matrices equivalent to the DNR method, where $\mathbf{w}_i$ is a weighting vector ($\mathbf{e}^\mathsf{T}\mathbf{w}_i = 1$) and α_i is a scaling parameter. For homogeneous media and blocks with planar faces,

the resulting method will be consistent independently of α. However, in practice, it is important to keep the condition number of $\mathbf{T}_i$ reasonable, and hence α can be used for this purpose; a good choice for α dependens on the choice of units. For a given $\mathbf{w}_i$ and α_i, the relation between the block-pressure and interface pressures is given by $p_i = \mathbf{w}_i{}^{\mathsf{T}}\pi_i + \alpha_i q_i$. Herein, we will use $\mathbf{w}_i = \mathbf{e}/(\mathbf{e}^{\mathsf{T}}\mathbf{e})$ or $\mathbf{w}_i = \mathbf{m}_i$ and $\alpha = 1/[\text{darcy}] \approx 1.0132 \cdot 10^{12}$. Another natural choice for α is by setting it equal to the coarse (average) block-pressure of the source shape function (6). This is justified by setting $\mathbf{v}_i = \mathbf{m}_i$ and $\pi_i = \mathbf{0}$ into (2), to obtain $\mathbf{m}_i = \mathbf{T}_i \mathbf{e} p_i = \frac{1}{\alpha_i}\mathbf{m}_i \mathbf{w}_i{}^{\mathsf{T}}\mathbf{e} p_i = \frac{p_i}{\alpha_i}\mathbf{m}_i$.

Boundary conditions in 2D. In flow-based upscaling, the choice of boundary conditions used to construct local solutions of (1) can have a significant impact on the accuracy of the method. To construct a discrete Dirichlet–Neumann map for a grid block Ω_i with n_i coarse interfaces, n_i linearly independent functions $g_k : \partial\Omega_i \to \mathbb{R}$ are needed that will be used as Dirichlet conditions for (5). The DNR method for upscaling [10] has been shown to produce quite accurate results when each g_k is a piecewise linear function in the curve length s measured along $\partial\Omega_i$ (from some reference) with joints at n_i points on $\partial\Omega_i$. These points are referred to as *pressure points*. The number and position of pressure points (and coarse interfaces) may be used to improve the accuracy of the Dirichlet–Neumann representation for heterogeneous models.

In [10], the authors specify pressure points along $\partial\Omega_i$ according to a rule that takes into account permeability. A partition of $\partial\Omega_i$ into a set of (non-overlapping) coarse-grid interfaces $\{\Gamma_j\}$ is then defined such that there is one pressure point on each interface Γ_j. Some technical conditions ensure that this partition is uniquely defined. If we assume that the curve length s takes the value s_j in the pressure point of Γ_j (i.e., the centroid of Γ_j) we require that $g_k(s_j) = \delta_{jk}$. This ensures that the basis functions $\{g_k\}$ forms a partition of unity on $\partial\Omega_i$, which is necessary for the Dirichlet–Neumann method to give zero (coarse and fine) flux for constant-pressure solutions. Furthermore, by choosing the pressure in each pressure point as degrees-of-freedom for the coarse-grid interface pressure, the matrix $\mathbf{M}_i$ is easily constructed from the interface fluxes of each shape function.

Herein, we have chosen a different approach: we take the partition of $\partial\Omega_i$ as given (say, from some algorithm that takes permeability into account) and choose the centroid of a fine-grid face $\Gamma_j^f \subset \Gamma_j$ in the middle of the coarse interface as the pressure point of interface Γ_j. This enables the use of grid amalgamation techniques (see e.g., [4]) to define the coarse interfaces in the coarse grid, which is straightforward to extend to 3D.

Boundary Conditions in 3D. To make a direct extension to 3D of the boundary conditions used for shape functions in 2D, we would need to construct piecewise linear functions in some parametrization of the surface $\partial\Omega_i$ of each coarse grid block Ω_i. This is a fairly complicated process and would require a triangulation of $\partial\Omega_i$ in parameter space. We have therefore attempted a different approach.

As in 2D, to construct the Dirichlet–Neumann representation for a grid block Ω_i whose outer surface $\partial\Omega_i$ is the union of n_i coarse-grid interfaces, we need n_i functions whose values will be used as Dirichlet boundary conditions in (5). The boundary conditions for (5) need only be given as piecewise constant on the fine-grid faces $\Gamma_j \subset \partial\Omega_i, j \in \mathcal{N} = \{1, \cdots, N\}$. We will therefore define n functions $\hat{g}_i$ taking the (local) face number as argument, $\hat{g}_i : \mathcal{N} \to \mathbb{R}$, that are roughly hat-shaped with local support.

To this end, we need a set of n pressure points, one for each interface $\Gamma_j \subset \partial\Omega_i$, that will position the apex of the corresponding $\hat{g}_i$. We assume that each pressure point is the centroid of fine-grid faces $f_i \in \mathcal{F} = \{f_1, \cdots, f_n\}$. The functions $\hat{g}_i$ should fulfill

$$\sum_i \hat{g}_i(j) = 1 \quad \text{for } j \in \{1, \cdots, N\} \quad \text{(partition of unity)}$$

$$\hat{g}_i(j) = \delta_{ij} \text{ for } j \in \{f_1, \cdots, f_n\} \quad \text{(local support)}.$$

In addition, we need to introduce a distance function $\text{dist}(i, j)$ between any two fine-grid faces in $\partial\Omega_i$.

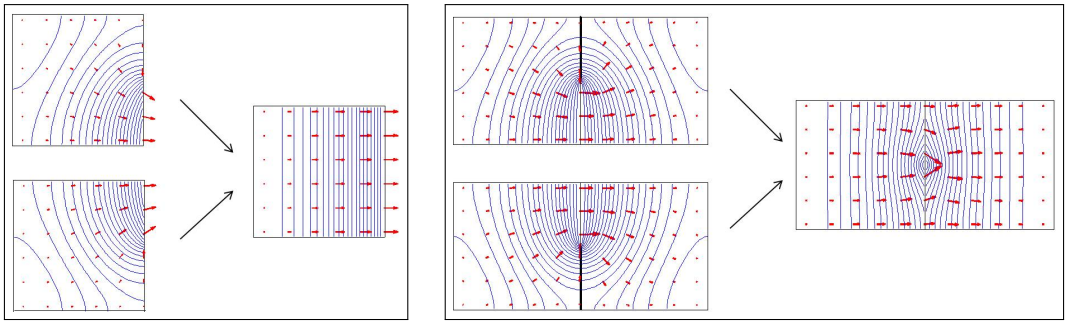

Figure 1 *The left plot shows two one-block MsMFE basis functions for a subdivided interface as well as the sum of the two basis functions reproducing* half *of a Raviart–Tomas basis function, and hence linear flux. The left plot shows two two-block MsMFE basis functions for a subdivided interface as well as the sum of the two that fails to reproduce linear flux.*

Then the boundary conditions are given by

$$\hat{g}_i(j) = \frac{\max(r_i - \mathrm{dist}(f_i, j), 0)}{\sum_k \max(r_k - \mathrm{dist}(f_k, j), 0)}, \tag{10}$$

where $r_i = \min\{\mathrm{dist}(f_i, f_k) | f_k \neq f_i, f_k \in \mathscr{F}\}$. Herein, we have defined the distance function to be the following graph distance: First, we construction the (fine-grid) surface grid over $\partial\Omega_i$ and build a graph $G(V,E)$ in which each face $f \in \partial\Omega_i$ is represented by a vertex in V and each edge between faces is represented by an edge in E. With abuse of notation, the distance $\mathrm{dist}(i, j)$ is then defined as the shortest graph distance between corresponding vertices in $i, j \in V$.

The Multiscale Mixed Finite-Element Method

The main idea of the multiscale mixed finite-element method [2] is to construct a special approximation space defined over the coarse grid in which each basis function represents the flow between two neighboring grid blocks and solves a local flow problem of the form (6), but with different boundary conditions used for localization, as will be explained below. Unlike the DNR method, the MsMFE method is designed to give conservative fluxes directly on the coarse and fine scale, as well as on any intermediate scale, by using the fine-scale resolution of the basis functions.

One-block basis functions. A MsMFE basis function represents the flow over an interface or sub-interface Γ_k between two neighboring blocks. Restricted to one block Ω_i, the basis function $\vec{\psi}_{i_k}$ is defined as the solution of

$$\vec{\psi}_{i_k} + \mathbf{K}\nabla\phi_{i_k} = 0, \quad \nabla\cdot\vec{\psi}_{i_k} = \omega_i, \tag{11}$$

in Ω_i with Neumann boundary conditions $\vec{\psi}_{i_k}\cdot\vec{n} = \nu_k$ on Γ_k and $\vec{\psi}_{i_k}\cdot\vec{n} = 0$ on $\partial\Omega_i \setminus \Gamma_k$ with $\int_{\Omega_i}\omega_i = 1$ and $\int_{\Gamma_k}\nu_k = 1$. If *global* information is available as a previously computed flux field $\vec{v}_0$, this is incorporated in ν_k by setting $\nu_k = \vec{v}_0\cdot\vec{n}_k / \int_{\Gamma_k}\vec{v}_0\cdot\vec{n}_k\,ds$. If no global information is available, we set $\nu_k = 1/\Gamma_k$. If more than one flux degree-of-freedom per interface is desired, this can be obtained by subdividing and assigning one basis function to each sub-interface. The computation of each basis remains the same; simply setting non-zero flux conditions on the sub-interface and zero flux conditions elsewhere. To the left in Figure 1, we have depicted two basis functions corresponding to a subdivided interface (with constant non-zero flux conditions) for a homogeneous domain. Taking the average of these functions we obtain the basis function for the original non-divided interface.

