INTEGRATION OF LOCAL-GLOBAL UPSCALING AND GRID ADAPTIVITY FOR SIMULATION OF SUBSURFACE FLOW IN HETEROGENEOUS FORMATIONS

M. GERRITSEN* AND J. V. LAMBERS†

Abstract. We propose a methodology, called multi-level local-global (MLLG) upscaling, for generating accurate upscaled models of permeabilities or transmissibilities for flow simulation on adapted grids in heterogeneous subsurface formations. The method generates an initial adapted grid based on the given fine-scale reservoir heterogeneity and potential flow paths. It then applies local-global (LG) upscaling for permeability or transmissibility [7], along with adaptivity, in an iterative manner. In each iteration of MLLG, the grid can be adapted where needed to reduce flow solver and upscaling errors. The adaptivity is controlled with a flow-based indicator. The iterative process is continued until consistency between the global solve on the adapted grid and the local solves is obtained. While each application of local-global upscaling is also an iterative process, this inner iteration generally takes only one or two iterations to converge. Furthermore, the number of outer iterations is bounded above, and hence the computational costs of this approach are low. We design a new flow-based weighting of transmissibility values in local-global upscaling that significantly improves the accuracy of LG and MLLG over traditional local transmissibility calculations.

For highly heterogeneous (e.g., channelized) systems the integration of grid adaptivity and local-global upscaling is shown to consistently provide more accurate coarse-scale models for global flow, relative to reference fine-scale results, than do existing upscaling techniques applied to uniform grids of similar densities. Another attractive property of the integration of upscaling and adaptivity is that process dependency is strongly reduced, that is, the approach computes accurate global flow results also for flows driven by boundary conditions different from the generic boundary conditions used to compute the upscaled parameters.

The method is demonstrated on Cartesian Cell-based Anisotropic Refinement (CCAR) grids, but can be applied to other adaptation strategies for structured grids, and extended to unstructured grids.

1. Introduction. The permeability of reservoir rocks is often highly variable and uncertain. However, it can greatly affect flow and transport in subsurface formations, and must therefore be adequately modeled. The effects of uncertain reservoir heterogeneity can be included either through stochastic modeling or by simulating a number of deterministic geostatistical realizations of the reservoir [9], which is the underlying assumption in this paper. Because the computational costs of solving for a large number of realizations are high, simulations are generally performed on grids that are coarse compared to the given geological models. The grid size is typically determined by computational considerations. These coarsened flow models should adequately represent key behaviors, such as the overall flow rate for given boundary conditions or critical connected flow paths.

In this paper, we focus on single-phase upscaling to test the effectiveness of our new approach. Assuming single-phase, steady and incompressible flow in a heterogeneous reservoir, the governing dimensionless equation, assumed valid on the fine geological scale, is obtained by combining the continuity equation $\nabla \cdot \mathbf{u} = \mathbf{0}$ and Darcy's law given by $\mathbf{u} = -\mathbf{K} \cdot \nabla p$, to get

$$(1.1) \qquad \nabla \cdot (\mathbf{k} \cdot \nabla p) = 0.$$

Here p is the pressure, $\mathbf{u}$ the Darcy velocity and $\mathbf{K}$ the permeability tensor, all of which are assumed non-dimensionalized by appropriate reference values. Although

*Department of Energy Resources Engineering, Stanford University, Green Earth Sciences Building, Stanford, CA 94305-2220, USA
†Department of Energy Resources Engineering, Stanford University, Green Earth Sciences Building, Stanford, CA 94305-2220, USA

homogenization of equations with variable coefficients generally yields coarse-scale equations of a different form than the original fine-scale equations, it is common practice to keep the same equation as (1.1) for the pressure on the coarse scale, with the permeability tensor replaced by the coarse-scale permeability tensor $\mathbf{K}^*$, also referred to as the effective permeability tensor, and p by the coarse-scale pressure. This has been found to lead to adequate coarse-scale pressure solutions [10, 27, 5]. We use Eqn. (1.1) indiscriminately of scale, and seek to find high quality representations of the global fine grid flow in the reservoir when the equation is solved on a coarse grid. The literature suggests three main approaches to control the quality of these representations: specialized upscaling, gridding and discretization strategies.

A large variety of single-phase upscaling techniques exist that are used to find representative permeability or transmissibility fields on the coarse grids. In purely local methods, the upscaled parameters are computed over fine-scale regions that correspond to the target coarse block under study. In extended local upscaling methods, regions neighboring the coarse block are included. In both approaches, the local flow is driven by generic Dirichlet, Neumann or periodic boundary conditions imposed on the fine-scale regions, see [10, 31, 27]. Although local methods may give satisfactory results for some reservoirs, they generally do not perform well when the spatial scales of important reservoir heterogeneities, such as high-permeability flow paths, are larger than the simulation grid, unless these grids are aligned with such paths. Global methods, which use global fine-scale flow simulations to determine coarse-scale parameters, can reduce errors in the presence of large connected flow paths, but they are generally expensive [19]. An attractive and relatively novel method to improve large-scale connectivity in upscaled permeabilities or transmissibilities is the local-global (LG) upscaling method [7]. Here, a standard extended local method is used to obtain an approximate coarse permeability or transmissibility field. Then, a coarse global solve is performed with generic boundary conditions. Interpolation of the coarse solution provides the Dirichlet boundary conditions for the next iteration of local solves. The process is repeated until the coarse field converges, which generally takes only a couple of iterations.

Gridding strategies include alignment of the grid with flow features and/or geological features, and local grid adaptation. Flow-aligned gridding, also referred to as flow-based gridding, in which grid cell boundaries are aligned with dominant flow paths, has been suggested by various authors, see, for example, [17] and references therein, to avoid smoothing of permeability contrasts straddled by the coarse grid. On such grids, local upscaling methods may lead to good representations of the reservoir heterogeneity. However, in strongly heterogeneous reservoirs, and especially in three dimensions, it is rather challenging to construct flow-based grids that are well aligned with the flow and have sufficient grid quality. An alternative is local grid adaptation strategies, in which the grid density is increased where needed to improve the representation of heterogeneity, for example, in regions with high permeability contrasts. Adaptive grids also allow important flow features with small spatial support, such as frontal regions, to be resolved accurately without sacrificing computational efficiency, see for example [29, 28, 14]. We think of adaptive grid strategies as part of the class of multi-scale methods, as discussed in, for example, [21] and references therein. Multi-scale methods, designed to capture multiple scales involved in the reservoir fluid flow processes, include approaches that upscale to a coarse simulation grid and reconstruct the solution on a finer scale within each coarse grid cell using a functional representation, but also approaches based on grid refinement. Whether representing the flow on

the fine-scale functionally or discretely, the success of the multi-scale method depends to a great extent on the chosen boundary conditions on the local fine regions.

Discretization strategies can also strongly impact the quality of the coarse solutions. With two-point flux approximations (TPFA), a flux through a given cell face is computed using pressure values from the cells that share the face, whereas with multi-point flux approximations (MPFA), additional cells may be involved. TPFA methods are attractive because of their simplicity, compactness and monotonicity. In many cases, the TPFA approximation is reasonable, but in cases with significant full-tensor effects, that may be introduced by the permeability field directly or by grid topology, larger errors should be expected, as assuming a TPFA is similar to neglecting the cross terms in upscaled permeability tensors. Interestingly, and perhaps surprisingly, TPFA methods combined with local-global procedures have been shown to provide good upscaled transmissibilities even in challenging cases [7, 6]. It is generally agreed that when strong full-tensor effects are present, MPFA methods are preferable. Many authors have investigated MPFA approaches for solving single-phase flow problems with full-tensor permeability that is assumed constant over grid cells [1, 12, 22]. Naturally, MPFA methods lead to larger stencils than TPFA. On structured grids in three dimensions, for example, the MPFA O-method developed by Aavatsmark [1] results in a 27-point stencil for the cell-centered pressure equation, as compared to a 7-point stencil for TPFA. MPFA methods have been found to suffer from non-monotonicity; that is, the pressure solution may contain non-physical oscillations [26]. Various new MPFA methods have been suggested that improve their compactness and robustness, including the L-method by Aavatsmark et al. [2], and the spatially varying multi-point flux upscaling method, VCMP [23]. Preliminary studies combining VCMP with a local-global approach are also reported in this study.

1.1. This work: Integrating upscaling and adaptivity.
The main questions we address in this work are

- how to combine local-global upscaling and local grid adaptivity to effectively represent large-scale heterogeneities on computational grids for single-phase flows,
- how to control the local grid adaptation during the upscaling process to improve the upscaling accuracy, and
- whether the proposed method reduces the process dependency that is generally introduced by upscaling methods.

