

# Flow-Based Grid Coarsening for Transport Simulations

Vera Louise Hauge

Knut–Andreas Lie

Jostein R. Natvig

July 15, 2010

## Abstract

Geological models are becoming increasingly large and detailed to account for heterogeneous structures on different spatial scales. To obtain simulation models that are computationally tractable, it is common to remove spatial detail from the geological description by upscaling. Pressure and transport equations are different in nature and generally require different strategies for optimal up-gridding. To optimize the accuracy of a transport calculation, the coarsened grid should generally be constructed based on *a posteriori error* estimates and adapt to the flow patterns predicted by the pressure equation. However, sharp and rigorous estimates are generally hard to obtain, and herein we therefore consider various ad-hoc methods for generating flow-adapted grids. Common for all, is that they start by solving a single-phase flow problem once and then continue to form a coarsened grid by agglomerating cells from an underlying fine-scale grid. We present several variations of the original method. First, we discuss how to include *a priori* information in the coarsening process, e.g., to adapt to special geological features or to obtain less irregular grids in regions where flow-adaption is not crucial. Second, we show how different algorithmic choices can simplify the matrix structure of the discretized system and hence lead to reduced computational complexity. Finally, we demonstrate how to improve simulation accuracy by dynamically adding local resolution near strong saturation fronts.

## Introduction

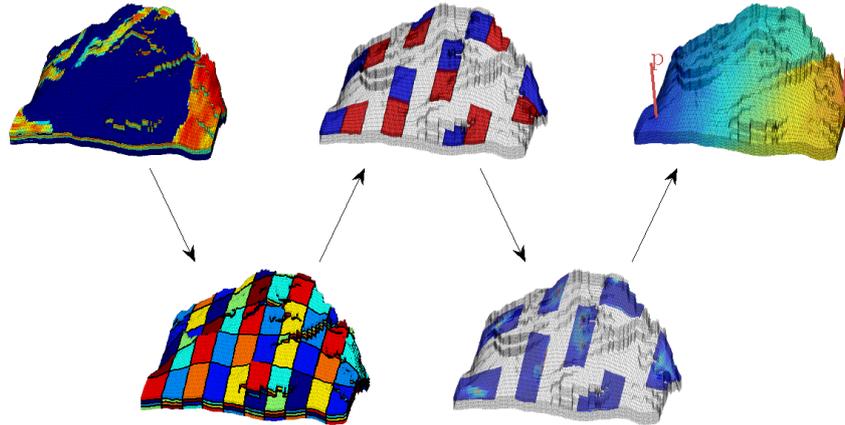
The main purpose of reservoir simulation is to provide predictions of the movement of hydrocarbon phases and water that will help oil and gas companies make better decisions on how to develop and produce their assets. The complexity of the workflows that lead to decisions is ever increasing, and advances in reservoir characterization, production optimization, and real-time reservoir management is leading to continued demand for faster and more advanced flow simulation tools. In particular, optimizing the recovery from mature and brown field assets will require multi-fidelity simulators that have a lot of flexibility and scalability to enable reservoir engineers to evaluate many (different) scenarios.

Coming up with a satisfactory solution is a challenging and daunting task and in this paper, we will only address a small part of the problem: development of multi-fidelity transport solvers to overcome the gap in resolution between geological and simulation models. Geo-cellular models resulting from structural and petrophysical modelling typically contain significantly more detail than what the reservoir engineer can afford if the simulation is to fit in memory and finish within a reasonable time frame of what he/she feels is necessary to capture the flow dynamics of a particular scenario with sufficient detail.

The traditional approach has been to use upgridding to create a new grid model with reduced spatial resolution and upscaling to bring petrophysical properties from the high-resolution geological description down to the new grid. A large number of different strategies have been developed to minimize the errors introduced in this model-reduction process, see e.g., (Durlafsky, 2003, 2005; Christie, 1996; Farmer, 2002). Upgridding and upscaling is generally a manual process and choosing the 'right' method and model resolution can be highly problem dependent and very time-consuming. The problem becomes more complicated as changes are introduced in the reservoir description to match observed (dynamical) data. Ideally, all changes should be made to the fine-scale geological model. However, because the turnaround time of traditional up- and downscaling processes is typically much larger than the man-hours allocated to the modelling project, one ends up with incompatible models at different spatial resolutions.

Herein, we will consider a multiscale approach to geological modelling. This approach differs from the traditional upgridding/upscaling approach in the sense that a fine-grid model is present at all times. Then it is up to the multiscale simulator to (automatically) coarsen the grid to reduce the number of degrees of freedom to a level that is sufficient to resolve flow physics and satisfy requirements on computational costs. The first component in such a simulator is a multiscale flow solver (Hou and Wu, 1997; Efendiev and Hou, 2009; Jenny et al., 2003) that captures the fluid flow as a linear combination of a set of numerically computed basis functions. As such, the multiscale flow solver can be considered either as a robust upscaling method, or as a single-step upscaling-downscaling method that delivers approximate fine-scale fluxes. The basis functions are computed by solving localized flow problems, and the main distinction between different multiscale methods is how the local flow problems are constructed. The methods presented in the following are developed to accompany a particular method (Aarnes et al., 2008), but all ideas presented can readily be combined with any multiscale flow solver that produces *conservative* fluxes on coarse, fine, and intermediate grids. Previous research (Aarnes et al., 2008) has shown that our particular multiscale flow solver gives the best performance when the associated coarse grid follows geological structures. For corner-point grids, this is typically achieved through a regular partitioning in index space (using the underlying logical  $ijk$  numbering). Figure 1 gives a visual illustration of a multiscale flow solver.

The second component of a multiscale simulator is an efficient transport solver that uses the flow field (pressure and total velocity) from the multiscale flow solver to evolve fluid saturations and compositions. To this end, there are several alternatives, depending upon whether one chooses to evolve the saturations on a coarse grid (e.g., as used in the multiscale flow solver), on the underlying fine grid, or on some intermediate grid. Herein, we will use the latter approach. Solving saturation equations on the coarse scale will in many cases be too inaccurate and solving it on the fine grid may be too costly. We will



**Figure 1** Illustration of a multiscale flow solver applied to a model from the SAIGUP study (Manzocchi et al., 2008). The coarse grid is generated as a regular partitioning of the underlying fine grid. Localized flow problems are set up (here, between each pair of blocks that share a common face) and then solved numerically. Finally, the basis functions are used to construct a global flow solution.

therefore generate a separate grid for the saturation solver that adapts to the flow patterns predicted by the flow solver. Ideally, such a flow-adapted grid should be constructed using a posteriori error estimates. However, obtaining sharp and rigorous estimates are generally hard and several authors have therefore developed ad-hoc methods (Durlofsky et al., 1997; Aarnes et al., 2007) that are able to capture flow and transport in highly heterogeneous reservoirs with good accuracy. In a flow-based method, the grid is aligned to capture high-flow regions and (clearly) distinguish between regions of high and low flow. By capturing these important characteristics of the flow, one is able to generate coarse grids with a high upscaling factor that still deliver good accuracy.

In this paper, we will combine several ideas for speeding up the transport solves in a multiscale simulator setting. That is, we will make a few algorithmic improvements for the nonuniform coarsening method of Aarnes et al. (2007) and borrow ideas from the adaptive multiscale method of Lee et al. (2009) and Zhou et al. (2009) and reordering methods of Natvig and Lie (2008a).