Two-block basis functions. In the two-block approach, the basis function for an interface Γ_k between blocks Ω_i and Ω_j is defined as the solution of

$$\vec{\psi}_k + \mathbf{K}\nabla\phi_k = 0, \quad \nabla\cdot\vec{\psi}_k = \begin{cases} \omega_i, & \text{for } \vec{x}\in\Omega_i, \\ -\omega_j, & \text{for } \vec{x}\in\Omega_j \end{cases} \tag{12}$$

in $\Omega_i \cup \Omega_j$ with zero Neumann boundary on $\partial(\Omega_i \cup \Omega_j)$. Solving the two-block problem for a basis function corresponding to a sub-interface means that the remaining part of the interface ($\partial\Omega_i \cap \partial\Omega_j \setminus \Gamma_k$) is considered as boundary and hence is shut through a zero Neumann condition. To the right in Figure 1, we have depicted two basis functions corresponding to a subdivided interface for a homogeneous domain. Because of the *inner boundary*, the flux field experiences a peak (singularity) at the inner corner. This means that the normal component of the flux over the open boundary is varying strongly and non-linearly, and that the average of the two basis functions differs from the basis function defined for the original non-divided interface.

The multiscale method. To approximate the solution of (1), we introduce the multiscale expansion

$$\mathbf{v}_{fs} = \Psi\mathbf{v} + \tilde{\mathbf{u}}_{fs}, \qquad \mathbf{p}_{fs} = \mathbf{I}\mathbf{p} + \tilde{\mathbf{p}}_{fs}, \qquad \pi_{fs} = \mathbf{J}\pi + \tilde{\pi}_{fs}. \tag{13}$$

Here, $\mathbf{v}$, $\mathbf{p}$, and π denote unknowns associated with the coarse grid $\{\Omega_i\}$, whereas $\mathbf{v}_{fs}$, $\mathbf{p}_{fs}$, and π_{fs} denote the same quantities reconstructed on the fine grid. The matrix Ψ represents the fine-scale reconstruction operator for fluxes and contains the flux basis functions $\vec{\psi}_{i_k}$. The matrices $\mathbf{I}$ and $\mathbf{J}$ are simple prolongation operators from coarse blocks and coarse faces to fine-grid cells and faces, respectively. Finally, $\tilde{\mathbf{u}}_{fs}$, $\tilde{\mathbf{p}}_{fs}$, and $\tilde{\pi}_{fs}$ denote reminder terms defined on the fine grid. To form a global system on the coarse grid $\{\Omega_i\}$, we insert (13) into (3), multiply by the compression operator $\mathrm{diag}(\Psi^{\mathsf{T}}, \mathbf{I}^{\mathsf{T}}, \mathbf{J}^{\mathsf{T}})$ from the left, and drop all remainder terms

$$\begin{bmatrix} \Psi^{\mathsf{T}}\mathbf{T}^{-1}\Psi & \Psi^{\mathsf{T}}\mathbf{C}\mathbf{I} & \Psi^{\mathsf{T}}\mathbf{D}\mathbf{J} \\ \mathbf{I}^{\mathsf{T}}\mathbf{C}^{\mathsf{T}}\Psi & \mathbf{0} & \mathbf{0} \\ \mathbf{J}^{\mathsf{T}}\mathbf{D}^{\mathsf{T}}\Psi & \mathbf{0} & \mathbf{0} \end{bmatrix} \begin{bmatrix} \mathbf{v} \\ -\mathbf{p} \\ \pi \end{bmatrix} = \begin{bmatrix} \mathbf{0} \\ \mathbf{I}^{\mathsf{T}}\mathbf{q} \\ \mathbf{0} \end{bmatrix}. \tag{14}$$

MsMFE as an upscaling method. To use MsMFE as an upscaling method, all we need to do is to compute inverse transmissibility matrices $\mathbf{T}_i^{-1}$ that consist of inner-products of the numerically computed basis functions $\vec{\psi}_{i_k}$ on the corresponding coarse blocks Ω_i. As above, the inverse transmissibility matrix $\mathbf{T}_i^{-1}$ describes the relation between the out-fluxes over the coarse interfaces of Ω_i and the block and interface pressures through (2). Given basis functions $\vec{\psi}_{i_k}$, the (k,ℓ) entry of matrix $\mathbf{T}_i^{-1}$ is given by

$$\mathbf{T}_i^{-1}(k,\ell) = \int_{\Omega_i} \vec{\psi}_{i_k} \cdot \mathbf{K}^{-1} \vec{\psi}_{i_\ell}.$$

The elements of $\mathbf{T}_i^{-1}$ can equivalently be derived from the computed pressure at the boundary $\partial\Omega_i$. First, observe that the pressure ϕ_{i_k} in (11) is only defined up to a constant and we may therefore add the condition $\int_{\Omega_i} \omega_i \phi_{i_k} = 0$ to close the equation. Using Gauss–Green's formula, we derive

$$\begin{aligned} \mathbf{T}_i^{-1}(k,\ell) &= \int_{\Omega_i} \vec{\psi}_{i_k} \cdot \mathbf{K}^{-1} \vec{\psi}_{i_\ell} = -\int_{\Omega_i} \vec{\psi}_{i_k} \cdot \nabla\phi_{i_\ell} = -\int_{\Omega_i} \nabla\cdot\vec{\psi}_{i_k}\phi_{i_\ell} + \int_{\partial\Omega_i} \phi_{i_\ell}\vec{\psi}_{i_k}\cdot\vec{n} \\ &= -\int_{\Omega_i} \omega_i\phi_{i_\ell} + \int_{\Gamma_k} \phi_{i_\ell} v_k = \int_{\Gamma_k} \phi_{i_\ell} v_k. \end{aligned} \tag{15}$$

This means that $\mathbf{T}_i^{-1}(k,\ell)$ is the (v_k-weighted) average of the pressure basis ϕ_{i_ℓ} on interface k. In particular, $\mathbf{T}_i^{-1} = [\phi_{i,1}, \ldots, \phi_{i,n_i}]$, where $\phi_{i,k}$ is the vector of average coarse interface pressures in the numerical solution of (11).

Numerical experiments

In this section, we report the results of several numerical experiments that were conducted to verify, validate, and compare the three numerical methods presented above (DNR, 1-block MsMFE, and 2-block MsMFE). First, we discuss two basic properties, reproduction of linear flow and robustness with respect to aspect ratios. Second, we use the SPE10 data set to investigate the robustness of the methods and how subdivision of coarse-block interfaces affect the accuracy of the methods. Third, we discuss how the accuracy of the methods can be improved by careful adaption of the coarse grid for high-contrast

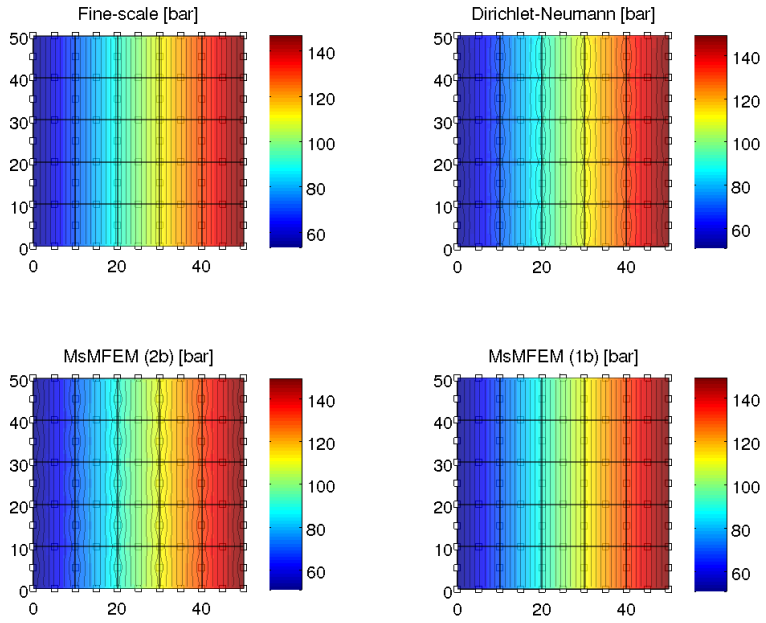

Boundary	Reference	DNR	2-block	1-block
East	9.869233e-07	9.689592e-07	9.404560e-07	9.869233e-07
West	-9.869233e-07	-9.689592e-07	-9.404560e-07	-9.869233e-07
North	-1.816886e-19	-1.475227e-19	2.309273e-15	-2.927903e-20
South	-1.079967e-19	-1.278490e-19	2.309294e-15	1.943271e-21

Figure 2 *Pressures and fluxes across the outer boundary for the linear flow example on a* 5×5 *coarse grid in which all coarse faces have been subdivided into two segments of equal length.*

media that contain barriers, high-flow channels, or combinations thereof. Finally, we show an example of application to a real-life reservoir modeled using a complex, 3D unstructured grid. All experiments are conducted using the MATLAB Reservoir Simulation Toolbox [7], Version 2011b.

Reproduction of linear flow is often used as a measurement for discretization methods. The DNR method does not automatically reproduce linear flow unless the coarse-grid faces are subdivided in certain ways. Likewise, the two-block version of the MsMFE method will not reproduce linear flow correctly if the coarse-grid interfaces are subdivided, as was illustrated in Figure 1. On the other hand, the 1-block version of MsMFE and the 2-block version with no face refinement will correctly reproduce linear flow.