While working on these research questions, we found that we can significantly improve local-global transmissibility upscaling by using a specialized flow-based weighting for local estimates of the transmissibilities for a given cell face. We will discuss this new weighting in this paper also.

We do not align the underlying computational grid with the flow, nor with the geology, instead focusing on the benefits that adaptivity and upscaling, specialized to adapted grids, can offer independently of any particular gridding strategy. For concreteness and simplicity, we apply the upscaling strategy to a cell-based anisotropic Cartesian refinement strategy. For similar reasons, we limit ourselves to TPFA methods. The proposed methodology is general enough, however, to be applicable to alternate gridding and discretization strategies. We focus on finite volume upscaling methods that lead to a cell-centered, coarse-scale pressure equation and conservative fluxes.

We continue the paper in Section 2 by describing Cartesian cell-based anisotropically defined (CCAR) grids and discretization of the single-phase pressure equation

on such grids. In Section 3, we present the general Multi-Level Local-Global (MLLG) upscaling strategy that we propose in this paper for permeability and transmissibility upscaling, and discuss in more detail several components of this algorithm. We present an indicator that we use to adapt the grid during the MLLG iteration to reduce flow solver and upscaling errors. The new transmissibility weighting is also described in this section. Numerical results are presented in Section 4. We apply our algorithm to a suite of two-dimensional test problems, including four channelized domains, and with flow driven in several directions. Although the numerical results are given in two dimensions only for ease of visualization, the method is fully three dimensional, see also [15]. Discussion and conclusions are presented in the last two sections.

2. Solving the Single-Phase Pressure Equation Using An Adaptive Cartesian Approach. In this section, we describe the strategies for gridding and discretization that we will use in conjunction with our upscaling method.

2.1. Cartesian Cell-based Anisotropically Refined (CCAR) grids. When designing an efficient, adaptive simulation approach, a grid topology must ideally facilitate fast transitions from dense to coarse grids without sacrificing grid and solver quality, allow efficient computation despite potentially high levels of refinement, and lead to fast grid generation. We have found a Cartesian grid topology with anisotropic cell-based refinement and coarsening quite attractive because it allows aggressive grid adaptation while maintaining grid quality. We use adaptivity, rather than flow-based gridding, to align the grid with important flow and geological features and assess its effectiveness in ensuring flow accuracy. The CCAR grid is formed by a number of anisotropic refinements of its cells; i.e., splitting a cell in half along a coordinate direction, to construct a suitably refined mesh. This is illustrated for three refinement levels in Figure 2.1.
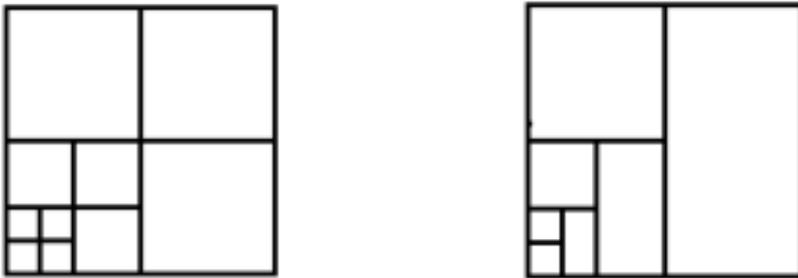

FIG. 2.1. *CCAR grid cells with 3 levels of refinement. Left: isotropic refinement. Right: anisotropic refinement.*

We impose the restriction that no cell can have more than four neighbors in each direction (or two neighbors in two dimensions) to maintain solution accuracy across interfaces [3]. The CCAR grids are stored using the unstructured data approach first suggested in [16], which leads to efficient look-ups. In this work, we denote cells in a CCAR grid, G, by $C_i, i = 1, \ldots, N_{cells}$, with N_{cells} the number of cells in G, and cell

faces by $F_i, i = 1, \ldots, N_{faces}$, with N_{faces} the total number of faces in G.

We use the compact finite volume discretization of the pressure equation 1.1 on CCAR grids developed in [15]. For homogeneous media, this discretization is formally second order. For heterogeneous media, the computation of the flux through a face with a hanging node requires careful consideration.
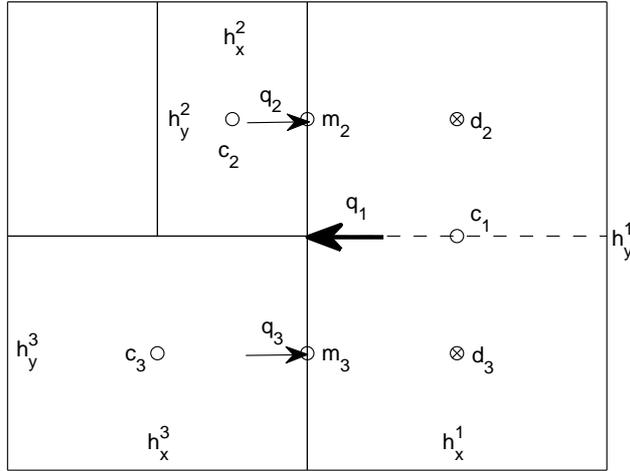

FIG. 2.2. *The points c_j, $j = 1, 2, 3$, d_j, $j = 2, 3$ and m_j, $j = 2, 3$, are used to compute the fluxes q_2 and q_3 across the half-faces between cell 1 and cell 2, and between cell 2 and cell 3.*

Consider Figure 2.2, showing three cells $C_i, i = 1, 2, 3$, with dimensions h_x^i and h_y^i, and diagonal permeability tensors with diagonal elements k_x^i and k_y^i. We will construct a flux for the face of cell 1 bordering cells 2 and 3. The area of the face of cell i that borders this face is denoted by A_i. The average pressures in the cells i are denoted by p_1, p_2 and p_3.

Requiring continuity of the flux across the half-faces, we obtain the equations

$$(2.1) \qquad \frac{k_x^j}{h_x^j}(p(m_j) - p(c_j)) = \frac{k_x^1}{h_x^1}(p(d_j) - p(m_j)), \quad j = 2, 3,$$

where the points m_j, c_j and d_j are as given in the figure, and the corresponding fluxes

$$(2.2) \qquad q_j = \frac{2h_y^j}{\frac{h_x^1}{k_x^1} + \frac{h_x^j}{k_x^j}}(p(d_j) - p(c_j)), \quad j = 2, 3.$$

5

We represent the pressure at the auxiliary points d_2 and d_3 in terms of the available pressures at nearby cell centers. To that end, we use a convex combination of two interpolations, one with a vertical orientation and one with a horizontal orientation, which are illustrated in Figure 2.3. For the vertically-oriented interpolation in d_2, we use bilinear interpolation at c_1, and the two closest cell centers besides c_2 and c_3, that form a triangle containing d_2. For the horizontally-oriented interpolation of d_2, we use bilinear interpolation at c_1, c_2, and the closest cell center that completes a triangle containing d_2. Similar triangles are constructed for d_3. We combine the horizontal and vertical interpolations with a weighting determined by the angle of the pressure gradient through the point d_i, $i = 2$ or 3, at which we are interpolating. This introduces a natural flow-based interpolation, and the resulting, wider, stencil at hanging nodes is reminiscent of the L-scheme suggested in [2].

FIG. 2.3. *Triangulations used to represent the pressure at d_2 and d_3 in terms of cell centers. The values at these points are obtained by bilinear interpolation of the values at the vertices of the horizontally, or vertically, oriented triangles containing them.*

The resulting matrix remains sparse with a reasonable condition number thanks to the compact finite volume discretization. Because we do not have a strict TPFA discretization at hanging nodes, the matrix is not guaranteed to be an M-matrix and hence positivity is not ensured, but the number of matrix elements with the wrong signs is typically very small. Solution of the system using an algebraic multigrid solver

(BoomerAMG [18]) is fast with almost linear convergence properties for our test cases.