### Mathematical Model and Numerical Discretization

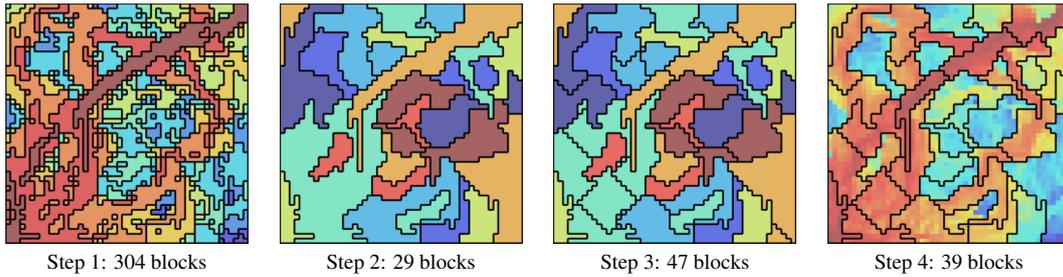
In general, multi-fidelity simulators need to cover a wide range of enhanced recovery processes including biological, chemical, electrical, gas-based, thermal, and water-based methods. However, in the following we will only consider the simplified setting of an incompressible, immiscible two-phase flow, described by a set of flow equations for the global pressure  $p$  and the total velocity  $\vec{v}$

$$\nabla \cdot \vec{v} = h_p, \quad \vec{v} = -\mathbf{K}\lambda(S)\nabla p, \tag{1}$$

and a transport equation for the saturation of one of the phases,

$$\phi \frac{\partial S}{\partial t} + \vec{v} \cdot f(S) = h_S. \tag{2}$$

These equations are defined over a singly-connected domain, represented by a grid that consists of a set of grid cells  $c_i, i = 1, \dots, n$ . No further assumptions are made on the geometry and topology of the grid, apart from the requirement of an explicit mapping  $\mathcal{N}(c)$  between cell  $c$  and its nearest neighbours. Most of the ideas we present in the following will therefore be applicable to any matching, unstructured, polyhedral grid. To keep the presentation as simple as possible, our examples will, with two exceptions, focus on 2D Cartesian grids, taken from individual layers of the widely used SPE 10 benchmark Christie and Blunt (2001). Likewise, as our interest is primarily in the transport solver, we make no assumptions



**Figure 2** Illustration of Steps 1 to 4 for a quarter five-spot simulation of a  $50 \times 50$  excerpt from Layer 68 from Model 2 of the SPE 10 benchmark (Christie and Blunt, 2001). The coarsening parameters are  $N_L = 20$  and  $N_U = 100$ . In the rightmost plot, the colours represent logarithm of the underlying fine-scale velocity field.

about the flow solver except that it produces mass-conservative fluxes on each cell. In a multiscale setting, our primary example of such a solver would be the multiscale mixed-finite method of Aarnes et al. (2008). In the following,  $v_{ij}$  will denote the flux over fine-cell interface  $\gamma_{ij}$  between cells  $c_i$  and  $c_j$ .

The remains of the paper will focus on transport solvers defined over a coarse grid that is constructed by grouping sets of cells into blocks  $B_\ell$ ,  $\ell = 1, \dots, N$ . The simplest way of representing such a coarse grid is by a partition vector  $p$  with  $n$  elements, for which element  $p_i$  assumes the value  $\ell$  if cell  $c_i$  is member of block  $B_\ell$ . Representing the coarse grid by a partition vector gives us great flexibility in the shape of individual grid blocks and also opens up for a simple (interactive) manual editing, if deemed necessary. Herein, however, we will *not* use any manual editing to improve grid quality. More details of the grid generation will be given from the next section and onward.

Having created a coarse grid, the next step is to construct a coarse-grid transport solver. To this end, we assume that the saturations are constant over each grid block, i.e.,  $S_\ell = |B_\ell|^{-1} \int_{B_\ell} S(x) dx$ . Then, a conservative coarse-grid discretization is obtained by summing a standard single-point upwind discretization for all cells in a block:

$$S_\ell^{n+1} = S_\ell^n + \frac{\Delta t}{\int_{B_\ell} \phi dx} \int_{B_\ell} h_S(S^{n+1}) dx - \frac{\Delta t}{\int_{B_\ell} \phi dx} \left[ f(S_\ell^{n+1}) \sum_{\gamma_{ij} \subset \partial B_\ell} \max(v_{ij}, 0) - \sum_{k \neq \ell} \left( f(S_k^{n+1}) \sum_{\gamma_{ij} \subset \Gamma_{k\ell}} \min(v_{ij}, 0) \right) \right]. \quad (3)$$

Notice that if there is bi-directional flow across  $\Gamma_{k\ell} = \partial B_k \cap B_\ell$ , then the phase-flux across  $\Gamma_{k\ell}$  is approximated using both  $S_\ell$  and  $S_k$ . Thus, although (3) stems from a single-point upwind scheme on the fine grid, we may obtain a two-sided upwind scheme on the coarse grid. Because the method uses fluxes computed on a finer scale to correctly propagate information over the faces of the coarse grid, it can be viewed as multiscale solver.

### Flow-Based Nonuniform Coarsening

We will use the nonuniform coarsening method by Aarnes et al. (2007) as our starting point for developing efficient transport solvers on flow-adapted coarse grids. This method partitions the fine grid into coarse blocks according to flow magnitude by separating regions of high and low flow.

The algorithm follows four steps:

1. *Compute an initial partitioning.* To this end, we will use the logarithm of the flow as our indicator function,  $g(c_i) \propto \log |\vec{v}(c_i)|$ , which is segmented into ten uniform bins:  $c_i \subset \tilde{B}_\ell$  if  $g(c_i) \in [g_\ell, g_{\ell+1})$ . Each bin  $\tilde{B}_\ell$  may consist of a multiply connected set of cells and must be postprocessed and split into singly-connected blocks.

2. *Merge small blocks.* If a block  $B'$  has too small volume,  $|B'| < \frac{N_L}{n} |\Omega|$ , it is merged with the neighbouring block  $B$  that has the closest  $g$ -value defined as  $g(B)|B| = \int_B g(c) dx$ .
3. *Refine blocks with too much flow.* If  $\int_B g(c) dx > \frac{N_U}{n} \int_\Omega g(c) dx$ , then
  - (a) Pick an arbitrary cell  $c_0$  belonging to  $\partial B$ .
  - (b) Find the cell  $c_1 \subset B$  that is furthest away from  $c_0$ , using, e.g., the distance between the cell centroids as a metric, and define  $B' = c_1$ .
  - (c) Progressively enlarge  $B'$  by adding cells surrounding  $B'$ ; that is, add  $c \in \mathcal{N}(B')$  if  $c \notin B'$ , as long as  $\int_{B'} g(c) dx \leq \frac{N_U}{n} \int_\Omega g(c) dx$ .
  - (d) Define  $B = B \setminus B'$  and continue to refine  $B$  if the upper bounds are still violated.
4. *Repeat Step 2.*

The four steps are illustrated in Figure 2. Two coarsening parameters determine the coarsening level:  $N_L$  gives a lower bound on the volume of blocks and prevents the algorithm from generating too small blocks, and  $N_U$  gives an upper bound on total amount of flow through each grid block and prevents the algorithm from generating too large blocks. In the original algorithm,  $\mathcal{N}(c)$  is defined as the face neighbours: that is, cells  $c$  and  $\tilde{c}$  are neighbours if they share a common face. For a Cartesian grid, this corresponds to the usual 5-point neighbour relation, and gives rise to the characteristic diamond-shaped cells seen in the two rightmost plots. To get blocks with a more regular shape, we will in the following use a 9-point neighbour relation. We refer the reader to (Hauge et al., 2010) for a more thorough discussion of neighbour relations in coarsening algorithms.