Example 1 (Linear flow) *We consider a 2D, horizontal reservoir that covers an area of* 50×50 m^2 *represented on a square grid with one meter resolution in each direction. Linear flow is specified by imposing a pressure drop of 150 to 50 bar from the west to the east boundary and no-flow boundary conditions on the north and south boundaries. Figure 2 shows the pressure for a homogeneous, isotropic permeability field of 100 mD computed by the DNR and the two MsMFE methods on a* 5×5 *coarse grid in which all coarse faces have been subdivided into two equal segments. The figure also reports the computed fluxes out of the four boundaries, for which the relative errors are 1.82% (DNR), 4.71% (2-block MsMFE), and 0 % (1-block MsMFE). It is obvious that both the DNR and the 2-block method fail to reproduce linear flow. For the DNR method, however, there is a simple fix [10]: it can be shown analytically that adding an extra segment with length equal one cell at the corners of each coarse block will reproduce linear flow exactly. For short, we will refer to the method as DNR-c in the following.*

To investigate the effect of face refinement, we also subdivided the coarse faces into five and ten segments; errors in the flux for the resulting computations are reported in Table 1. For comparison, the table also reports the corresponding errors for a model with lognormal permeability for which neither of the schemes will reproduce the correct flow unless refined to the fine grid. As expected, dividing the

Table 1 *Errors in coarse-scale fluxes measured in the relative L^1 norm for linear flow in a 2D reservoir with homogeneous or lognormal permeability represented on a 50×50 grid. The approximate solutions are computed on a coarse 5×5 grid in which each face is subdivided into n uniform segments, except for the last row for which the faces are divided nonuniformly into three segments with a segment one cell wide at each corner.*

		Homogeneous			Lognormal		
Segments	bases	DNR	2-block	1-block	DNR	2-block	1-block
1	60	1.04e-01	1.50e-13	1.62e-13	9.68e-02	1.59e-02	4.66e-02
2	120	2.47e-02	9.01e-02	1.40e-13	2.51e-02	9.93e-02	5.38e-02
5	300	8.97e-02	3.16e-02	1.67e-13	8.87e-02	3.47e-02	2.28e-02
10	600	1.05e-13	6.25e-13	2.05e-13	6.51e-14	1.12e-12	3.41e-13
corner	180	7.96e-14	4.77e-02	1.60e-13	8.83e-03	5.72e-02	5.31e-02

faces into two or three segments gives a significant increase in the error for the 2-block MsMFE method. On the other hand, although this method is not a consistent discretization, it produces (by far) the lowest error when using no face refinement for the lognormal permeability. For the DNR method, we observe that adding extra segments at the corners of the coarse grid reduces the error by one order of magnitude for the cost of three times as many basis functions. Finally, the table verifies that all methods converge as the coarse faces are fully refined using ten segments.

Real-life reservoir models often have high aspect ratios. In the next example, we will therefore investigate how the accuracy of the DNR, DNR-c and the MsMFE methods is affected by increasing the aspect ratio of the cells in the underlying fine grid.

Example 2 (Aspect ratio) *We consider a $100 \times (100/L)$ m^2 homogeneous reservoir subject to two different types of boundary conditions: linear flow making an angle θ with the x-axis or Dirichlet boundary conditions with pressure $p(x,y) = 1000\sin(\pi x/L)$ on the north and east faces and zero pressure otherwise. Table 2 reports errors for the DNR, DNR-c, and the two MsMFE methods for aspect ratios $L = 1,2,3,4,6,12$. The DNR method has significantly larger errors than the other three methods and the error increases with increasing aspect ratio. The error also increases for the MsMFE methods for the analytic case, but this can mainly be attributed to inaccurate representation of the Dirichlet boundary conditions; if subdivide all boundary faces, the error does not increase with increasing aspect ratio. Overall, the DNR-c method is the most accurate, but also has three times as many basis functions as the MsMFE methods.*

Although not widely reported in the literature, the multiscale finite-volume method [5] is also known to loose accuracy for increasing aspect ratios. We believe that the explanation is the following in the case of Cartesian grids: In the MsMFE methods, the basis functions are constructed by considering flow in each

Table 2 *Errors in coarse-scale fluxes measured in the relative L^1 norm for a 2D reservoir of dimensions $100 \times (100/L)$ with flow driven by two different boundary conditions. Approximate solutions are computed on a 12×12 coarse grid overlying a fine 60×60 grid.*

	Linear, $\theta = \pi/3$				Analytic			
Ratio	DNR	DNR-c	2-block	1-block	DNR	DNR-c	2-block	1-block
1:1	8.46e-02	3.00e-13	1.63e-09	2.94e-13	8.71e-02	3.11e-03	4.95e-03	4.95e-03
2:1	2.53e-01	2.02e-13	1.32e-09	2.03e-13	2.30e-01	3.47e-03	7.71e-03	7.71e-03
3:1	3.29e-01	3.18e-13	1.10e-09	3.12e-13	3.05e-01	4.39e-03	1.03e-02	1.03e-02
4:1	3.66e-01	1.82e-13	9.15e-10	1.76e-13	3.45e-01	5.46e-03	1.31e-02	1.31e-02
6:1	4.02e-01	3.50e-13	6.70e-10	3.39e-13	3.90e-01	8.06e-03	1.79e-02	1.79e-02
12:1	4.40e-01	2.06e-13	3.60e-10	1.87e-13	4.41e-01	1.21e-02	2.60e-02	2.60e-02

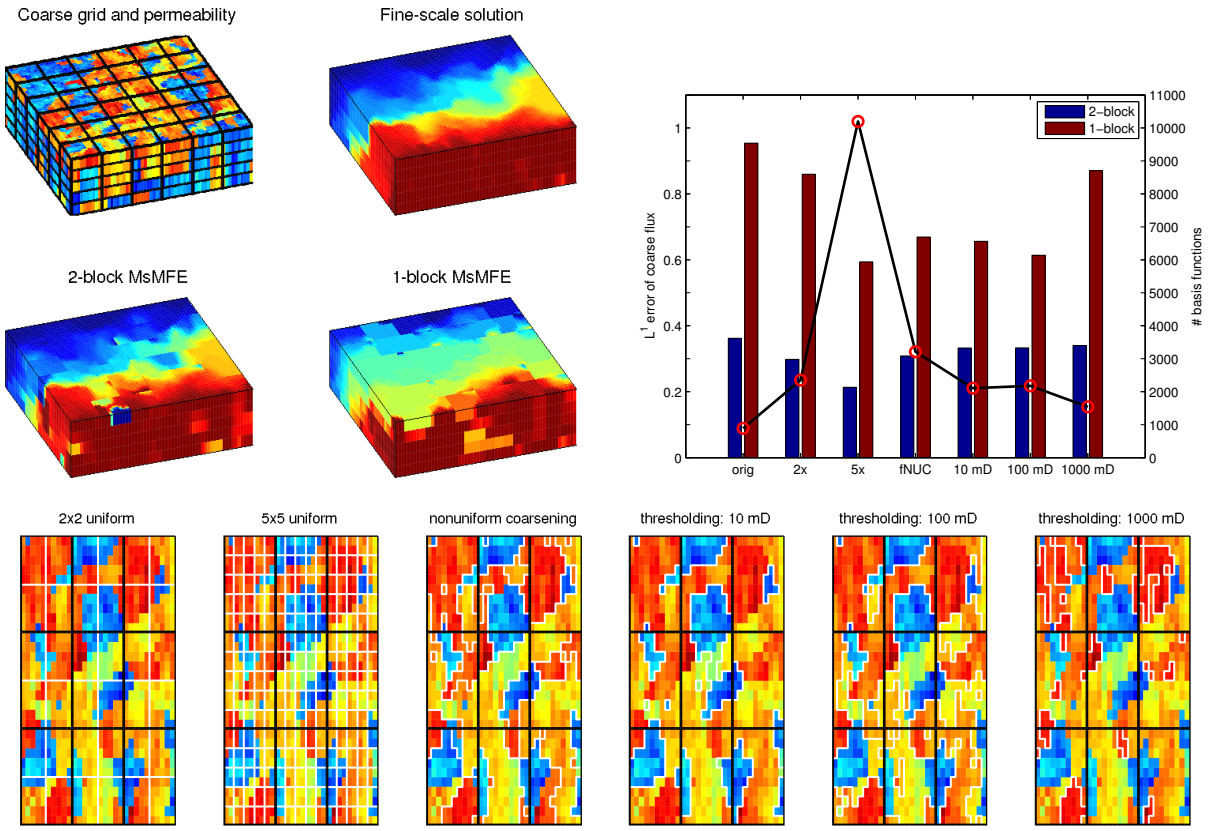

Figure 3 *A* $[1:60]\times[1:60]\times[61:65]$ *cut-out from the SPE10 model. The upper left plots show approximate pressure solutions on a base case with a* $6\times6\times5$ *coarse grid. The upper right plot shows the number of coarse blocks involved in the computation of basis functions and* L^1 *errors of the coarse-scale fluxes for the face refinements shown in the lower plots:* 2×2 *and* 5×5 *uniform subdivision, nonuniform coarsening using permeability values [4], and permeability thresholding with values* 10, 100 *and* 1000 *mD.*

axial direction, which is represented by a separate basis function. In the MsFV and the DNR methods, on the other hand, basis functions are constructed by imposing a unit pressure at one coarse-grid vertex and zero pressures at the other vertices. Each basis function will therefore introduce a certain coupling of the flow in the axial directions, and this effect *may* cause the error to increase significantly with increasing aspect ratios. Introducing extra degrees of freedom at the corners (in the form of segments that are one cell width wide) breaks the coupling and makes the DNR method robust with respect to aspect ratios.

Next, we study several cut-outs from Model 2 of the 10th SPE Comparative Solution Project [3] that was designed to benchmark various upscaling methods and has later become one of the most popular data sets to use when validating multiscale methods. In the literature, the most common test is to use horizontal slices in which the permeability is isotropic. We have therefore conducted a systematic study in which we have run all slices orthogonal to each of the axial directions using linear and analytic boundary conditions as in the previous example. In all runs, we used two different types of face refinement: uniform refinement of each coarse face into n segments or nonuniform refinement of each coarse face based upon a two-bin segmentation of permeability values for a given threshold value. The only conclusion we can draw from these experiments is that both the DNR and the 2-block MsMFE method are able to exploit channelized or patchy patterns in the permeability in a relatively robust manner to produce accurate flow field. Apart from that, the results are somewhat inconclusive. In some cases, the DNR method is more accurate than the 2-block MsMFE method and vice versa. Likewise, using nonuniform refinement of the block interfaces will give significant improvements in some cases and not in others.