3. Effective upscaling for CCAR grids. As mentioned in the introduction, it has been shown that for Cartesian grids, the local-global (LG) method leads to improved results as compared to local, or extended local, methods [7]. For each coarse grid cell, effective permeabilities or face transmissibilities are computed by solving a local flow problem on the fine geo-cellular grid directly around the coarse cell. This local flow problem is driven by boundary conditions interpolated from a coarse global solve. The resulting coupling between the local flow problems and the coarse global solve leads to an iterative scheme that generally converges rapidly. Because of its effectiveness, LG upscaling serves as the cornerstone of the local-global upscaling algorithm for CCAR grids, which we now present.
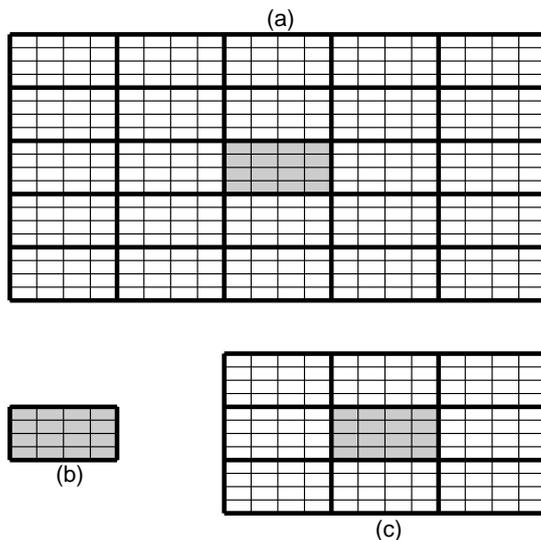
Fig. 3.1. *Schematic showing (a) fine- and coarse-scale grids, (b) fine-scale local region and (c) fine-scale extended local region, with $r = 1$ for permeability upscaling. For transmissibility upscaling, the fine-scale local region is centered around a cell-face rather than the cell itself.*

3.1. The general Multi-Level Local-Global algorithm. In permeability upscaling, we associate with each cell $C_i, i = 1$, in the CCAR grid G an extended region $E_i = C_i \cup B_i$, where B_i is a border region of constant width r around C_i. In transmissibility upscaling the (extended) local region is centered around a coarse cell face F_i, rather than the cell itself. In each region E_i we discretize the pressure equation on the uniform geocellular grid, which here we assume to be Cartesian. We note that in the local-global upscaling method in [7] the width of each dimension of the border region is chosen to be a multiple of the dimension of the corresponding grid cell instead. This model works well on a uniform grid, but in a CCAR grid, the coarsest cells do not such a large border because the fine-scale permeability within such cells tends to

be smoother relative to the coarse grid scale. In permeability upscaling our border region scaling is also attractive for efficiency reasons: if a cell C_i with corresponding extended region E_i should be bisected during refinement, the resulting new cells will have extended regions whose union coincides with E_i. This allows pressure values on the boundary of E_i to be re-used for the new regions.

The general MLLG algorithm is given as:

Construct the base CCAR grid G using static refinement indicators
while grid adaptation not completed **do**
    while not converged **do**
        for each cell C_i (permeability upscaling) or each face in G **do**
            Select boundary conditions on E_i:
                if this is the first iteration **then**
                    Use generic boundary conditions: constant pressure in one
                        coordinate direction, no-flow in the other coordinate direction
                else
                    Use bilinear interpolation of global solutions p_x and p_y to
                        approximate pressure on the boundary of E_i
                    Use Dirichlet boundary conditions using these pressure values
                end
            Solve the pressure equation on E_i with selected boundary conditions
            Compute fluxes between the fine grid cells in E_i and pressure values at their cell
                centers
            Use these fluxes and pressure values on E_i to compute new diagonal permeability
                tensor K_i for C_i , or a new transmissibility for F_i
        end
        Solve generic flow problems on G to obtain pressure fields p_x and p_y
        Check for convergence
    done
Check if any of the cells in G must be refined
done

The first step in the MLLG algorithm is the construction of an initial base CCAR grid. This base grid is designed with the aim of preserving important high permeability flow paths. After the base grid is generated, we apply local-global upscaling, as in the Cartesian case, to find representative permeability values on the cells C_i of the CCAR grid, or transmissibility values for the faces F_i of the CCAR grid. Once the local-global step has converged, which generally only takes one or two iterations, we check if further grid refinement is necessary, and repeat the local-global iteration on the updated grid. Thus, the MLLG algorithm has an inner iteration, which is a single application of local-global upscaling for CCAR grids, and an outer iteration with the number of outer iteration steps bounded above by the number of levels of refinement between the coarsest scale of the base grid, and the scale of the fine grid.

In the next subsection we discuss the main components of the MLLG algorithm in more detail.

3.2. Main Components of the MLLG Algorithm.

3.2.1. Creation of the Base CCAR Grid. To create a base CCAR grid, we apply a static refinement indicator based on the method proposed by Younis and Caers in [32]. The method is based on a "leaky" connected-set approach that
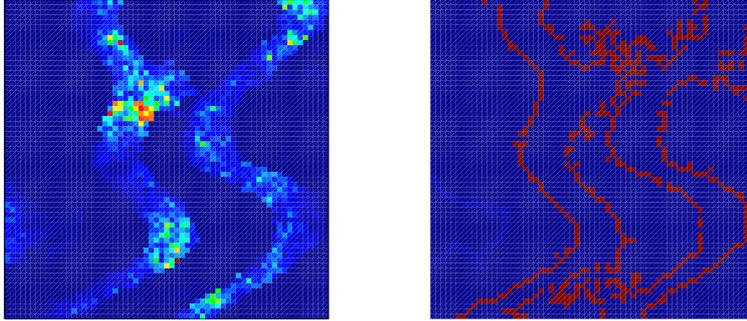
FIG. 3.2. *Permeability field (left), and boundaries of connected flow paths detected by leaky connected-set approach (right)*

delineates clusters of fine cells with similar permeability. Heuristic measures are used to color each connected set, indicating the level of the importance of capturing the set by refinement. Refinement flags are then obtained by explicit boundary detection. An illustration of the process appears in Figure 3.2. In the current implementation, a minimum level of refinement is set based on the resolution of the geocellular grid and efficiency considerations, and a CCAR grid G is generated by iteration until these level constraints are met. Such a grid is illustrated in Figure 3.3.

In our experience, this approach has given us reasonable CCAR grids to start our MLLG process. We note that the grid is further adapted during MLLG.

3.2.2. Boundary Data for Local Solves. For each cell C_i of the CCAR grid G, we must solve local flow problems (in two dimensions) over the corresponding extended region E_i. In order to obtain the Dirichlet boundary data for each such problem, we use bilinear interpolation, with the interpolation points obtained from each of the global coarse-scale solutions computed during the previous iteration of local-global upscaling. While it is possible to use a higher-order interpolation scheme, it is not advisable to do so because the interpolant would exhibit oscillatory behavior, as opposed to the monotonicity we would expect from a pressure field.

We illustrate the interpolation for permeability upscaling. In Figure 3.4, a cell C_i is shown, along with its extended region E_i. For each of the points p_{ij}, $j = 1, 2, \ldots, 8$ along the boundary of E_i, we find the three cells of G whose centers are as close to p_{ij} as possible, and form a triangle containing p_{ij}. We then use barycentric coordinates [30] to approximate the pressure at p_{ij}.

3.2.3. Local Flow-based Refinement of the Base Grid. The fluxes across a face, from each global pressure solve, can be used to determine whether cells bordering the face should be refined.

Specifically, let F_i be a face with area A_i, and let q_i be the flux across the face obtained by a pressure solve on the CCAR grid with flow driven in the direction perpendicular to the face. If Q is the global flow, and A is the area of the face of the global flow domain perpendicular to the direction of the flow, then we determine that the cells sharing the face should be refined if

(3.1) $$|q_i| > \alpha \frac{A_i}{A} |Q|,$$

9

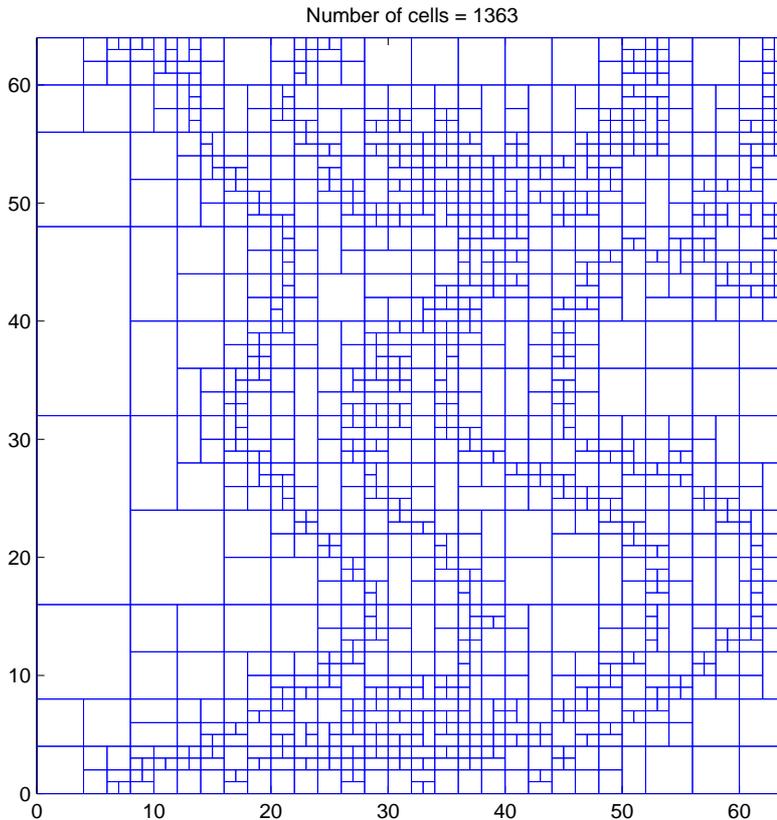Number of cells = 1363

FIG. 3.3. *base CCAR grid generated from permeability field and connected sets from Figure 3.2*

where α is a proportionality constant that depends on the ratio of areas, i.e. $\alpha = \alpha\left(\frac{A_i}{A}\right)$. The term $A_i|Q|/A$ can be interpreted as the face's "fair share" of the global flow.