### Imposing *A Priori* Shapes

The algorithm presented in the previous section is fully automated in the sense that the degree of coarsening is determined by two user-supplied parameters ( $N_L$  and  $N_U$ ) and the shape of the blocks is determined by the neighbourhood definition, so that the algorithm creates coarse grids with block boundaries that are aligned with distinct features in the permeability field as reflected in the computed velocity field. For highly heterogeneous reservoirs, and in particular for strongly channelized reservoirs, this results in grids that capture the dominating flow patterns very accurately, even for high upscaling factors. This was demonstrated in (Aarnes et al., 2007), for both structured and unstructured grids, and in particular for Model 2 from the 10th SPE Comparative Solution Project Christie and Blunt (2001), for which flow-based coarsening significantly reduces saturation errors and errors in water cuts compared with a regular Cartesian coarsening for the fluvial part of the model (the 50 layers of the Upper Ness formation). Moreover, Aarnes et al. (2007) demonstrated that the method is robust with respect to coarsening degree and varying flow patterns, as long as the flow field is dominated by the underlying heterogeneity.

Generally, the nonuniform coarsening algorithm is not as robust and efficient as it was somewhat optimistically reported in (Aarnes et al., 2007). On the upper layers of the SPE10 model (the somewhat smoother Tarbert formation), the flow-based grids do not give better accuracy than a straightforward Cartesian coarsening. In fact, we have run numerous experiments with varying coarsening factors for the SPE10 and other models that all indicate that Cartesian coarsening predominantly gives slightly better accuracy in saturation fields and water cuts than the flow-based grid for cases with small or moderate heterogeneity. Altogether, our experiments suggest that the original algorithm, as presented in the previous section, has a tendency of exaggerating the effect of the underlying velocity pattern and thus creates grids that are more irregular than what is needed.

The flow-based grids typically have more irregular blocks with more neighbours and coarse-block interfaces and this tends to increase the coupling in the discrete nonlinear system. The number of couplings in the nonlinear system will typically affect how costly it is to solve, and hence it is desirable to increase

the regularity of the blocks if it does not significantly affect their ability to resolve flow patterns; we will come back to this discussion in the next section. In many cases, the user will have expert knowledge of what has the most influence on accuracy and may want, e.g., to impose *a priori* information on the (local) shape of the coarse blocks. Likewise, there can be other geological features that the user may want to use to create grids that better adapt to the underlying geology. In the rest of this section, we will demonstrate how to use such *a priori* information to create grids that give improved accuracy.

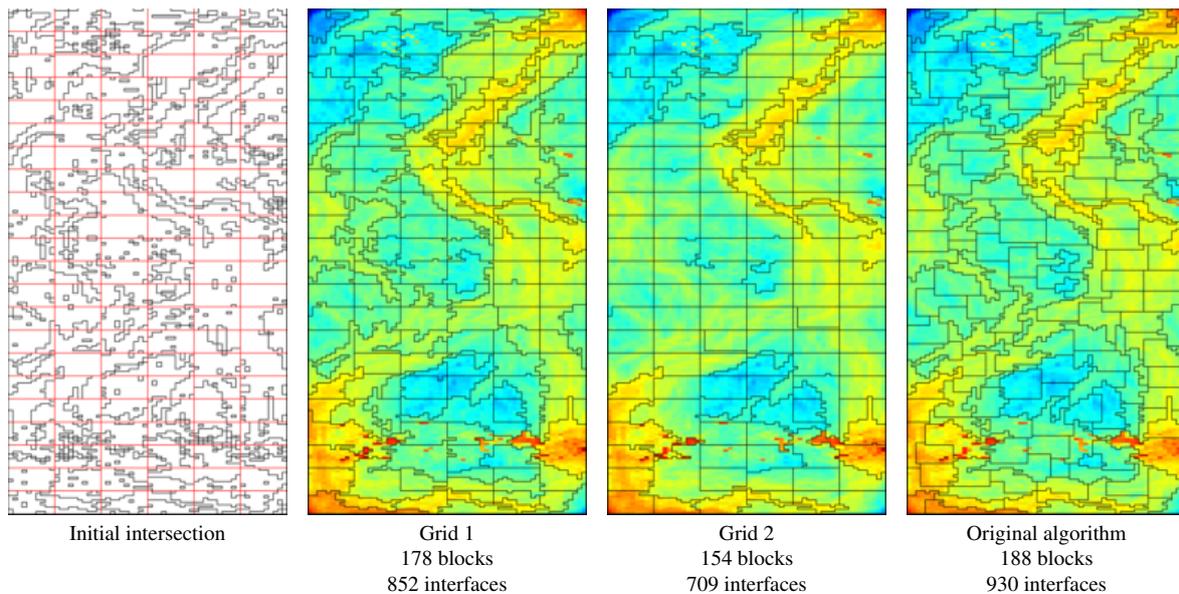
First of all, we propose an additional step in the original algorithm that consists of intersecting the initial flow-based colouring of cells in the first step with an *a priori* partitioning. This intersection will then be the basis for the rest of the steps, which remain unchanged from the original algorithm. We point out that this additional step is applicable to *any* grid, for which the user is able to specify an *a priori* partition vector  $p_a$ . Second, to reduce the influence of the underlying heterogeneity on the coarsening, we look at the initial colouring of cells. Reducing the number of bins means that we increase the size of the blocks resulting from the flow-based colouring, and hence to a large extent preserve the *a priori* partition. If the sizes of the *a priori* blocks are within the bounds specified by the  $N_L$  and  $N_U$  parameters, a large number of these blocks will be left intact by Steps 2–4.

To illustrate the effect of the extra step and the adaptive number of initial bins, we will consider an example in which we seek to impose a regular Cartesian partitioning on the flow-based gridding process.

**Example 1 (Layer 1 of SPE10)** *The leftmost plot in Figure 3 shows the intersection of a regular  $6 \times 22$  Cartesian coarsening with the initial flow-based colouring. Altogether, this generates a finer partitioning as the starting point of the merging and refinement steps. The added interfaces are straight lines that will typically be preserved in low-flow regions, as seen by comparing Grid 1 with the grid generated by the original algorithm. In the original algorithm and for Grid 1, we have used 10 bins in the initial colouring. This amounts to approximately one bin per order of magnitude in the underlying velocity field for the fluvial parts of the SPE10 model. For the layers in the Tarbert formation, as considered here, the logarithmic span in the velocities is significantly smaller. For Grid 2, we therefore have chosen the number of initial bins equal to the logarithmic span in the underlying fine-scale velocity field. As a result, more blocks from the *a priori* partitioning remain unchanged throughout Steps 2–4, and the resulting grid has a much more regular structure than Grid 1. On the other hand, for both Grid 1 and Grid 2, we observe that the high-flow channels are distinctly outlined and corresponds to the ones detected by the original algorithm. In regions of low flow, we also observe that the original grid has a much more complex grid structure, with a larger number of neighbouring connections.*

We have run a large number of different studies using the original and our improved algorithm. Choosing the number of bins according to the logarithm span in velocity (or possibly permeability) seems to be a good choice. The other parameters, however, must be chosen with some care and possibly be fine-tuned to give optimal results. For small and moderate heterogeneities, it seems particularly important to balance the choice of the two partitioning mechanisms. Choosing a coarse *a priori* partitioning and small partitioning parameters ( $N_L, N_U$ ), means that the flow-based partitioning will dominate and any advantages from the *a priori* partitioning disappears. Likewise, choosing a fine *a priori* partitioning and large values for  $N_L$  and  $N_U$ , implies that most cells are merged into large blocks in Step 2 and then refined in Step 3. The resulting grid will hence have the characteristics of the original algorithm, and the effects of the *a priori* partitioning disappears.