We have also run several 3D subsets of the full model using the two MsMFE methods, e.g., as shown in Figure 3. With no face refinement, the 2-block method reproduces most of the qualitative features of the fine-scale solutions although it sometimes exhibits certain numerical artifacts. The 1-block method, on the other hand, will in many cases fail to resolve the fine-scale solution qualitatively. Introducing various forms of face refinement improves the resolution of both methods but the improvements may not be large enough to justify the increased computational cost induced by the extra basis functions.

In all simulations reported so far, the face refinement was imposed *a priori*, i.e., without using any information about the flow pattern at all. In practice, it may be possible to adapt the coarse grids by identifying regions or individual cells that are responsible for non-monotone and other erroneous behavior. In some cases it may be possible to use a combination of experience and geological knowledge to introduce *a priori* adaption in targeted regions. In other cases, one may have one or two global fine-scale solutions available for diagnostics. Herein, we use a much simpler approach: knowing that the analytical solution should be monotone, we can refine coarse faces and recompute the corresponding basis functions in regions where the approximate solution falls outside the minimum and maximum values from the prescribed boundary condition. The refinement can be introduced in several ways, e.g., as shown in the next example.

Example 3 (SPE10: Adaption) *To demonstrate the effect of* a posteriori *face adaption, we consider sixty x-slices with* 120×85 *cells and 220 y-slices with* 60×85 *cells extracted from the SPE10 data. Flow is driven by linear boundary conditions with* $p = 150$ *bar on the west and* $p = 50$ *bar on the east boundary or by analytic boundary conditions with* $p(x,y) = 1000\sin(\pi x/L)$ *on the north and east boundaries and zero pressure otherwise. As base cases, we will use a* 6×5 *coarse grid with 2D analogues of the first two and the last three face refinements shown in Figure 3: coarse faces divided uniformly into one, two, or five uniform segments and or divided nonuniformly based upon permeability thresholding with values* 10, 100, *and* 1000 *mD. We then pick all slices in which the relative* L^1 *error exceeded unity and recompute the approximate solution with three different strategies for refining faces: In the first approach, we subdivide the coarse faces uniformly in all blocks in which the pressure computed from the coarse-scale system falls outside the prescribed pressure range. For the base grids with uniform segments, we subdivide faces into ten segments, while for the base grids with nonuniform segments we subdivide coarse faces into five segments. In the second approach, we subdivide the coarse faces uniformly, but only in those cells (along the face) in which the reconstructed fine-scale pressure falls outside the prescribed pressure range. In the third approach, we subdivide the coarse faces uniformly in all blocks in which a certain percentage of the reconstructed fine-scale pressures fall outside the prescribed pressure range.*

Figure 4 shows the corresponding reduction in L^1 *error and increase in the number of coarse blocks involved in the computation of basis functions for y-slices with analytic boundary conditions for three of the different base cases. For the* 6×5 *base case, the only approach that seems to work is the third approach using a zero percentage, i.e., refining in all cells that fall outside the pressure range. For the other two base cases, all our three approaches for* a posteriori *adaption give significant improvements in the error for most of the slices at the cost of less than a doubling of the number of blocks involved in the computation of basis functions. However, it is worth noting that improved solutions are not obtained in all cases even if the number of basis functions is increased. The adaptive methods can for this reason not be considered as fully robust.*

To investigate the efficiency of a posteriori *adaption, we include an experiment in which we compare the various* a priori *face refinements for a* 6×5 *coarse grid with similar* a posteriori *face refinements for a* 3×3 *coarse grid. As expected, the scatter plots in Figure 5 show that adding degrees-of-freedom* a posteriori *to a relatively coarse basis grid gives significant improvements of the error (and hence the robustness of the method).*

Similar a posteriori *adaption can of course be introduced for the MsMFE methods as well. Table 3*

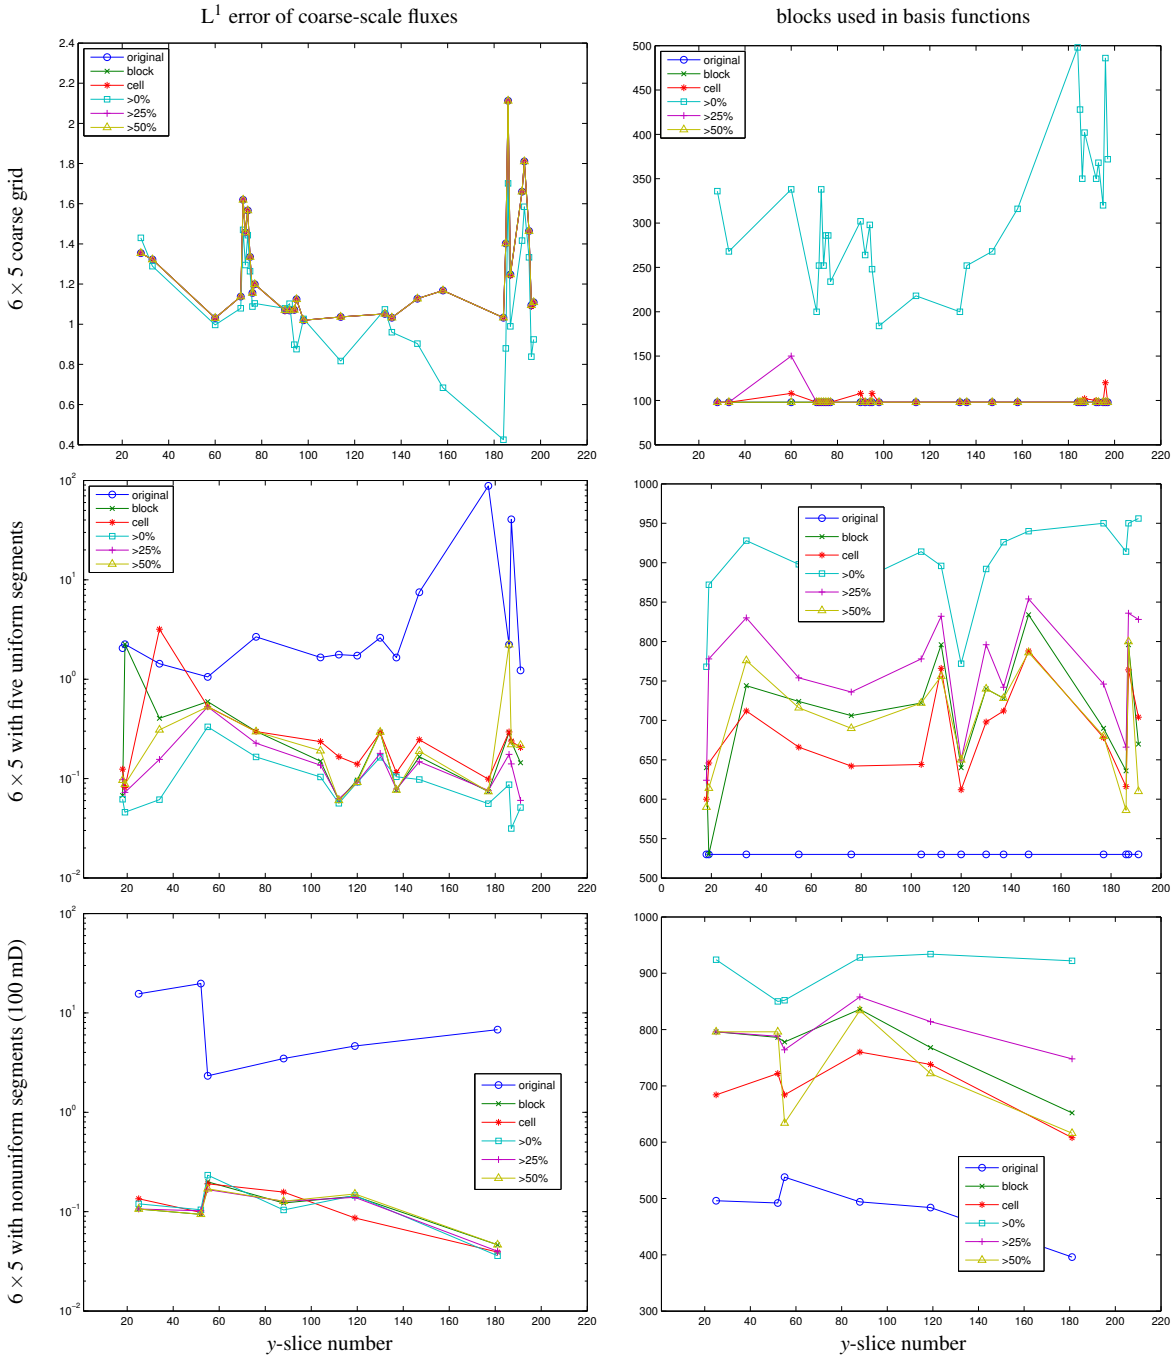

Figure 4 *Use of* a posteriori *face adaption for the DNR method on y-slices of the SPE10 data set. The plots to the left show L^1 errors of coarse fluxes, the plots to the right show the number of coarse blocks used to compute basis functions. Each row corresponds to a different base grid that was used before introducing the* a posteriori *adaption.*