To choose $\alpha(A_i/A) > 1$, we use the following heuristics. First, when A_i is the area of a face of a cell at the coarsest scale in the grid, then $\alpha(A_i/A)$ should be approximately 1, so that refinement will always take place except in areas that are receiving less than their fair share of flow. This low tolerance is necessary to avoid overlooking high-flow paths that occur within large regions that otherwise have very low flow. On the other hand, as A_i approaches the size of a face in the fine grid, then $\alpha(A_i/A)$ should be approximately A/A_i, so that refinement will not take place unless there is very high flow; that is, nearly all of the global flow is crossing the face F_i. This progression from low to high tolerance is used to avoid superfluous refinement while not missing any finer-scale regions of high flow. Under the assumption that $\alpha(A_i/A) = C(A_i/A)^{-p}$ for some positive real number p, we obtain

$$\alpha(A_i/A) = \frac{\delta}{(A/A_c)^p}\left(\frac{A_i}{A}\right)^{-p},$$
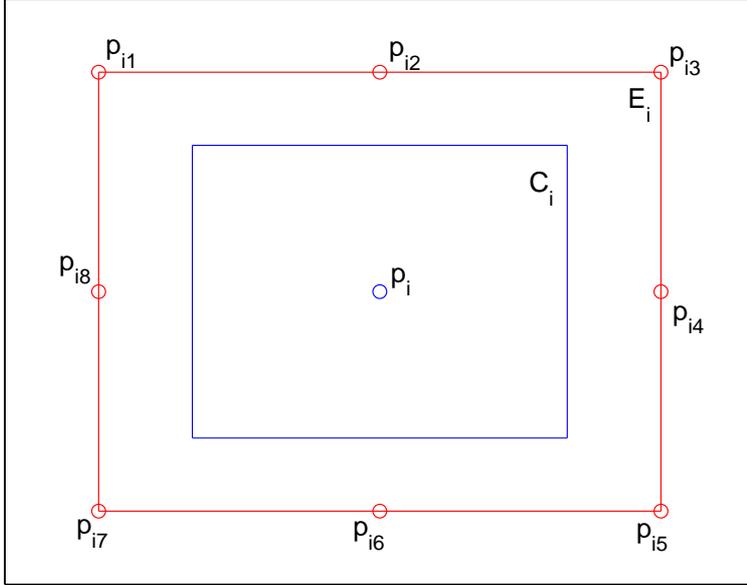
FIG. 3.4. *points* p_{ij}, $j = 1, 2, \ldots, 8$ *on the boundary of an extended region* E_i, *shown for permeability upscaling, at which global coarse-scale pressure is interpolated. The region* E_i *corresponds to the cell* C_i *with center* p_i.

where A_c is the area of the largest face in the grid parallel to F_i, $\delta \approx 1$ is the chosen value of $\alpha(A_c/A)$, and p is given by

$$p = \frac{\ln \frac{A}{\delta A_f}}{\ln \frac{A_c}{A_f}},$$

where A_f is the area of each fine grid face that is parallel to F_i.

If the face F_i borders only two cells, then both are refined. Otherwise, the face borders two hanging cells and one non-hanging cell; only the non-hanging cell is refined (e.g., cell 1 in Figure 2.2). All refinement is in the direction that is orthogonal to the face.

We can use additional refinement criteria to reduce the number of "hanging" nodes that occur in high-flow regions. First, we apply the above flow-based criterion to each "half-face" that is shared by a hanging cell and a non-hanging cell. If the flow across this half-face is more than its fair share, then the non-hanging cell is refined along the direction that is orthogonal to the face. This refinement also takes place if too much of the flow across the entire face (75%, in practice) is through one hanging node.

We note that in transmissibility upscaling, we also refine if the estimated transmissibility for a face is negative. This is further discussed in the next subsection.

3.2.4. A New Local Gradient-based Weighting of Transmissibilities. Referring back to Figure 2.2 we choose the extended region E_i centered at the face

F_i for which we compute the flux to be large enough to contain the cells centered at the points c_1, c_2 and c_3. We solve the pressure equation on E_i with Dirichlet boundary conditions as in permeability upscaling. The computation of the fluxes given by equation 2.2 from the local solution, along with the approximate pressure differences $\Delta p_j = p(d_j) - p(c_j)$ obtained by averaging the local pressure field over the appropriate regions, can be used to compute the transmissibilities T_2 and T_3 corresponding to the fluxes q_1 and q_2 indicated in Figure 2.2. These transmissibilities are then used in the stencils centered at c_2 and c_3. In the stencil centered at c_1, we use the transmissibility $T_1 = T_2 + T_3$.

The original local-global upscaling algorithm in [7] computed upscaled transmissibility across a face by considering only data obtained from a global flow driven by a pressure difference in the direction perpendicular to the face. But, since we are using two generic flow problems, we actually obtain two fluxes for each face. These fluxes should indeed not have equal weight: a flow that is nearly parallel to the face should not contribute much, if anything, to the computed face transmissibility. If it does, this may result in a stronger process dependency, which we would like to avoid. Rather than choosing only one flow, we designed a new weighting process for the transmissibilities, which we now describe.

Consider the simple case of a face shared by only two cells. Let $q^{(1)}$ and $q^{(2)}$ denote the fluxes across the face obtained from the two local pressure solves, obtained either in an extended local or a local-global method, and let $\Delta p^{(1)}$ and $\Delta p^{(2)}$ denote the differences in the cell averages of the pressure from these solves. From these values, we obtain two approximations for the transmissibility,

$$(3.2) \qquad\qquad T^{(j)} = -\frac{q^{(j)}}{\Delta p^{(j)}}, j = 1, 2.$$

Let $\mathbf{n}$ be a unit vector that is perpendicular to the face. We find a unique transmissibility T using weighted least squares, with the weights determined by the average pressure gradients near the face, $\nabla p^{(1)}$ and $\nabla p^{(2)}$, where the averages may be cell or line averages (it remains to be determined which approach is best). This approach yields the system $WAT = W\mathbf{b}$, where

$$W = \begin{bmatrix} \omega^{(1)} & 0 \\ 0 & \omega^{(2)} \end{bmatrix}, \quad A = \begin{bmatrix} \Delta p^{(1)} \\ \Delta p^{(2)} \end{bmatrix}, \quad \mathbf{b} = -\begin{bmatrix} q^{(1)} \\ q^{(2)} \end{bmatrix},$$

and

$$\omega^{(i)} = \frac{|\nabla p^{(i)} \cdot \mathbf{n}|}{\|\nabla p^{(i)}\|}, \quad i = 1, 2.$$

The rationale for choosing the weights $\omega^{(i)}$ in this manner is to determine, in an average sense, the orientation of the flow from the two local solves. If either of these flows is determined to be nearly parallel to the face, then it is not be assigned much significance in the computation of the transmissibility across the face. On the other hand, if the flow is nearly perpendicular to the face, then the ratio of the negative flux to the pressure difference is believed to represent the coarse-scale transmissibility accurately.

The weighted least squares problem has the solution

$$T = \frac{A^T W^T W \mathbf{b}}{A^T W^T W A} = -\frac{[\omega^{(1)}]^2 q^{(1)} \Delta p^{(1)} + [\omega^{(2)}]^2 q^{(2)} \Delta p^{(2)}}{[\omega^{(1)}]^2 [\Delta p^{(1)}]^2 + [\omega^{(2)}]^2 [\Delta p^{(2)}]^2}.$$

In the case of a face adjacent to hanging nodes, we apply this same approach to compute the transmissibilities across the half-faces. If either of the equations in (3.2) would yield a negative transmissibility, then it is not included in the computation of T. If both equations would yield negative transmissibilities, then we set T equal to the sum of the fine-scale transmissibilities along the face, and flag the cells sharing the face for refinement.