*A priori* information can also be used to distinguish different geological features that need to be taken into account and/or preserved during the coarsening. In the next example, we demonstrate how one can use facies numbers as the initial partitioning to ensure that the coarse blocks do not cross facies boundaries. This can be very useful if the facies have different relative permeabilities or capillary pressure curves. By making sure that each coarse block consists of only one facies, one can avoid cumbersome upscaling of relative permeabilities and capillary pressure functions to the coarse grid.

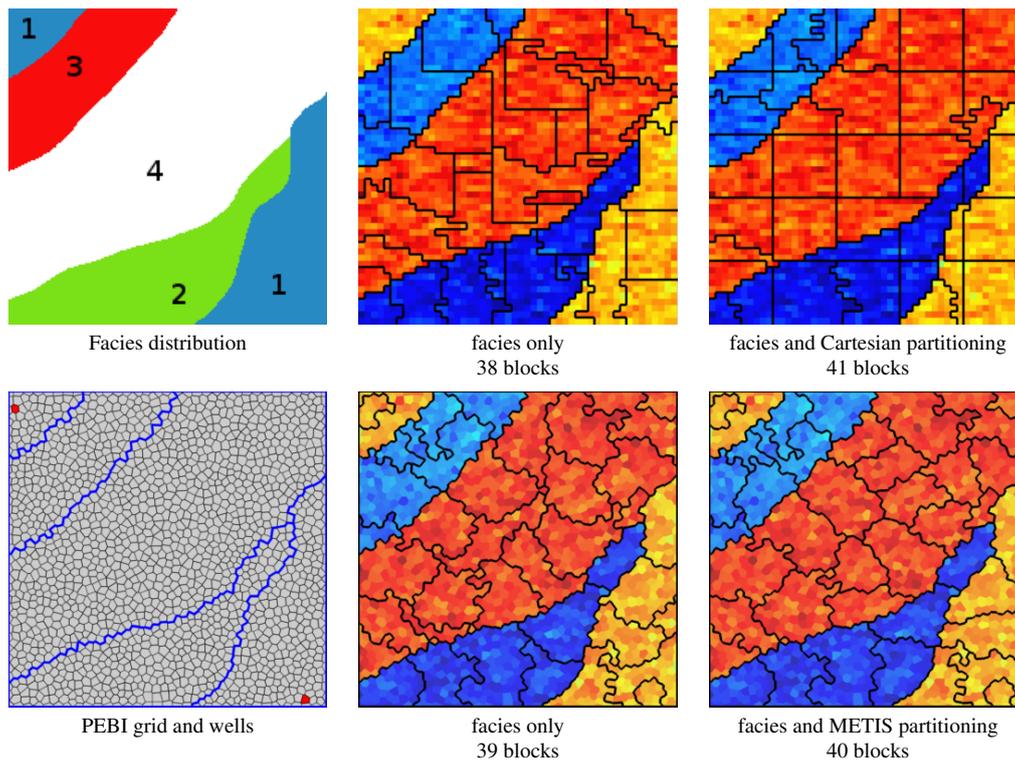


**Figure 3** Nonuniformly coarsened grids for Layer 1 of Model 2 from SPE10. The leftmost plot outlines the intersection of the initial partitioning, where red lines represent the a priori partitioning and black lines the partitioning into bins. Grid 1 is generated with 10 initial bins and Grid 2 with the number of bins adapted to the flow, and for both we used coarsening parameters  $N_L=25$ ,  $N_U=125$  and a uniform  $6 \times 22$  a priori partitioning. The rightmost plot shows the grid generated by the original algorithm with parameters  $N_L = 15$ ,  $N_U = 75$  and a 9-point neighbourhood relation. All three grids are outlined on top of the logarithm of the fine-scale velocity field, computed by solving a quarter five-spot problem.

**Example 2 (Facies model)** We consider a rectangular domain with a facies distribution as shown in the upper-left plot of Figure 4. The permeability distribution follows a lognormal distribution inside each facies with mean values of 400, 20, 35, and 800 mD, respectively. The permeability distribution is sampled on a regular  $50 \times 50$  Cartesian grid. The reservoir is produced by an injector-producer pair located near the upper-left and lower-right corners of the domain, respectively; the wells are shown as red dots in the lower-left plot in Figure 4. For illustration purposes, we generate two coarse grids: for the first one, we only use the facies distribution as our a priori partition vector, and for the second one, we impose a regular Cartesian partitioning, in addition, as discussed for Grid 2 in Example 1. The two grids are shown in the middle and left plot in the upper row of Figure 4. From the plots, we clearly see that the coarse blocks are confined inside a single facies. By additionally imposing an a priori Cartesian partitioning, we increase the number of blocks, but also get more regularity in our coarsening.

To demonstrate that our method is not restricted to Cartesian grids, we consider the same facies and permeability distribution (re)sampled on the PEBI grid (unstructured Voronoi grid) with 1697 cells shown in the lower-left plot of Figure 4. The middle and left plots in the lower row of Figure 4 show two coarse grids by our algorithm. The difference in the two grids lie in how we perform the refinement in Step 3 of the algorithm. In the middle plot, we have used the refinement method outlined above with the unstructured equivalent of the 9-point neighbour definition. Also in this case, the blocks are clearly confined to a single facies, as desired. However, because of the unstructured connections in the fine-scale grid, the coarse blocks are significantly more irregular than in the Cartesian case. In the left plot, we have instead used METIS to perform the refinement step, this in an attempt to improve the regularity of the blocks.