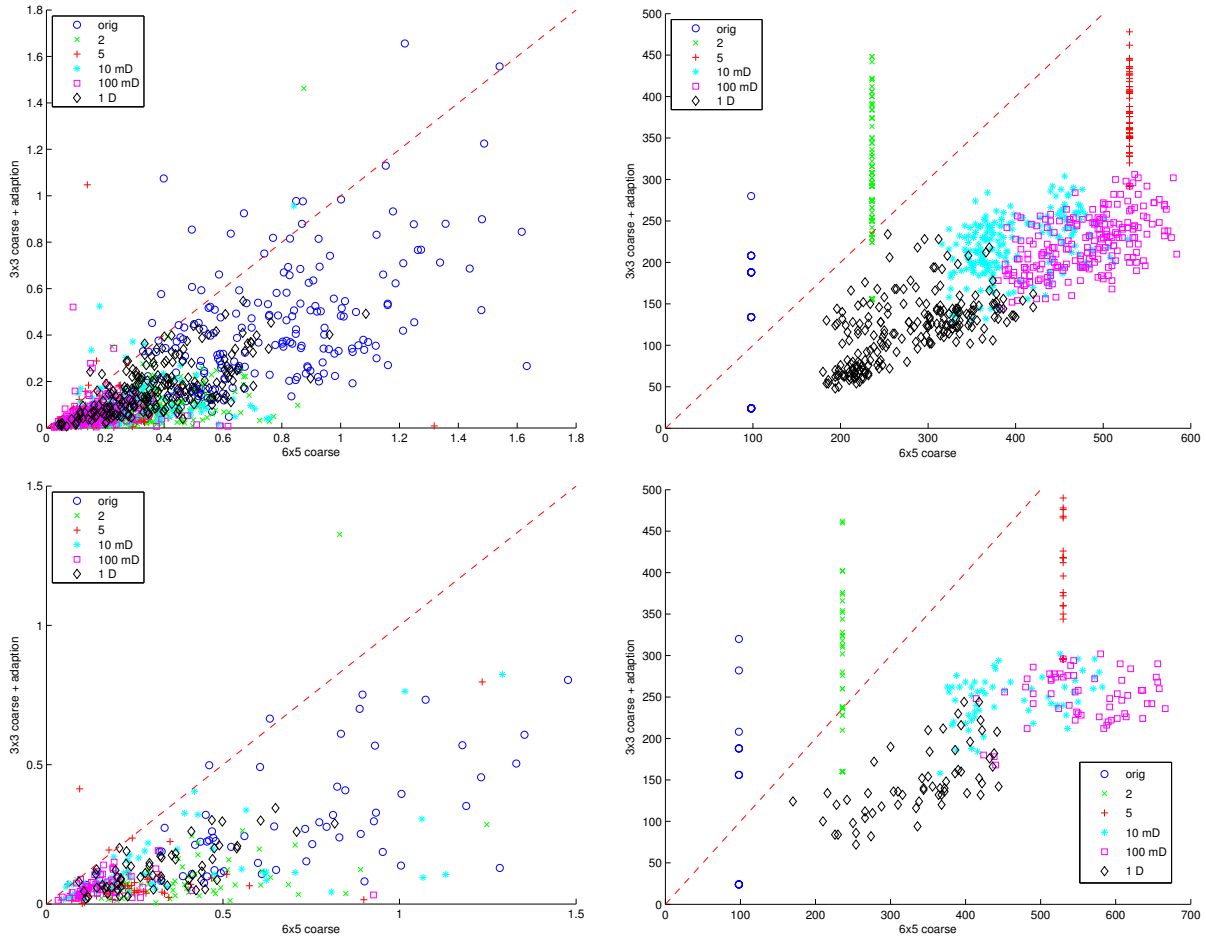

Figure 5 *The left column shows scatter plot of the L*1 *error on a set of* 3×3 *coarse grids with* a posteriori *face refinement versus the* L^1 *error on a* 6×5 *coarse grid with* a priori *face refinement for y-slices of SPE10 with analytic boundary conditions (upper row) and x-slices with linear boundary conditions. The plots in the right column show the corresponding number of coarse blocks used in the calculation of basis functions.*

Table 3 *Ratios in the relative* L^1 *error and in the number of coarse blocks involved in computing basis functions for simulations with and without* a posteriori *adaption in the 2-block MsMFE methods.*

Model	Boundary	base case		2×2		5×5		10 mD		100 mD		1 D	
	condition	err	bf	err	bf	err	bf	err	bf	err	bf	err	bf
x-slice	linear	0.34	9.12	0.49	3.36	0.73	1.45	0.85	1.28	1.02	1.17	0.92	1.69
x-slice	analytic	0.36	8.71	0.44	3.20	0.60	1.47	0.86	1.25	1.08	1.13	0.94	1.60
y-slice	linear	0.42	8.89	0.62	3.15	0.83	1.38	0.93	1.35	1.04	1.26	0.85	1.94
y-slice	analytic	0.31	8.61	0.48	3.01	0.79	1.39	0.92	1.28	1.03	1.20	0.82	1.86

reports the average reduction in error and corresponding increase in the number of basis functions for the 2-block MsMFE method when we refine the faces of all blocks that have at least one fine-scale pressure value outside the prescribed range. Results for the 1-block method are similar and hence not reported.

Extensive numerical experiments have shown that although multiscale methods may provide high accuracy for media with very strong permeability contrasts (high-permeable streaks or barriers) [1, 9], one can also easily construct cases where the methods fail to provide accurate solutions [1, 6]. Previous experience has show that adapting the coarse grid to barriers in the permeability field may give improved resolution [1]. Whereas it is easy to implement support for arbitrarily shaped coarse grids in the MsMFE methods, we have not yet found a good method to do so for the DNR method in 3D. In the following, we will therefore only use the two MsMFE methods when considering 3D examples of high-contrast media.

Example 4 (Multiple streaks) *We consider a rectangular domain that contains several high or low-permeable streaks described on a uniform* $50 \times 50 \times 10$ *fine grid. Flow is driven by Dirichlet boundary conditions equal 300 bar on the south side and 100 bar on the north side. To systematically study the behavior of the two MsMFE methods, we generated three different permeability masks and run cases with low-permeable barriers in a high-permeable background, or high-permeable streaks in a low-permeable background. In both cases, we study the effect of increasing the contrast between the background and the foreground permeability. In the barrier cases, the background permeability is held fixed at* $K_b = 1$ *D while the foreground permeability of the barriers varies seven orders of magnitude,* $K_f = 10^{-3}, \ldots, 10^3$ *mD. In the high-flow channel cases, the channel permeability is held fixed at* $K_f = 1$ *D, while the background permeability varies seven orders of magnitude,* $K_f = 10^{-3}, \ldots, 10^3$ *mD. Figure 6 reports the observed errors in coarse-scale fluxes for the 2-block method. Here, we can make several interesting observations.*

Starting with the barrier cases, we observe very large errors for the base case with a $5 \times 5 \times 1$ *coarse grid when the media contrast spans five to seven orders of magnitude, in particular for the third mask. Refining the coarse-grid faces reduces the error somewhat, and the reduction in error is almost the same regardless of what type of face refinement is introduced. In all cases, the lowest number of basis functions is obtained for the nonuniform face refinement, and this should therefore be our method of choice. The biggest reduction of the error, however, is obtained when volumetric adaption to the barriers is introduced for cases with strong media contrasts. On the other hand, as the media contrast decreases, the error of the volumetric adaption increases because of numerical errors introduced by not using rectangular coarse blocks. This is particularly evident for parameter* $K_f = 1000$ *that corresponds to a homogeneous medium, for which we clearly observe strong numerical artifacts also for grids having more than a single basis function associated with each coarse interface. One possible source of this error is that the 2-block does not reproduce uniform flow. The 1-block method, on the other hand, reproduces uniform flow but has errors that are equal or larger for all media contrasts except for the homogeneous case. This indicates that boundary conditions and geometrical complexity of the coarse blocks may be more important sources of errors than the inability to produce uniform flow. The increase in error is also most pronounced for the third mask, for which the streaks have the most complicated geometry.*

Moving on to the case with high-permeable streaks, we first of all observe that the volumetric subdivision is clearly unsuitable. This is particularly evident for the second case in which the cells inside the streaks in many places are only connected through edges (and not faces). This will result in a large number of coarse blocks and an overall set of basis functions that does not produces the almost linear flow in a good way. Secondly, it does not seem to matter much what type of face refinement one chooses for the two first masks. For the third mask, subdividing faces according to foreground and background permeability (denoted 'fNUC' in the figure) gives surprising small improvement in the error and further