We are able to prevent most instances of negative transmissibility from arising in the first place by computing flux across each face as follows. Assume that a face F_i in the coarse grid consists of n fine-scale faces f_j, $j = 1, \ldots, n$. We compute the flux using

$$q_i = \sum_{j=1}^{n} T_j \Delta p_j,$$

where, for each j, T_j is the transmissibility across f_j, and Δp_j is the pressure difference across the face. To compute this pressure difference, we first use the pressure values in the two fine-scale cells c_j and d_j bordering f_j. Then, we compute a pressure difference $\Delta p'_j$ using the pressure values in the two cells c'_j and d'_j that are one cell removed from f_j along the direction orthogonal to f_j. If Δp_j and $\Delta p'_j$ have opposite sign, then we assume that Δp_j is an isolated fluctuation and compute the flux across f_j using $\frac{1}{3}\Delta p'_j$; otherwise we use Δp_j. The cells involved in the computation are shown in Figure 3.5.
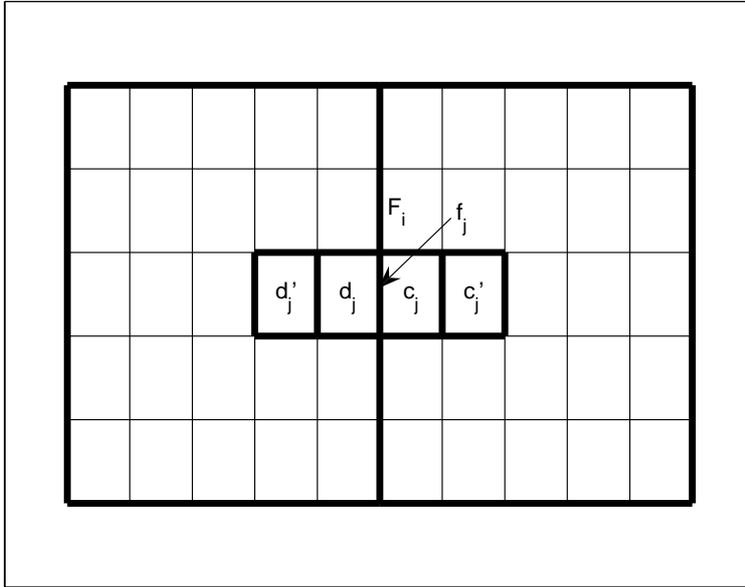

FIG. 3.5. *The cells c_j, d_j, c'_j and d'_j used to compute the pressure differences Δp_j and $\Delta p'_j$ across the face f_j that is a portion of the face F_i, the flux across which is used to compute the upscaled transmissibility across F_i.*

4. Numerical Results. In this section, we demonstrate the effectiveness of our upscaling method on a variety of permeability fields and boundary conditions.

Abbreviation	Description
EL	Extended local, on uniform Cartesian grid
EL-W	EL with new transmissibility weighting
LG	Local-global, on uniform Cartesian grid
LG-W	Local-global with new transmissibility weighting, on uniform Cartesian grid
EL-CCAR	Extended local, on CCAR grid
MLLG	Multi-level local-global, on CCAR grid
MLLG -W	Multi-level local-global with new transmissibility weighting, on CCAR grid

TABLE 4.1
Abbreviations used to refer to upscaling algorithms in Section 4.3.

4.1. The Test Suite. Our test suite consists of three types of fine permeability fields that are used in the experiments described in this section. All have dimension $L_x \times L_y$, where $L_x = 256$ and $L_y = 64$, and uniform grid step size 1 in both spatial dimensions. The fields are:

- 10 realizations with long correlation lengths of $\ell_x = 50$, $\ell_y = 10$, aligned with the grid.
- 10 realizations with long correlation lengths of $\ell_x = 75$, $\ell_y = 5$, not aligned with the grid
- Four channelized domains, given by layers 44, 48, 76 and 80 of the SPE10 Comparative Project (see [8]) that are modified to fit the domain dimensions $L_x = 256$ and $L_y = 64$.

Examples of realizations for the permeability fields with long correlation lengths are shown in Figure 4.1. For these cases, which are generated using 2-point statistics, we use mean 3.0 and standard deviation 1.735, which corresponds to a log-normal distribution. The fine-scale permeability fields of the channelized domains are shown in Figure 4.2.

The main purpose of the realizations with long correlation lengths is to show the effectiveness of the new transmissibility upscaling obtained with the new weighting scheme. The results also show the accuracy of local-global upscaling vs. extended local upscaling, but we refer to [7] for more extensive comparisons. For these cases, we use the upscaled transmissibilities to solve the pressure equation on a Cartesian grid with 32 cells in the x-direction, and 16 cells in the y-direction. The dimensions of the cells are $\Delta x = 8$ and $\Delta y = 4$.

The channelized cases are used to assess the effectiveness of MLLG and MLLG-W. Channelized reservoirs are generally difficult to upscale accurately because of the strong and large-scale flow paths. For the channelized cases, we use a CCAR grid obtained by iterative refinement of an initial 8×4 uniform Cartesian grid, with $\Delta x = 32$ and $\Delta y = 16$. The refinement is performed according to the criteria described in Section 3.2.3.

For each permeability field and each grid, we solve the pressure equation with the following six sets of boundary conditions, in order to measure how effective each algorithm is in reducing process dependency resolving the fine-scale velocity:

- Constant pressure/no-flow boundary conditions in the x- and y-directions, respectively. For the first set, we prescribe $p(0,y) = 1$, $p(L_x,y) = 0$, and
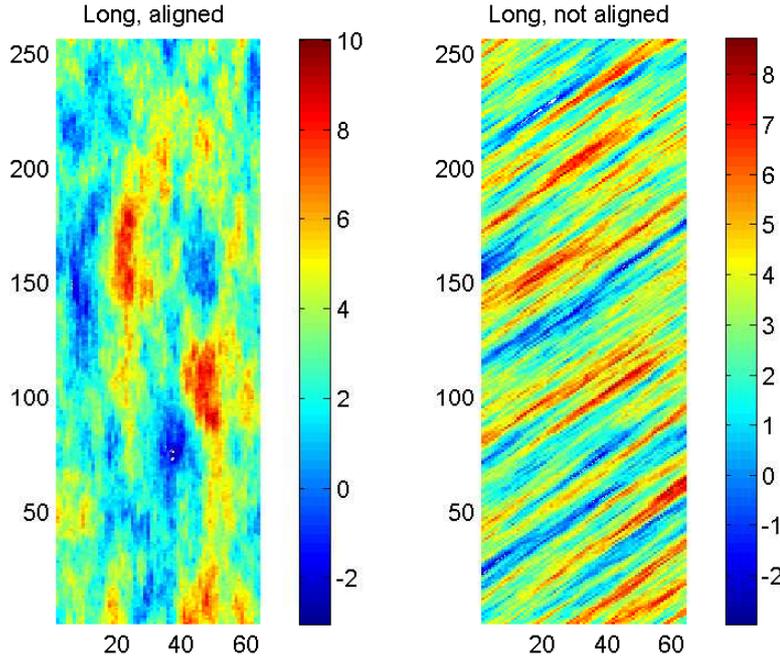
FIG. 4.1. *Fine-scale permeability fields used as test cases in Section 4.1. Left: long correlation length, aligned to grid. Right: long correlation length, not aligned to grid. The grids contain 256×64 cells, with $\Delta x = \Delta y = 1$.*
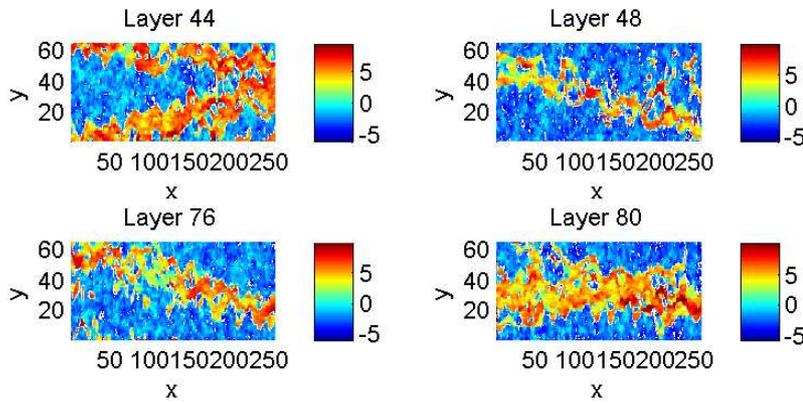

FIG. 4.2. *Fine-scale permeability fields for modified version of SPE10 layers 44, 49, 76 and 80. The grids contain 256×64 cells, with $\Delta x = \Delta y = 1$.*

$p_y(x, 0) = p_y(x, L_y) = 0$. For the second set, we prescribe $p(x, 0) = 1$, $p(x, L_y) = 0$, and $p_x(0, y) = p_x(L_x, y) = 0$. These are the generic boundary conditions used for the global coarse solves in local-global upscaling, and also for the local fine solves in local or extended-local upscaling.