Although the results presented in the previous example are encouraging for PEBI grids, we believe that more research is needed to come up with a better algorithm for refining large blocks in Step 3, particularly in 3D, in a way that imposes more regularity on the resulting blocks. The same problem arises when generating coarse grids to be used in multiscale flow solvers. Because of the complex



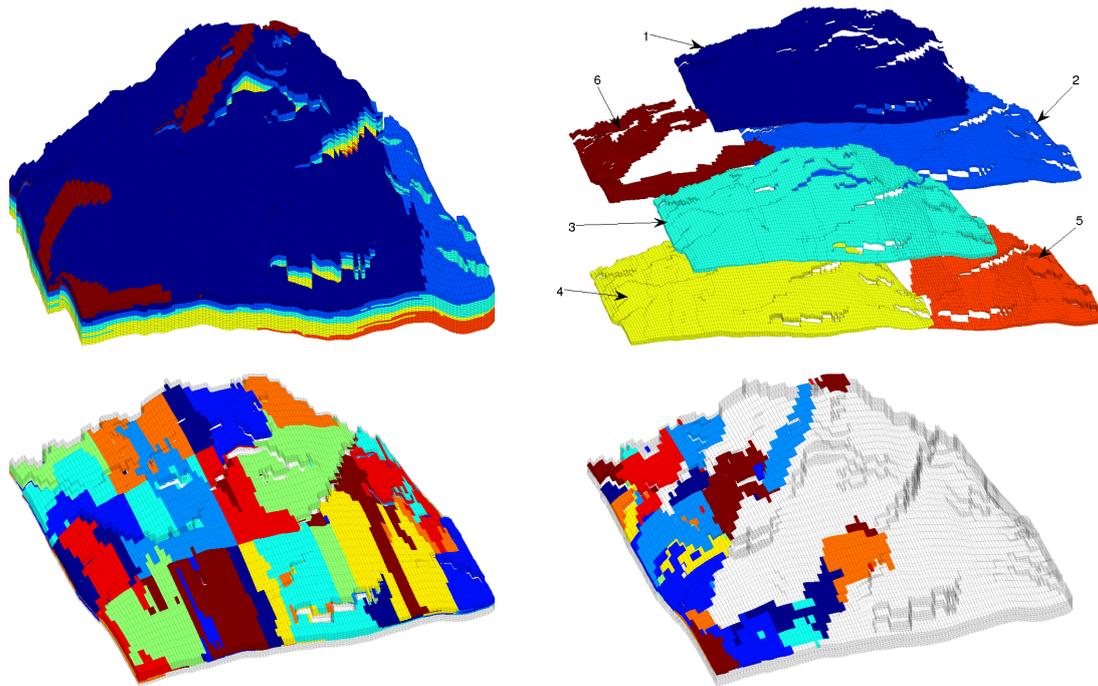
**Figure 4** Coarsening using a priori information of facies distribution. In the upper row, the underlying fine grid is a regular  $50 \times 50$  Cartesian grid. In the lower row, the fine grid is a fully unstructured PEBI grid with 1697 cells.

connection pattern of the underlying unstructured grid, it is generally difficult to come up with coarse blocks that do not have irregular boundaries. One possibility would be to include a smoothing step, but this would need to be carefully designed to distinguish between irregular faces caused by flow adaption and irregular faces induced by the unstructured connections.

In the next example, we will study a realistic 3D model represented in the industry-standard corner-point format, i.e., as a grid that consists of a set of hexahedral cells that are topologically aligned in a Cartesian fashion so that the cells can be numbered using a logical  $ijk$  index. From a coarsening perspective, the underlying  $ijk$  index is very useful and can e.g., be utilized to impose a regular *a priori* partitioning as in Example 1. Here, however, we will use saturation regions as our *a priori* partitioning.

**Example 3 (SAIGUP)** We consider one of the faulted models from the SAIGUP study (Manzocchi et al., 2008), which we have already used in Figure 1 to illustrate the key steps of a multiscale flow solver. The petrophysical parameters for the model were generated on a regular  $40 \times 120 \times 20$  Cartesian grid and then mapped onto a structural model described using the corner-point grid format. The left plot in Figure 5 shows the structural model. The colors represent the six different saturation regions (Eclipse keyword SATNUM), which may or may not correspond to different facies or rock types. Because the main purpose of the example is to illustrate the gridding capabilities on a model with realistic geometry and petrophysical properties, we use a simple injector–producer pair (see Figure 1) and create a relatively coarse flow-based grid.

The coarse grid was created by imposing the six saturation regions as an *a priori* partitioning. Moreover, in Steps 2 and 4, we restricted the neighbourhood definition to only include cells that were part of the same saturation region. As we see from Figure 5, the coarse blocks have complicated shapes but seem to follow the saturation regions; this is particularly evident in region six. The plot may be slightly deceiving with respect to the connection between blocks: blocks that appear to be multiply connected are, in fact, singly connected through cells in deeper layers that are not visible in the plot.



**Figure 5** Flow-based coarsening of a faulted SAIGUP model. The upper-left plot shows the structural model. The upper-right plot shows the six saturation regions. The lower plots show the coarse-grid blocks in saturation regions three and six, respectively.

Another example of flow-based gridding on corner-point grids was presented by Krogstad et al. (2009), who used such grids to accelerate forward simulations in a production optimization workflow.

### Computational Efficiency

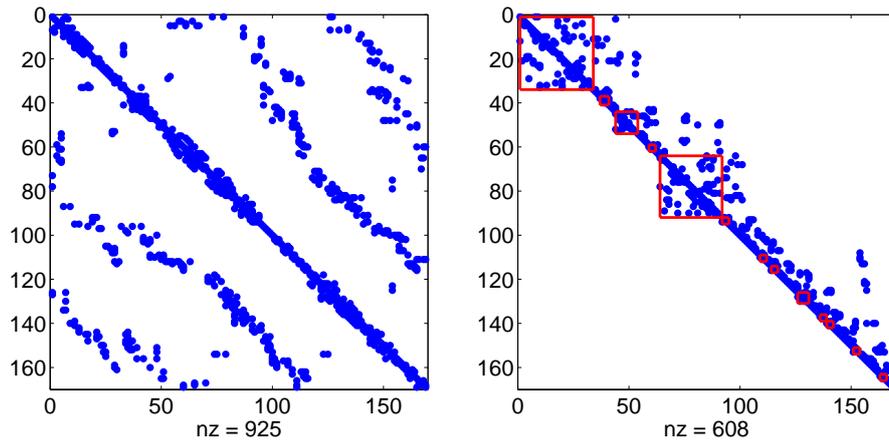
The discretization (3) leads to a system of nonlinear equations that will typically be solved by a Newton–Raphson type method. The computational efficiency of the nonlinear solver will to a large extent depend on the structure and the condition number of the system. The best we can hope for, is that the discrete system has an upper (or lower) triangular form, because then we can use a nonlinear Gauss–Seidel solution procedure and compute the unknown block saturations by backward (or forward) back-substitution. This is clearly possible for a 1D problem. Likewise, in the absence of gravity and capillary forces, (2) has an inherent causality principle that is utilized in streamline methods to transform the multi-dimensional transport equation into a family of 1D problems along streamlines.

Natvig and Lie (2008a) recently demonstrated how this causality principle can be used to compute an optimal flow-based ordering that renders the system in a block-triangular form. If the flow solver is monotone, each diagonal matrix block will correspond to only one grid block and the solution can be computed block-by-block, moving gradually downstream from wells or other fluid sources. For non-monotone flow solvers, there will be some circulation in the discrete fluxes, which will lead to larger matrix blocks that contain grid blocks that are circularly dependent. Still, the system can be solved by an efficient block-wise back-substitution procedure, in which the circularly dependent grid blocks (henceforth called connected components) are solved for simultaneously. Furthermore, if the linearization is performed *locally* on each matrix block, we gain local control over the nonlinear (Newton–Raphson) iterations and thereby obtain highly efficient and (near) optimal nonlinear solvers (Natvig and Lie, 2008a,b). Finding the connected components and the optimal ordering are standard and efficient  $\mathcal{O}(N)$  operations from graph theory that are easy to implement.