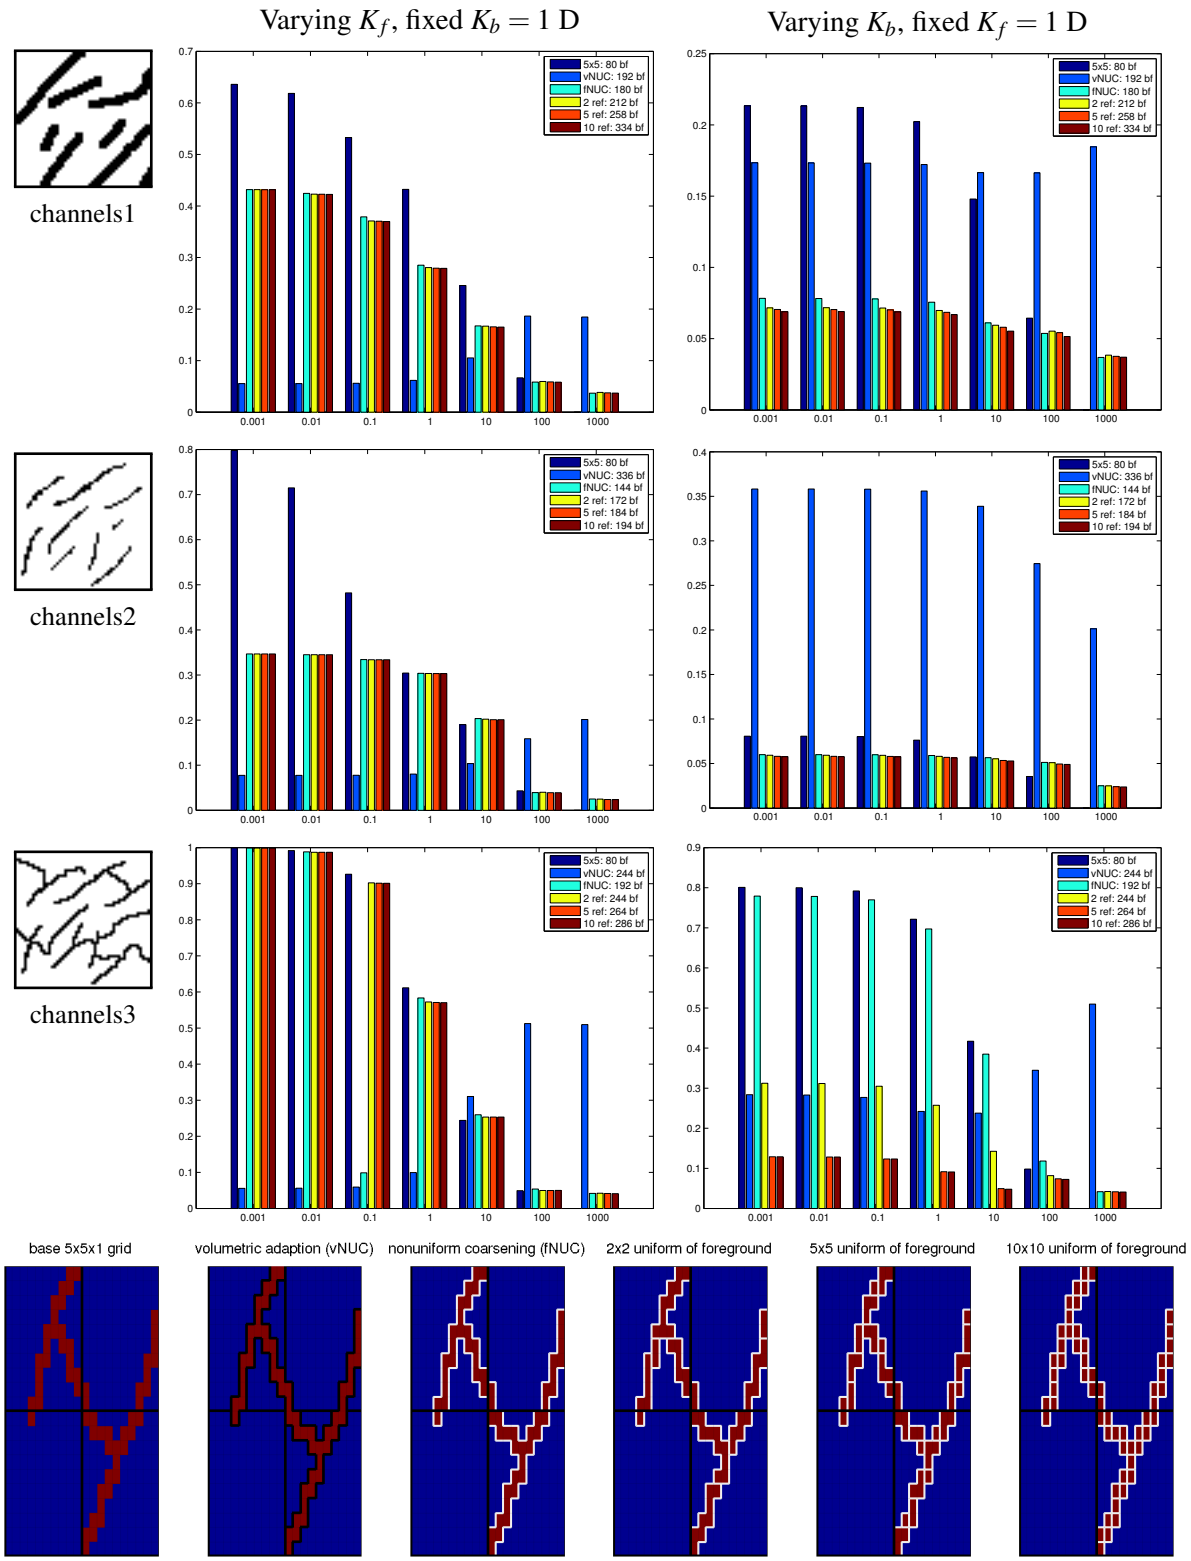

Figure 6 *Relative L^1 error in coarse-scale fluxes for the 2-block MsMFE method for a set of cases with barriers or high-permeable streaks. The small plots in the left column show the permeability masks with black color denoting the foreground permeability and white color denoting the background permeability. The lower plots show a zoom of the six different grids used: red color is foreground permeability, black color is coarse-block faces before subdivision, and white color subdivision lines.*

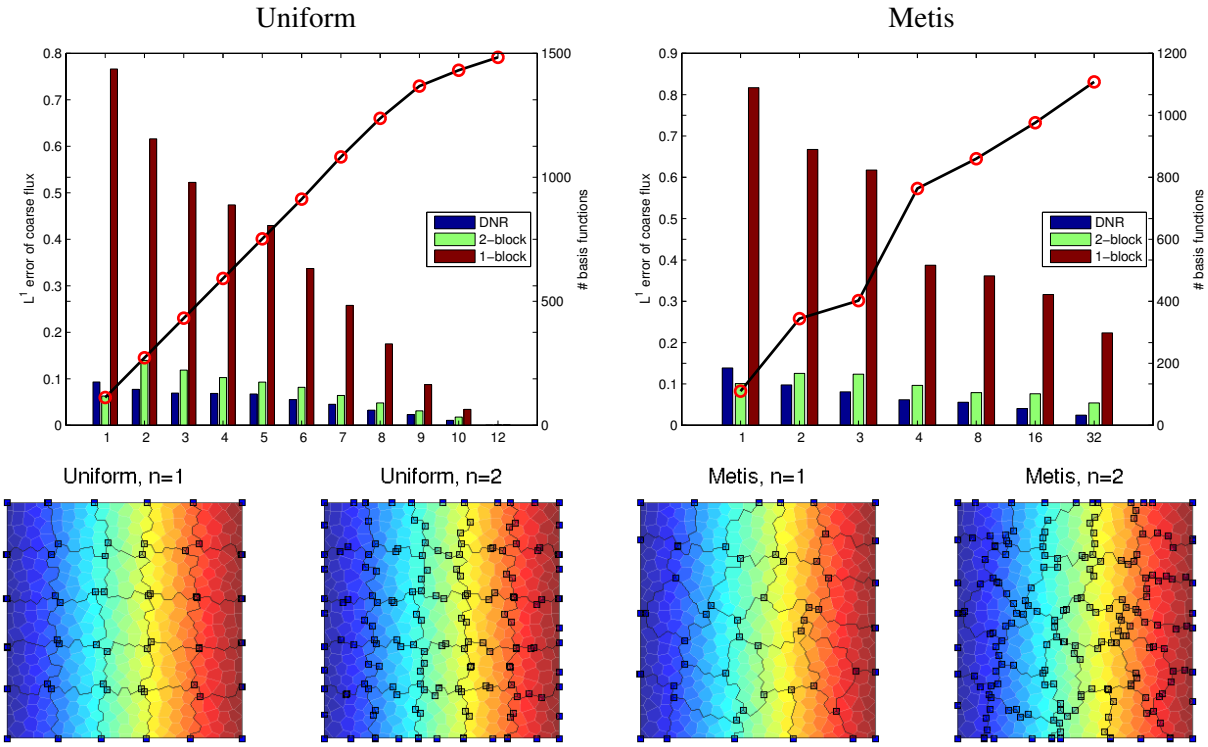

Figure 7 *L^1 errors of the coarse-scale fluxes and the number of coarse blocks involved in the construction of basis functions for a homogeneous medium with linear flow represented on a 2D PEBI grid. The upper left plot shows results for an almost uniform 5×5 coarse grid with face refinement derived from a $5n \times 5n$ Cartesian grid. The upper right plot shows results for a 25-block coarse grid generated by Metis with face refinement derived from a 25n-block Metis partition.*

subdivision of the high-permeability streaks is needed to reduce the error.

In the next example, we apply the DNR and MsMFE methods to a 2D unstructured test case.