• Flow from each of the four corners of the domain. For example, for flow

15

from the lower left corner, we prescribe $p(0, y) = \exp(-(2y/L_y)^4)$ (shown in Figure 4.3), $p(x, 0) = \exp(-(2(L_x - x)/L_x)^4)$, and $p(L_x, y) = p(x, L_y) = 0$. The boundary conditions for the other three corners are defined similarly.
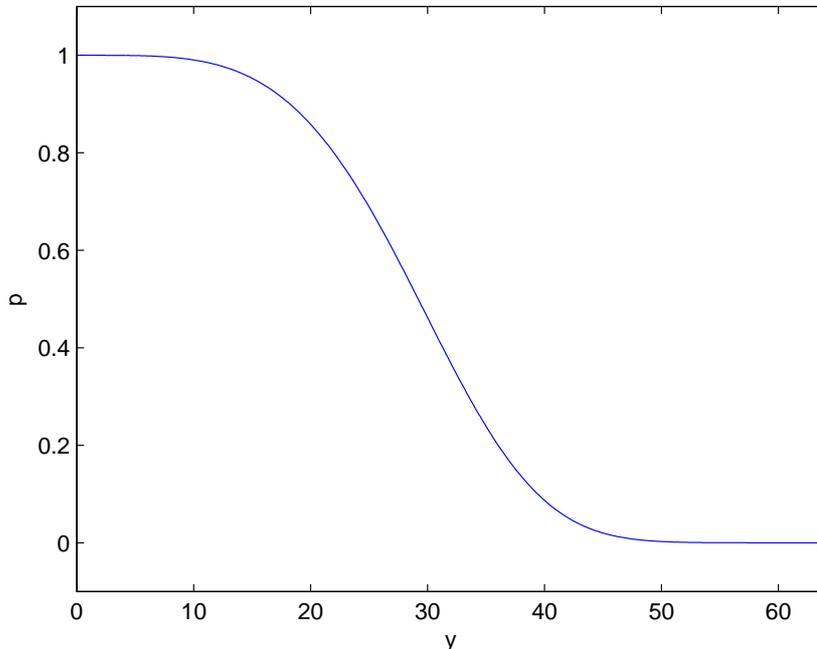

FIG. 4.3. *The pressure $p(0, y) = \exp(-(2y/L_y)^4)$ prescribed on the left boundary, for the set of boundary conditions in which flow emanates from the lower left corner of the domain.*

We will compare the total flow to that obtained from a fine-scale solution, for each set of boundary conditions, and for each upscaling method. For cases generated from 2-point statistics, we present the mean and standard deviation of relative errors in total flow across all realizations. In selected cases, we will also examine the velocity fields, compared to the fine-scale velocity field averaged over each coarse grid cell.

4.2. Non-channelized cases. We first report results for the test cases involving permeability fields with long correlation lengths.

4.2.1. Long Correlation Length, Aligned with Grid. The relative errors in total flow are given in Table 4.2 for the 10 realizations. Clearly, both methods outperform EL. This is expected, and in agreement with what is reported in [7], as EL cannot accurately represent permeability contrasts at a scale larger than the coarse grid size. We see that LG and LG-W are quite comparable, due to the fact that the gradient-based weighting suggested in Section 3.2.4 is designed to behave like LG in regions where the flow is oriented with the grid, which is the case with these realizations.

4.2.2. Long Correlation Length, Not Aligned with Grid. Table 4.3 lists the relative errors in the total flow for the 10 realizations with long correlation lengths $\ell_x = 75$ and $\ell_y = 5$, not aligned to the grid. This time, both EL and LG perform

	Fine		EL		LG		LG-W	
Flow	Mean	Std Dev	Mean	Std Dev	Mean	Std Dev	Mean	Std Dev
In x-direction	53.97	6.11	1.36	0.31	0.45	0.36	0.44	0.39
In y-direction	53.97	6.11	0.50	0.19	0.58	0.22	0.63	0.19
From lower left	18.83	6.00	1.00	0.56	1.17	0.71	1.16	0.78
From upper left	−19.62	7.67	1.58	0.72	1.55	0.84	1.50	0.73
From lower right	16.97	5.69	1.39	0.69	1.09	0.66	1.16	0.62
From upper right	−17.63	6.36	0.80	0.42	0.80	0.46	0.85	0.40

TABLE 4.2

Mean percentage errors in total flow, and corresponding standard deviations, after applying various upscaling methods to ten realizations with long correlation length, aligned with the grid, using a 32×16 coarse grid.

poorly. LG-W has the best overall performance, and is quite accurate for all boundary conditions, except for the low-flow case in which flow is driven in the x-direction. Such sensitivity of the upscaling accuracy in weaker flow cases is fairly typical for TPFA methods on permeability fields that are not grid-aligned. In such cases, it is indeed preferable to go to MPFA methods, see, for example [23].

	EL		LG		LG-W	
Flow	Mean	Std Dev	Mean	Std Dev	Mean	Std Dev
In x-direction	6.38	4.58	22.10	5.26	29.82	5.91
In y-direction	14.94	5.95	6.35	3.72	3.15	1.74
From lower left	26.84	7.13	16.07	4.52	4.53	2.94
From upper left	18.94	6.73	9.59	3.88	3.63	2.77
From lower right	12.99	4.76	6.03	3.38	3.27	4.02
From upper right	26.58	5.95	15.52	4.84	4.69	2.21

TABLE 4.3

Mean percentage errors in total flow, and corresponding standard deviations, after applying various upscaling methods to ten realizations with long correlation length, not aligned with the grid, using a 32×16 coarse grid.

Figure 4.4 shows portions of the velocity field computed by EL and LG-W. The portions chosen are those with the strongest flow. It can be seen from these figures that LG-W captures both the direction and magnitude of the fine-scale velocity vectors more accurately than EL, which yields velocity vectors that, locally tend to have a fixed orientation and not change direction as rapidly as the field generated by LG-W.

4.3. Channelized cases. Table 4.4 summarizes the results for layer 44 (see Figure 4.2) for the global flow, comparing multi-level local-global upscaling with the new transmissibility weighting (MLLG-W) to extended local upscaling on a Cartesian grid (EL), local-global upscaling on a Cartesian grid (LG), local-global upscaling with the new weighting on a Cartesian grid, (LG-W), and multi-level local-global upscaling without the new weighting (MLLG). We see that in all cases, the new transmissibility weighting improves the results significantly. Not only are the errors in total flow reduced strongly, process dependency is also greatly reduced. MLLG yields comparable accuracy to the Cartesian methods, while using a CCAR grid that contains significantly fewer cells due to coarsening in low-flow regions.