Rather than considering a specific nonlinear solver, we will consider the degree to which the discrete

**Table 1** Comparison of the matrix structure on two layers of SPE10.

	Layer 1			Layer 37		
	Grid 2	Original	Cartesian	Grid 2	Original	Cartesian
# grid blocks	154	188	220	169	205	220
# scalar components	1	0	75	1	0	16
# connected components	4	1	50	1	2	24
largest component	89	188	10	168	203	48
# off-diagonal elements	507	798	503	756	911	574



**Figure 6** Matrix structure for the flux vector of Grid 2 on Layer 37. The left plot shows the matrix with the ordering coming from the coarsening algorithm. The right plot shows the reordering when the bi-directional fluxes are replaced by net fluxes, in which case 65 of the saturation values depend only on their upwind neighbours, whereas the remaining 104 have some circular dependence and are part of one of the thirteen connected components. The corresponding matrix blocks are marked in red: two large, one intermediate, and ten small.

system can be reordered into a triangular form as a measure of the efficiency of the grid. Our idea is that if such a structure exists, any efficient nonlinear solver should ideally be able to exploit it.

**Example 4 (Layers 1 and 37 of SPE10)** Continuing from Example 1, we consider a quarter five-spot problem on two layers from the SPE10 model: Layer 1 from the Tarbert formation and Layer 37 from the fluvial Upper Ness formation. For each layer, we will use three different grids: a  $10 \times 22$  Cartesian grid, a grid generated by the original algorithm from (Aarnes et al., 2007), and Grid 2 from Example 1 and its equivalent on Layer 37. Table 1 reports the corresponding number of scalar components, number of connected components, number of blocks in the largest connected component, and number of off-diagonal entries.

Let us first look in detail on a few of the grids. We start by the Cartesian grid for Layer 1. In the discrete system there are 50 connected components that contain at most ten grid blocks and 75 scalar components. Next, consider Grid 2 on Layer 37: out of the 169 blocks in the grid, all blocks except one have circular dependence and are part of a single connected component. Hence, all saturation values must be computed simultaneously by solving a  $168 \times 168$  nonlinear system. Preliminary studies indicate that reducing the bi-directional connections arising from the fine-scale fluxes into net fluxes at the block interfaces, will in general reduce the number of entries in the matrix, see Figure 6. In addition, several connections will disappear, thereby breaking up most of the circular dependencies. The rightmost plot gives an indication of the simplified structure using net fluxes. More research regarding accuracy using bi-directional and net fluxes is in progress.

## Dynamically Adaptive Grid

For displacements with strong displacement fronts, the majority of the projection error is associated with inaccurate representation of the fluid front. Think of a typical Buckley–Leverett profile: in the unswept area ahead of the displacement front, the solution is constant and can be accurately represented on a relatively coarse grid. Likewise, behind the displacement front, the solution is smooth and slowly varying and can hence be evolved on a coarse grid. In the absence of capillary forces, or other second-order terms in the transport equation (2), the displacement front is a discontinuity that needs high grid resolution to be accurately approximated. Motivated by these observations, we will in this section demonstrate how the simulation accuracy can be significantly improved by dynamically adding local resolution near strong saturation fronts. Somewhat similar ideas were used by Lee et al. (2009) and Zhou et al. (2009) in their adaptive multiscale finite-volume method.

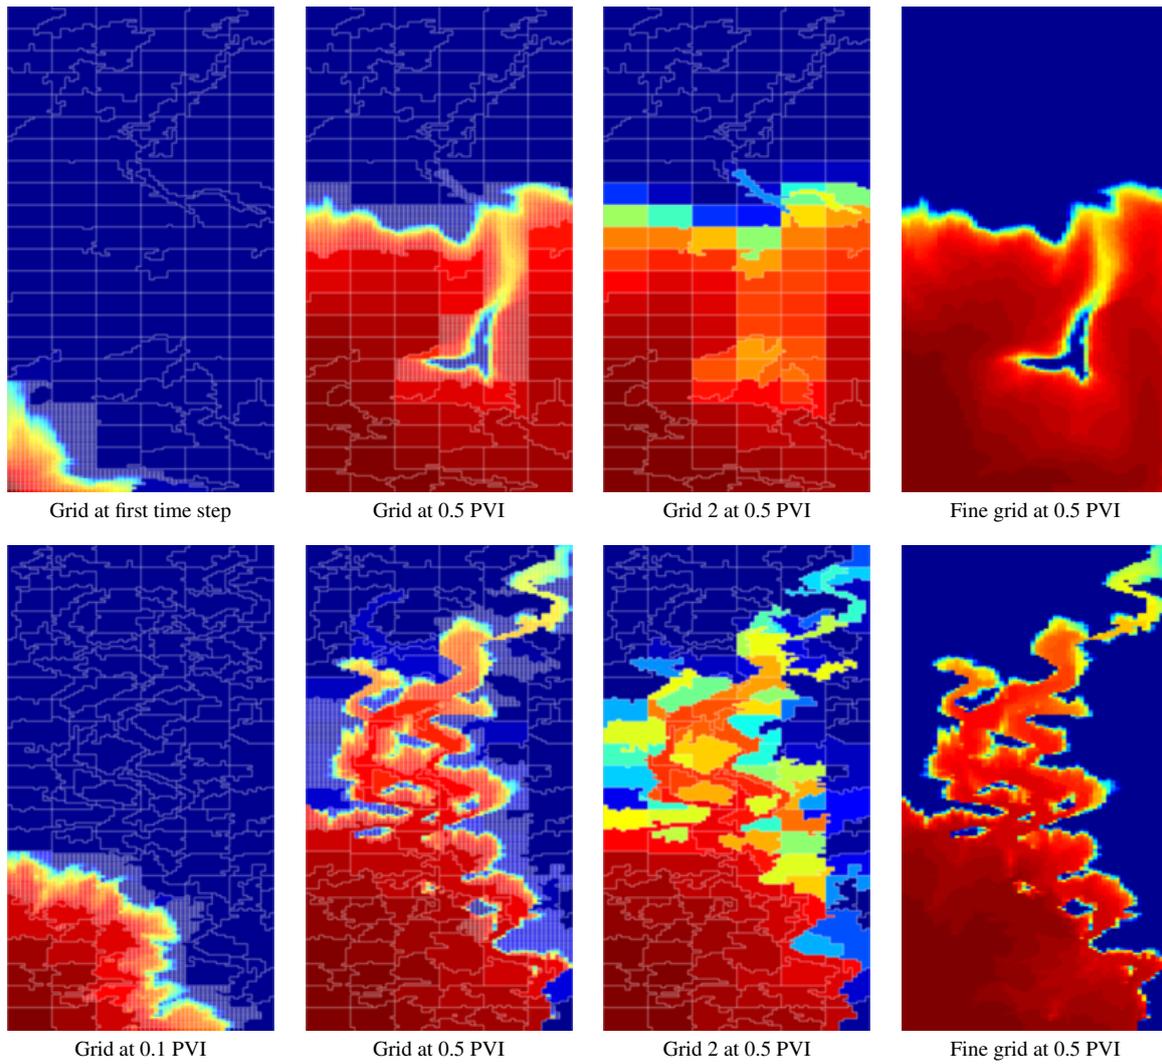
Because all grids considered herein are obtained by coarsening an underlying fine grid, it is relatively straightforward to add local refinement by manipulating the partition vector  $p$ , giving a local resolution that may be less or equal that of the fine grid. Moreover, this refinement can be added or removed dynamically provided we have good indicators of when to do so. Herein, we rely on the simplest approach possible, namely to compute each saturation step twice: once with a coarse resolution to estimate the movement of the front and once (locally) with higher resolution to resolve the movement more accurately. After the first step, we mark all blocks in which the saturation change from the previous time step exceeds a prescribe tolerance. These blocks are then refined. Likewise, after the second saturation step, we go through all refined blocks and mark those where the total fine-scale saturation changes are below another prescribed tolerance. In each marked block, the saturations are averaged back onto the original coarse block, and the refinement is removed by manipulating the partition vector.