Example 5 (2D PEBI) *Unstructured grids will naturally impose coarse faces that do not follow the axial directions. To investigate the effect this has on the solution quality, we consider a square domain with a homogeneous permeability field represented on a 2D unstructured PEBI grid. We consider two different coarse grids: an almost uniform grid obtained by sampling from a 5×5 Cartesian partition and a k-way Metis partition [8] with 25 blocks. Figure 7 shows a convergence study of the errors in coarse-scale fluxes obtained on a sequence of coarse grids with varying degree of face refinement for the DNR and the two MsMFE methods. As should be expected, the errors of the DNR and the 1-block methods decay with increasing face refinement. The error in the 1-block method is particularly high because constant flux is imposed along coarse faces that do not align with the principal axis of flow. For the 2-block method, the error* increases *significantly when the coarse faces are subdivided and we need to introduce a quite fine partition before the error comes back to the level of the base case, in particular with the uniform partition. Results for a set of lognormal permeabilities are qualitatively the same and thus not reported.*

Figure 8 reports a comparison of various face refinements for two high-contrast models: 'channels3' from Figure 6 with high-permeable streaks on a low-permeable background or low-permeable barriers in a high-permeable background. For the barrier cases, the highest accuracy is obtained if the faces are adapted to fit the barrier exactly (denoted '10 mD' in the figure). The nonuniform method ('fNUC' in the figure) segments permeability and then merges small blocks to generate a nonuniform volumetric partition that has approximately the same number of blocks as the base case. The face partition is

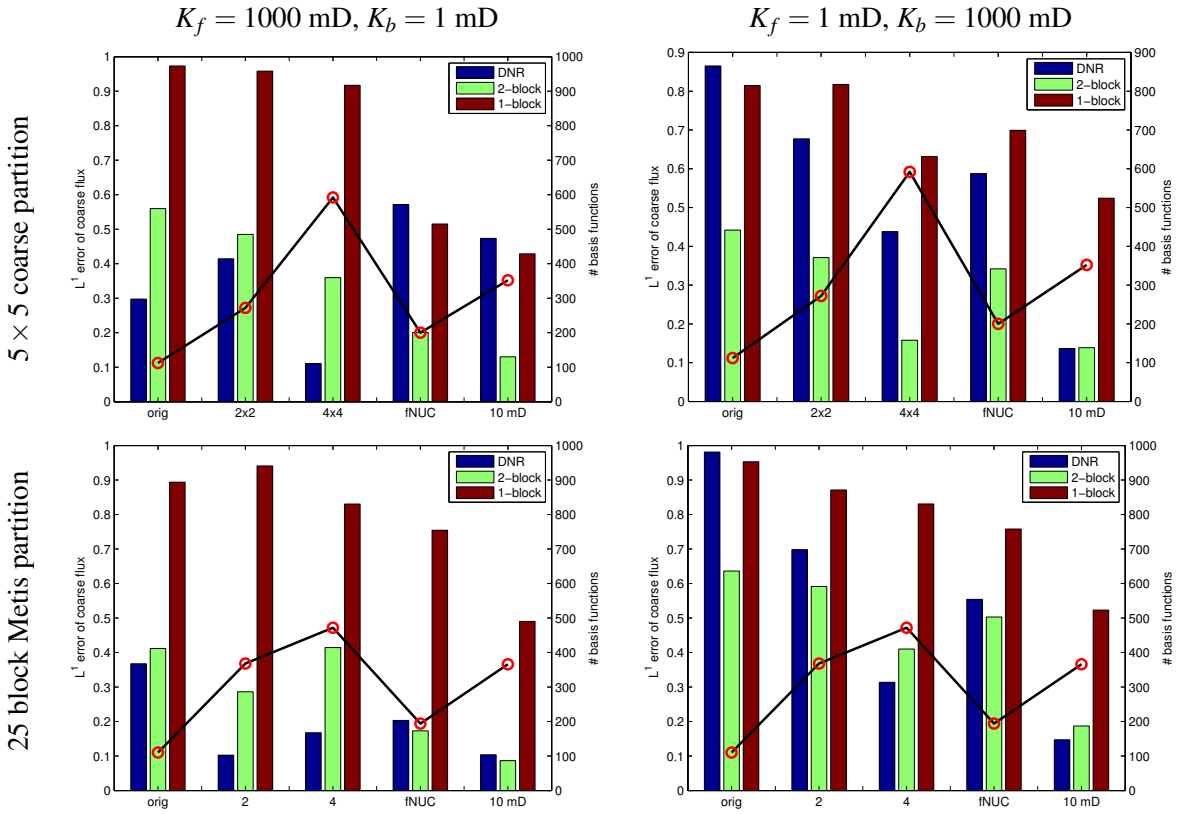

Figure 8 *L^1 errors of the coarse-scale fluxes and the number blocks involved in the computation of basis functions for various types of face refinement on model 'channels3' with linear flow.*

then generated by intersecting the two volumetric partitions and will hence not necessarily capture the barriers exactly, which explains why this method gives much smaller improvement in accuracy. For high-permeable streaks, the 2-block method still has its lowest error when the grid is adapted to the high-permeable streaks, whereas the DNR method has significantly lower error when using a refinement that does not adapt to the structure of the permeability field.

Figure 9 reports a similar refinement study for two models: one with lognormal permeability and one with lognormal permeability and low-permeable barriers. For the first model, we observe three interesting trends: (i) uniform partitions systematically give higher accuracy than Metis partitions; (ii) whereas a uniform face partition improves the DNR solutions, nonuniform partitions adapted to the permeability do not; and (iii) all face refinements diminish *the accuracy for the 2-block method. The flow pattern in the second model will to a large extent be dictated by the low-permeable barriers and the best accuracy is obtained if we add face refinement that adapts to the barriers using either a single threshold value of 10 mD or a partition of permeability into four levels.*

As derived herein, the DNR method is in principle applicable to unstructured grids also in 3D. In practice, however, the method proved to be significantly more difficult to implement than the two MsMFE methods. To build pressure boundary conditions, one must assume that coarse faces are connected and do not have holes (in 3D). Both requirements are needed to generate approximate piecewise linear boundary conditions with apex in a face centroid. This means that compared with the MsMFE methods, the DNR method requires more complex data structures and quite intricate algorithms to construct coarse grids. In 2D, we have implemented methods that guarantee that both requirements are fulfilled. In 3D, we have only been able to ensure that the requirements are fulfilled in certain special cases with a fine grid. In the last example, we will therefore only use the two MsMFE methods.

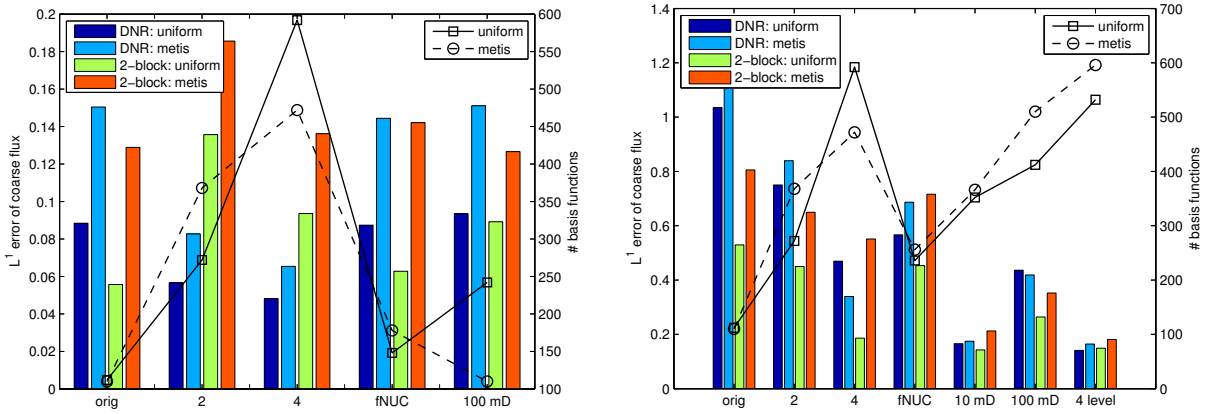

Figure 9 *The left plot shows a comparison of face refinement for uniform and Metis partitions for a model with lognormal permeability and linear flow. The right plot shows a similar lognormal model with low-permeable barriers (0.01 mD) imposed according to mask 'channels3'.*

Example 6 (Realistic model) *We consider a realistic model of a deepwater environment that has been generated using process simulation instead of a geostatistical method. The model is characterized by a large number of small void spaces that are scattered throughout the reservoir volume, resulting in poor vertical (and lateral) connection and making the model somewhat reminiscent of a danish pastry. Altogether, the model contains 150 vertical layers, but many of the layers are very thin and have partially been eroded away, resulting in increased geometrical complexity.*

To partition the grid, we use a simple strategy (see Figure 10): first, we introduce a 6×6 *partition in the lateral direction by using the centroids of the cells to sample from a* $6 \times 6 \times 1$ *partition of the model's bounding box, followed by a simple majority-vote algorithm to ensure that each vertical stack of cells is assigned to a single coarse block. Each coarse block is then post-processed into a set of singly-connected blocks. In a similar way as in Example 3, we can introduce* a posteriori *grid adapting by refining the faces of all blocks that contain cells in which the pressure does not fall inside the span of the boundary values. The faces of all such blocks are subdivided by computing the face intersection with a* $3n \times 3n$ *volumetric partition.*

Figure 11 compares the fine-grid solution and the 2-block MsMFE solution for a 6×6 *partition with and without face refinement. In both cases, there are 673 blocks and for a fine grid with* 90467 *cells this gives an upscaling factor of 134. The corresponding number of blocks involved in the computation of 2-block basis functions are 3164 and 5254, respectively. Introducing face refinement in some of the blocks* a posteriori *improves the approximate solution, as can clearly be seen in the plot of pressure discrepancy. Using face refinement, we can obtain (significantly) better resolution for the same number of basis functions, especially for the pressure.*

Concluding remarks

We have presented two new methodological developments. First, the Dirichlet–Neumann representation (DNR) method [10] has been extended to 3D and fully unstructured grids. Second, the multiscale mixed finite-element (MsMFE) method [2] has been extended to coarse grid having subdivided coarse faces so that more than one basis function may be associated with each pair of coarse blocks. The 2-block formulation of the MsMFE method, does not reproduce linear flow when coarse faces are subdivided. Subdividing coarse faces will therefore not improve the error if the flow field has a dominant linear component. Likewise, the DNR method only reproduces linear flow if extra segments one cell wide are added at each vertex of the coarse grid. On the other hand, although reproduction of linear flow is important for smooth heterogeneities, it may be less important for channelized and high-contrast media.