To demonstrate that both local-global upscaling and adaptivity contribute to accurate resolution of flow, we compare our algorithm to a modification of the algorithm
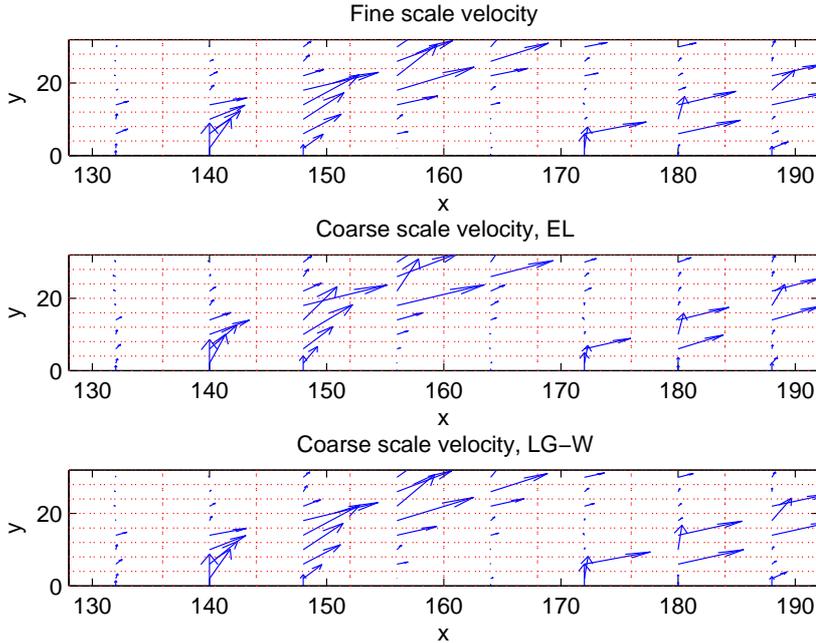
FIG. 4.4. *Velocity fields for permeability field with long correlation length, not aligned with the grid, with flow in the y-direction. Top plot: fine-scale velocity, averaged over each cell of a 32 × 16 coarse grid. Bottom plot: coarse-scale velocity, obtained using method LG.*

Flow	Fine-scale total flow	EL	LG	LG-W	MLLG	MLLG-W
In x-direction	7.8544	13.68	3.55	3.61	0.09	3.59
In y-direction	35.7709	7.20	2.68	3.57	2.31	3.51
From lower left	7.4063	8.51	1.78	0.44	1.57	1.46
From upper left	-1.346	0.96	12.53	0.62	20.31	1.73
From lower right	21.0316	3.95	3.97	6.54	3.40	5.93
From upper right	-28.0695	4.86	3.27	4.64	2.60	4.26

TABLE 4.4

Fine-scale total flow (first column) and percentage errors in total flow for modified SPE10 Layer 44 for various upscaling methods and boundary conditions. The EL, LG and LG-W methods use a 32 × 16 Cartesian grid, MLLG uses a 365-cell CCAR grid, and MLLG-W a 371-cell CCAR grid. The difference in these CCAR grids is due to the differences in transmissibility calculations, thus differences in local flow calculations, which impact the refinement criteria.

in which extended local upscaling is used for each CCAR grid, as opposed to using extended local only for the first inner iteration with the initial grid. Table 4.5 lists global flow errors for both algorithms. In all but two cases, local-global upscaling yields significantly more accuracy, even with fewer cells.

In Figure 4.5, we demonstrate the convergence of the global flow as the grid is refined, using the criteria defined in Section 3.2.3. We can see that for both generic global flow problems, with flow driven by pressure differences in the x-direction and

Flow direction	% error in global flow, extended local, 353-cell CCAR grid	% error in global flow, local-global, 371-cell CCAR grid
In x-direction	17.39	3.59
In y-direction	7.18	3.51
From lower left	8.25	1.46
From upper left	1.22	1.73
From lower right	5.31	5.93
From upper right	6.05	4.26

TABLE 4.5

Total flow error from applying extended local upscaling and MLLG on CCAR grids, with various boundary conditions, to the modified SPE10 layer 44

y-direction respectively, convergence to the fine-scale global flow is obtained after just two passes of refinement. Furthermore, a subsequent local-global iteration on the final base grid converges immediately to the desired accuracy. The grids from the iterative refinement process are shown in Figure 4.6. Similar results are obtained for the other layers.
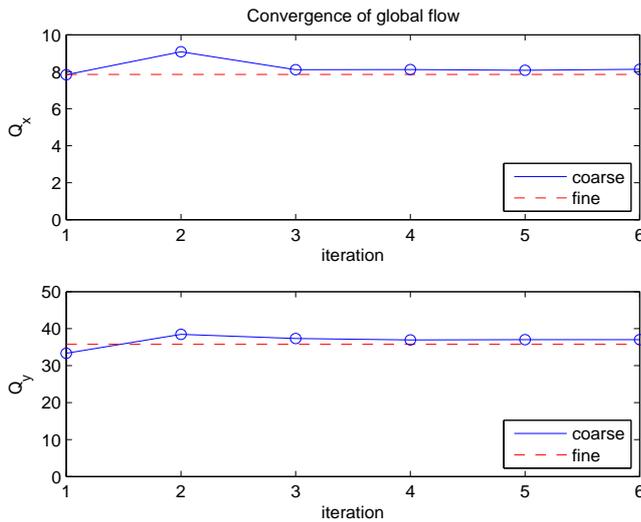

FIG. 4.5. *Convergence of the global flow from the coarse-scale solutions to the pressure equation with permeability defined by a modified SPE layer 44, and generic boundary conditions. In both cases, the five outer iterations correspond to an initial iteration using extended local upscaling on a 32-cell Cartesian grid, three iterations of local-global upscaling on CCAR grids containing 32, 116, and 389 cells, respectively, followed by a final iteration of local-global upscaling on the 389-cell grid. Top plot: global flow in the x-direction. Bottom plot: global flow in the y-direction.*

We repeat the comparisons of the upscaling methods with layer 48. Table 4.6 lists the errors in total flow for the same upscaling methods and boundary conditions. MLLG-W yields comparable or greater accuracy for all six boundary conditions. It is again interesting to note that for both Cartesian and CCAR grids, the new transmissibility weighting significantly improves accuracy for the highest-flow case (flow in the x-direction), without much loss of accuracy for the other boundary conditions, due
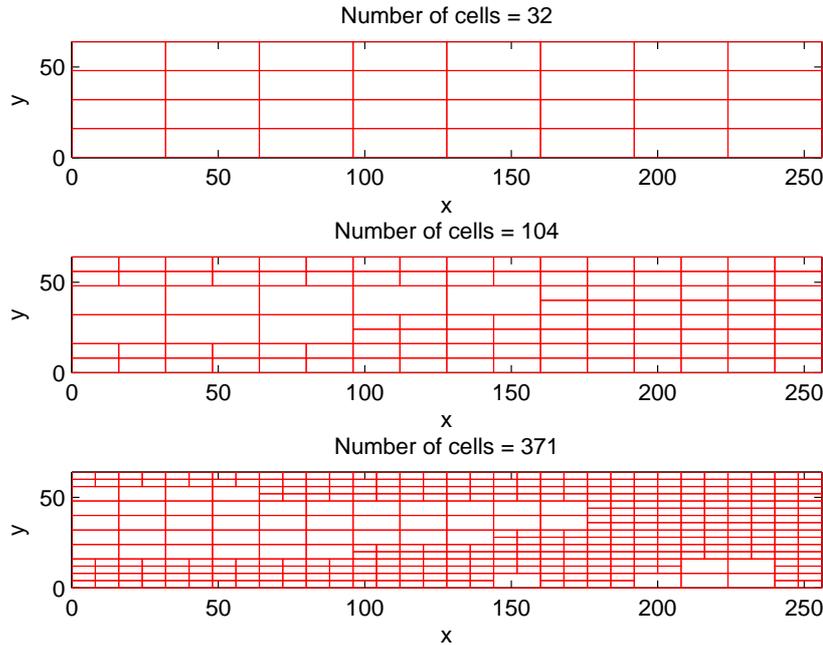
Fig. 4.6. *Grids obtained by successive refinement, based on the distribution of the flux obtained from solutions of the pressure equation, with permeability defined by a modified SPE layer 44 and generic boundary conditions.*

to its improved reduction of process dependency. Furthermore, applying the same upscaling method to a CCAR grid instead of a Cartesian grid results in improved accuracy in most cases for fewer grid cells. Finally, note the particularly dismal performance of extended local upscaling, compared to layer 44 above. It is important to remember that these results were obtained with a TPFA method.

Flow	Fine scale	EL	LG	LG-W	MLLG	MLLG-W
In x-direction	3.1182	65.11	6.72	4.60	5.50	3.37
In y-direction	1.6742	19.13	6.46	7.12	6.29	6.92
From lower left	0.44651	36.38	9.34	7.80	7.14	5.13
From upper left	1.2255	89.84	11.31	13.63	11.71	14.03
From lower right	-1.8434	87.43	13.31	16.17	14.76	17.03
From upper right	-0.27058	26.11	34.71	36.91	34.36	34.04

TABLE 4.6

Errors in total flow for modified SPE10 Layer 48, for various upscaling methods and boundary conditions. The EL, LG and LG-W methods use a 32×16 Cartesian grid, MLLG uses a 383-cell CCAR grid, and MLLG-W a 386-cell CCAR grid. The difference in these CCAR grids is due to the differences in transmissibility calculations, thus differences in local flow calculations, which impact the refinement criteria.