**Example 5 (Layers 1 and 37 from SPE10)** *We revisit the two models studied in Example 4. In the fluid model we use quadratic relative permeabilities with a viscosity of 1 cP for the displacing fluid (water) and a viscosity of 0.2 cP for the displaced fluid (oil). This favourable displacement ratio will lead to a sharp front, for which local refinement of the grid is needed to avoid excessive numerical smearing when using a relatively coarse grid.*

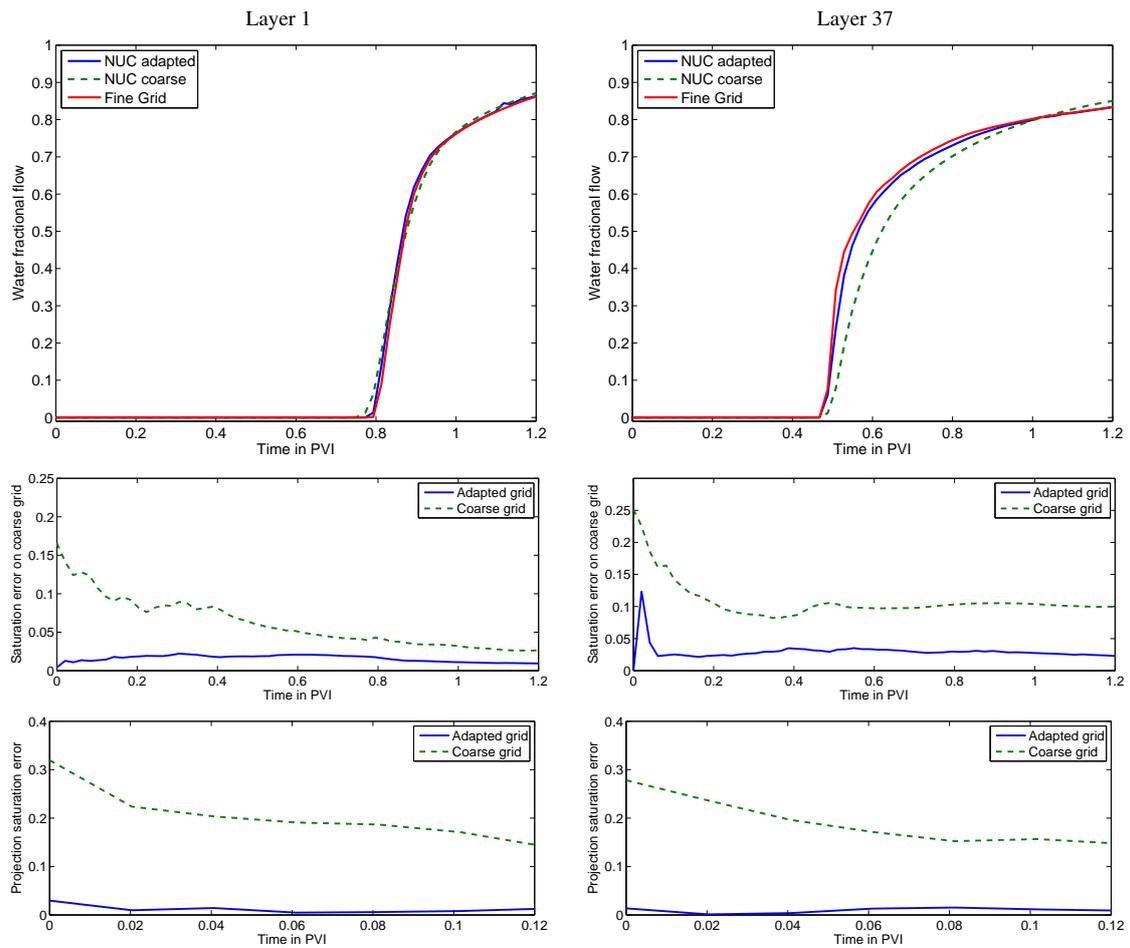
*Figure 7 shows the solution at dimensionless time 0.5 PVI computed on a coarse grid (Grid 2 from Figure 3), a coarse grid with local adaptive refinement, and on the original  $60 \times 220$  Cartesian grid. For Layer 1, we notice the excessive smearing at the fluid front, midway through the reservoir, but also that the coarse grid completely fails to capture the pocket of bypassed oil. The adaptive grid, on the other hand, is in good correspondence with the fine-scale reference solution. For Layer 37, most of the flow is confined to an intertwined pattern of narrow high-flow channels. Outside these channels, the static grid has relatively coarse blocks and once fluid enters these large low-flow blocks, the saturations get spread over a large area, causing excessive numerical diffusion.*

*Figure 8 shows the corresponding errors as a function of time. The coarse-scale saturation error is largest initially and decays toward water breakthrough; the qualitative behaviour of the fine-scale error is almost identical and the corresponding curves are therefore not reported. It may come as a surprise how well the static flow-based grid captures the water-cut curve for Layer 1, given the large initial error; but this result is in correspondence with previous observations (Aarnes et al., 2007; Krogstad et al., 2009) both for Cartesian and corner-point models. The most interesting result, however, is how much both the coarse-scale and projection errors are reduced by adding dynamical refinement.*

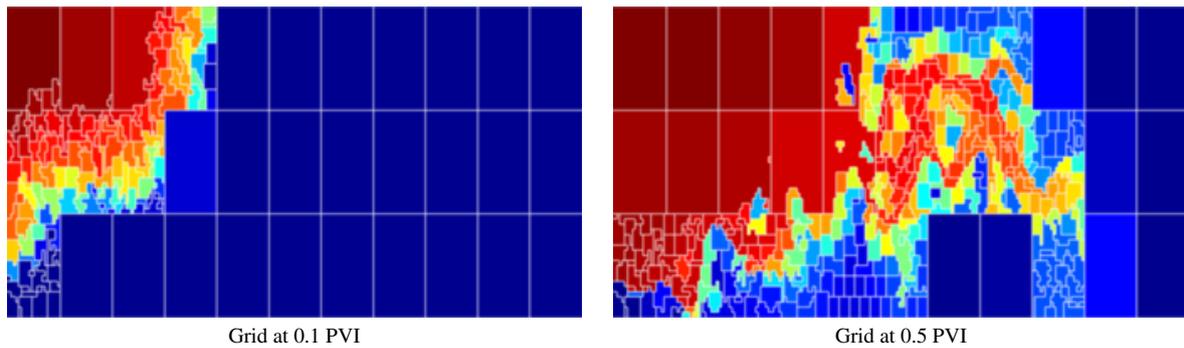
The previous example used a simple refinement approach in which all blocks marked for refinement were replaced with the underlying fine grid. We have also experimented with more advanced options, like adding an intermediate resolution and using flow-based coarsening with finer thresholds  $N_L$  and  $N_U$  in the refinement areas. Two examples of such grids are shown in Figure 9, where we (for illustration



**Figure 7** Examples of locally adapted grids for Layer 1 (top row) and Layer 37 (bottom row) of the SPE10 model.



**Figure 8** Water-cut curves and coarse-scale and projection errors as function of time for Layers 1 and 37 of the SPE10 model.