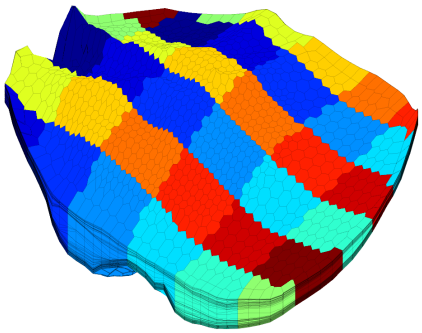

a) lateral 6×6 partition: 35 blocks

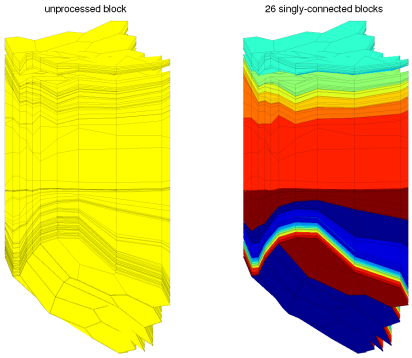

b) split lateral block into singly-connected blocks

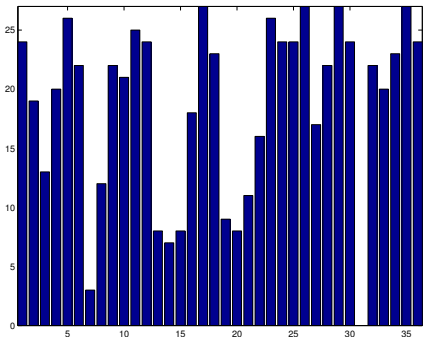

c) number of new blocks in each lateral block

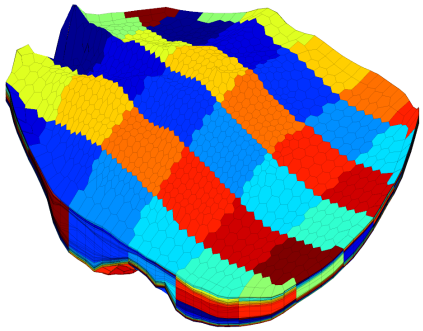

d) *a priori* partition: 673 blocks, 2255 faces

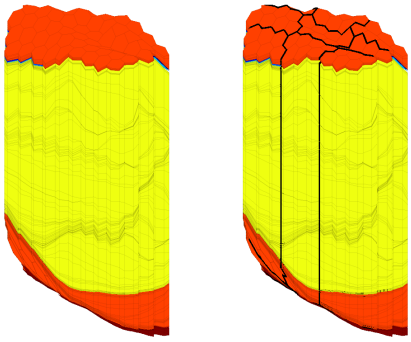

e) subdivision of faces using $3n \times 3n$ partition

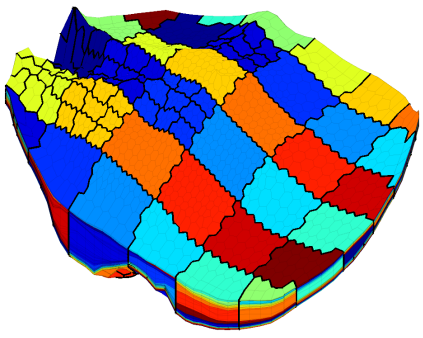

f) *a posteriori* partition: 673 blocks, 4361 faces

Figure 10 *Grid partition of the real-life model based upon a* 6×6 *lateral partition.*

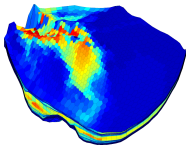

Horizontal permeability

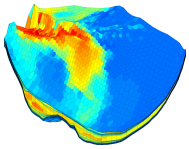

Porosity

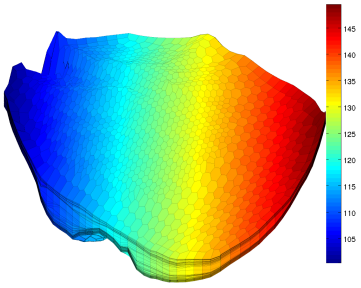

pressure (fine-scale solver)

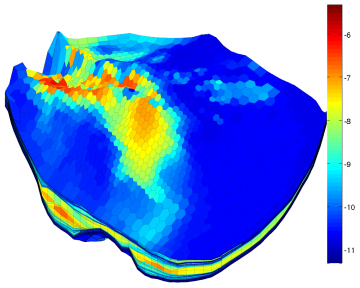

logarithm of cell velocity (fine-scale solver)

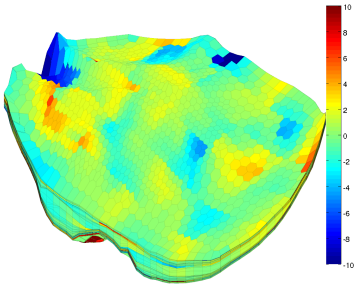

pressure discrepancy (MsMFE)

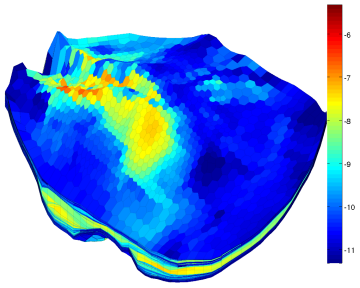

logarithm of cell velocity (MsMFE)

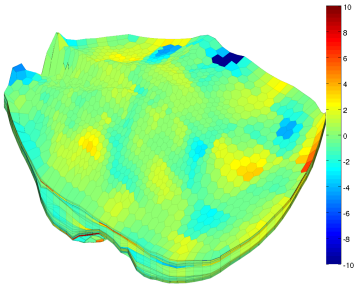

pressure discrepancy (MsMFE-fr)

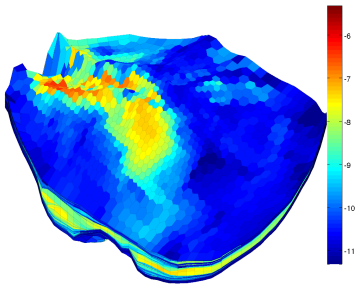

logarithm of cell velocity (MsMFE-fr)

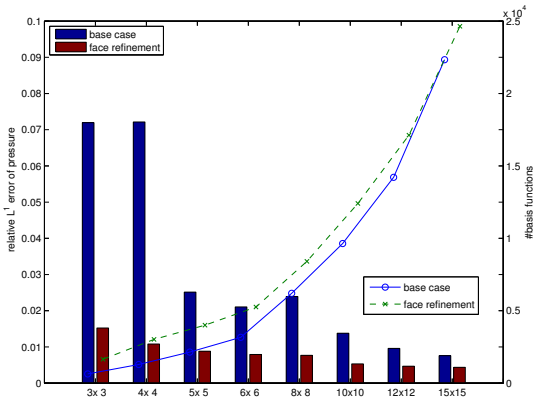

pressure discrepancy $n \times n$

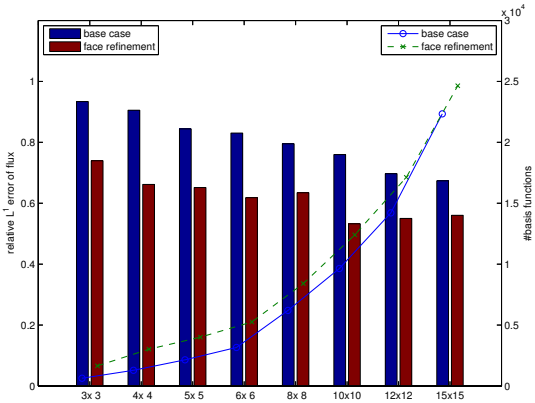

flux discrepancy $n \times n$

Figure 11 *Fine-scale and 2-block MsMFE solutions (with and without face refinement) for the real-life reservoir model subject to linear boundary conditions with the pressure drop making a $\pi/4$ angle with the x-axis. The last row reports discrepancies in flux and pressure for a series of coarse grids based upon a lateral $n \times n$ partition.*

Through a series of numerical experiments, of which a few have been reported herein, we have shown that subdividing coarse faces to capture high-flow effects will increase the resolution of the DNR and the MsMFE methods both for synthetic and realistic cases. Barriers and other no-flow effects, on the other hand, are captured more accurately by using volumetric grid adaption. Altogether, the DNR and MsMFE methods are quite robust and generally produce low errors compared with standard upscaling methods. However, the methods will occasionally lead to large errors for specific combinations of coarse grids and face refinements. One solution to cure this problem would be to use *a posteriori* adaption; a simple approach in this direction was presented for subsets of the SPE10 data set.

Acknowledgment

The authors are grateful to Adedayo Oyerinde for building the real-life model and helpful discussions.

References

[1] 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), ISSN 1540-3459.

[2] Chen, Z. and Hou, T.Y. [2003] A mixed multiscale finite element method for elliptic problems with oscillating coefficients. *Math. Comp.*, **72**, 541–576.

[3] 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/`.

[4] Hauge, V.L., Lie, K.A. and Natvig, J.R. [2012] Flow-based coarsening for multiscale simulation of transport in porous media. *Comput. Geosci.*, **16**, 391–408, ISSN 1420-0597, doi: 10.1007/s10596-011-9230-x.

[5] 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.

[6] 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, ISSN 1420-0597, doi:10.1007/s10596-007-9074-6.

[7] Lie, K., Krogstad, S., Ligaarden, I.S., Natvig, J.R., Nilsen, H.M. and Skaflestad, B. [2012] Open-source MATLAB implementation of consistent discretisations on complex grids. *Comput. Geosci.*, **16**, 297–322, ISSN 1420-0597, doi:10.1007/s10596-011-9244-4.

[8] Metis [2012] Metis – serial graph partitioning and fill-reducing matrix ordering. Url: `http://glaros.dtc.umn.edu/gkhome/views/metis`.

[9] Natvig, J.R. et al. [2011] Multiscale mimetic solvers for efficient streamline simulation of fractured reservoirs. *SPE J.*, **16**(4), doi:10.2018/119132-PA.

[10] Yang, Y., Li, D., Parashkevov, R. and Wu, X.H. [2011] A Dirichlet Neumann representation method for simulating flows in reservoirs. *2011 SPE Reservoir Simulation Symposium, The Woodlands, Texas, USA, 21-23 February 2011*, doi:10.2118/141512-MS.