**Figure 9** Uniform Cartesian coarsening with adaptive flow-based refinement for Layer 37 of SPE10. The figures show the dynamic grid at dimensionless time 0.1 and 0.5 PVI, respectively.

purposes) have used a uniform Cartesian coarse grid and added flow-based refinement dynamically along the sharp displacement fronts. Likewise, one can relatively easily implement multilevel approaches. However, for the relatively simple water-flooding scenarios we have considered, these ideas have so far not been worth the (slight) increase in algorithmic complexity.

In the previous section, we briefly discussed the need for efficient communication between a multiscale flow solver and a transport solver working on a flow-adapted grid. Introducing local refinements would potentially reintroduce the need to communicate through the fine-scale grid to dynamically provide fine-scale fluxes in refined blocks. However, it has previously been demonstrated by Kippe et al. (2008) that to accurately capture the dynamic changes in the flow field, it is sufficient to update the multiscale basis functions only when a strong saturation front passes through a coarse block. Basis functions will therefore typically be updated in the regions where the transport grid is refined, and hence all the necessary fine-scale fluxes will be available. After some time steps, the strong fronts will have left the region, the coarse blocks have been reintroduced, and we can go back to use precomputed saturation mappings and sparse representation of fine-scale fluxes. For adverse mobility ratios, weak saturation fronts and smooth changes in the saturation implies that neither a dynamic refinement nor dynamic updates of basis functions are necessary.

### Concluding Remarks

In this paper, we have shown that flow-based coarsening is a versatile method to develop efficient transport solvers that can be used in combination with multiscale flow solvers. In particular, we have started with a method proposed by Aarnes et al. (2007) and shown how the partitioning computed by this method can be improved by including *a priori* information about block shapes or geological structures. Moreover, we observed that the sparsity pattern of the nonlinear discrete system preserves the causality of the underlying transport equations better when using net fluxes rather than bi-directional fine-scale fluxes across the interfaces in the coarse grid. Use of net fluxes may therefore significantly reduce the computational cost of transport solves, but it remains to see what the effect is on the accuracy. Finally, we have demonstrated how dynamical adaptivity easily can be included in the method to improve both the evolution and representation errors of strong saturation fronts. To simplify the presentation, most of the examples have focused on the widely-used SPE10 model. However, we have also included two examples that demonstrate that the ideas also apply to corner-point and PEBI grids.

### Acknowledgments

The research was funded in part by Shell Norge AS and the Research Council of Norway through grants no. 175962 and 186935. The authors thank Arne Skorstad (Norwegian Computing Center) for providing access to the SAIGUP model.

## References

- Aarnes, J.E., Hauge, V.L. and Efendiev, Y. [2007] Coarsening of three-dimensional structured and unstructured grids for subsurface flow. *Adv. Water Resour.*, **30**(11), 2177–2193.
- Aarnes, J.E., Kippe, V. and Lie, K.A. [2005] Mixed multiscale finite elements and streamline methods for reservoir simulation of large geomodels. *Adv. Water Resour.*, **28**(3), 257–271.
- Aarnes, J.E., Krogstad, S. and Lie, K.A. [2008] Multiscale mixed/mimetic methods on corner-point grids. *Comput. Geosci.*, **12**(3), 297–315, ISSN 1420-0597, doi:10.1007/s10596-007-9072-8.
- Christie, M.A. [1996] Upscaling for reservoir simulation. *J. Pet. Tech.*, **48**(11), 1004–1010, doi:10.2118/37324-MS.
- 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/>.
- Durlofsky, L.J., Jones, R.C. and Milliken, W.J. [October 1997] A nonuniform coarsening approach for the scale-up of displacement processes in heterogeneous porous media. *Adv. Water Resour.*, **20**, 335–347, doi: 10.1016/S0309-1708(96)00053-X.
- Durlofsky, L.J. [2003] Upscaling of geocellular models for reservoir flow simulation: A review of recent progress. Presented at 7th International Forum on Reservoir Simulation Bühl/Baden-Baden, Germany, June 23–27, 2003.
- Durlofsky, L.J. [2005] Upscaling and gridding of fine scale geological models for flow simulation. Presented at 8th International Forum on Reservoir Simulation Iles Borromees, Stresa, Italy, June 20–24, 2005.
- Efendiev, Y. and Hou, T.Y. [2009] *Multiscale Finite Element Methods*, vol. 4 of *Surveys and Tutorials in the Applied Mathematical Sciences*. Springer Verlag.
- Farmer, C.L. [2002] Upscaling: a review. *Int. J. Numer. Meth. Fluids*, **40**(1–2), 63–78, doi:10.1002/fld.267.
- Hauge, V.L., Lie, K.A. and Natvig, J.R. [2010] Grid coarsening based on agglomeration for multi-fidelity transport solvers. *submitted*.
- Hou, T.Y. and Wu, X.H. [1997] A multiscale finite element method for elliptic problems in composite materials and porous media. *J. Comput. Phys.*, **134**, 169–189.
- 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.
- 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.
- Krogstad, S., Hauge, V.L. and Gulbransen, A.F. [2009] Adjoint multiscale mixed finite elements. *SPE Reservoir Simulation Symposium, The Woodlands, TX, USA, 2–4 February 2009*, doi:10.2118/119112-MS.
- Lee, S., Zhou, H. and Tchelepi, H. [2009] Adaptive multiscale finite-volume method for nonlinear multiphase transport in heterogeneous formations. *J. Comput. Phys.*, **228**(24), 9036 – 9058, ISSN 0021-9991, doi: 10.1016/j.jcp.2009.09.009.
- Manzocchi, T. et al. [2008] Sensitivity of the impact of geological uncertainty on production from faulted and unfaulted shallow-marine oil reservoirs: objectives and methods. *Petrol. Geosci.*, **14**(1), 3–15.
- Natvig, J.R. and Lie, K.A. [2008a] Fast computation of multiphase flow in porous media by implicit discontinuous Galerkin schemes with optimal ordering of elements. *J. Comput. Phys.*, **227**(24), 10108–10124, ISSN 0021-9991, doi:10.1016/j.jcp.2008.08.024.
- Natvig, J.R. and Lie, K.A. [2008b] On efficient implicit upwind schemes. *Proceedings of ECMOR XI, Bergen, Norway, 8–11 September*, EAGE.
- Natvig, J.R. et al. [2009] Multiscale mimetic solvers for efficient streamline simulation of fractured reservoirs. *SPE Reservoir Simulation Symposium, The Woodlands, TX, USA, 2–4 February 2009*, doi:10.2118/119132-MS.
- Zhou, H., Lee, S.H. and Tchelepi, H.A. [2009] Multiscale finite volume formulation for the saturation equations. *SPE Reservoir Simulation Symposium, The Woodlands, TX, USA, 2–4 February 2009*, doi:10.2118/119183-MS.