Figures 4.7 and 4.8 show a portion of the velocity field computed by LG-W and MLLG-W, respectively, for the case in which flow is driven in the x-direction. Both

methods accurately resolve the fine-scale velocity, but MLLG-W does so with fewer cells.
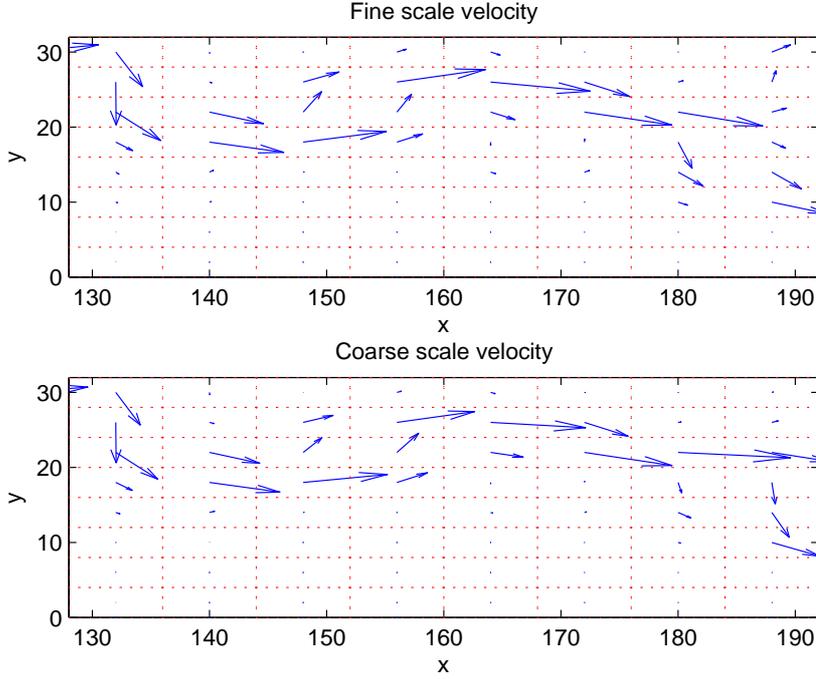

FIG. 4.7. *Velocity fields for channelized domain (layer 48), with flow driven in the x-direction. Top plot: fine-scale velocity, averaged over each cell of a 32×16 Cartesian grid. Bottom plot: coarse-scale velocity, obtained using method LG-W.*

Table 4.7 shows the percentage errors in total flow for layer 76. This layer poses severe challenges to TPFA methods. Indeed, none of the methods are able to match the fine-scale flow with reasonable accuracy in the cases of flow from the upper left or lower right corner, so they are omitted here. For the remaining boundary conditions, we observe that (1) using a CCAR grid preserves the accuracy obtained using a Cartesian grid with more cells, and (2) gradient-based weighting reduces the process dependency for both Cartesian and CCAR grids, making MLLG-W the best all-around choice.

Flow	Fine scale	EL	LG	LG-W	MLLG	MLLG-W
In x-direction	4.2935	32.91	5.69	3.39	5.93	3.99
In y-direction	1.5661	2.46	4.20	4.63	3.83	4.52
From lower left	0.42436	38.23	36.81	26.79	33.69	26.25
From upper right	-0.099821	9.30	16.96	19.46	17.26	20.50

TABLE 4.7
Errors in total flow for modified SPE10 Layer 76, for various upscaling methods and boundary conditions. The EL, LG and LG-W methods use a 32×16 Cartesian grid, LG-CCAR uses a 425-cell CCAR grid, and MLLG uses a 422-cell CCAR grid.

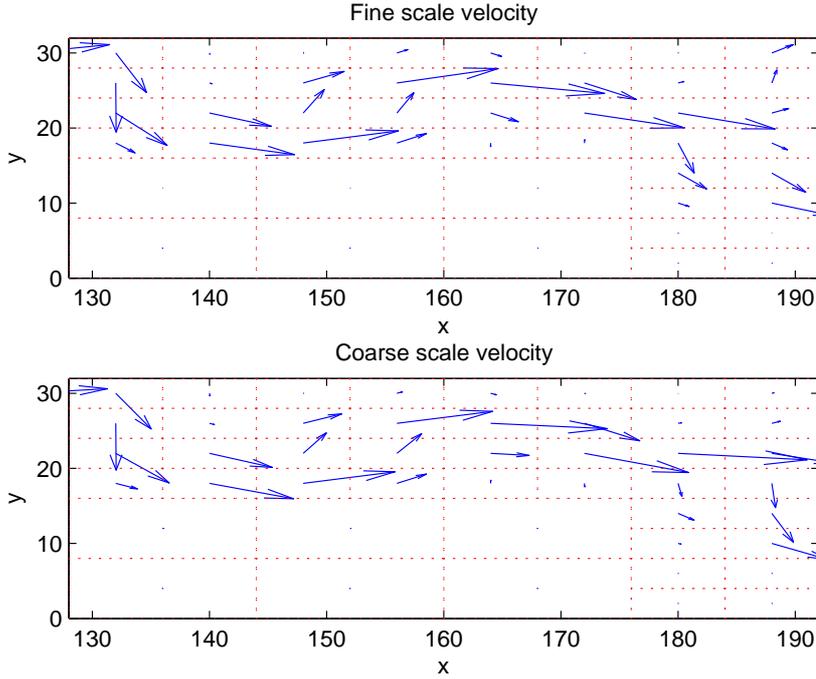Table 4.8 shows the percentage errors in total flow for the same upscaling methods

FIG. 4.8. *Velocity fields for channelized domain (layer 48), with flow driven in the x-direction. Left plot: fine-scale velocity, averaged over each cell of a 371-cell CCAR grid. Bottom plot: coarse-scale velocity, obtained using MLLG-W.*

for layer 80. In this layer, the main channel is oriented along the grid. Because of this orientation, the upscaling methods that employ gradient-based weighting do not provide as much benefit, but they still exhibit reduction in process dependency. In particular, note the smaller errors for the low-flow cases.

Flow	Fine scale	EL	LG	LG-W	MLLG	MLLG-W
In x-direction	14.6725	12.23	6.22	4.87	5.88	6.10
In y-direction	3.0696	4.07	0.11	1.24	0.02	1.29
From lower left	13.2104	6.81	6.69	6.77	4.28	6.21
From upper left	-8.9824	6.53	5.02	7.08	0.37	4.60
From lower right	-7.8925	5.31	9.66	6.83	10.05	8.28
From upper right	-0.60702	0.30	17.33	12.39	15.27	14.24

TABLE 4.8

Errors in total flow for modified SPE10 Layer 80, for various upscaling methods and bounday conditions. The EL, LG and LG-W methods use a 32×16 Cartesian grid, MLLG uses a 431-cell CCAR grid, and MLLG-W uses a 419-cell CCAR grid.

5. Discussion. This work shows that it is possible to obtain good upscaling accuracy on challenging permeability fields, such as the channelized cases used in this paper, on Cartesian adapted grids (CCAR grids) and two-point flux approximations (TPFA) using the proposed multi-level local-global method (MLLG). We compared

global flow errors for channelized systems, which are generally challenging for upscaling methods, produced by upscaled CCAR grids to errors produced by local-global upscaling on Cartesian grids of similar density. In [7] it was shown that the uniform local-global method on these channelized systems strongly outperformed local or extended local upscaling methods. For the generic boundary conditions used in the upscaling process the upscaled CCAR grids cut these global flow errors again, often as much as 50% or more. Process dependency is also reduced significantly. These results suggest that refinement is an effective means to control upscaling errors and reduce process dependency. A large contributor to the improved accuracy presented in this paper is the new gradient-based weighting of transmissibilities, which we strongly recommend for local-global upscaling strategies.

We are especially encouraged by the accuracy we obtain with our method in the local flow. The combination of adaptivity and local-global upscaling retains important connected flow paths and accurately compute local flow directions and speeds. This is promising for displacements with high mobility contrasts, such as gas injection processes, where the mobile fluid will seek these high permeability paths. We are implementing the CCAR methods in a compositional solver to further investigate this potential benefit.

The computational overhead of MLLG is relatively low, because global solves are performed on coarse grids, the sizes of extended local regions are limited, and generally very few steps are required for the local-global iteration to converge.

The effectiveness of MLLG is good news, as it means that it is perhaps not as essential to align computational grids, thus simplifying gridding strategies, or to always use more expensive multi-point flux approximations. Multi-point flux approximations are necessary in cases with very strong permeability tensor effects and/or weak flows. If such areas are local and can be identified, this opens the possibility of using adaptive stencils that employ TPFA where possible and revert to MPFA in the critical areas. One such method is introduced in [23], and other related work can be found in [2].

Future work will include the investigation of other possible adaptivity criteria, including criteria for coarsening, and alternative methods of handling occurrences of negative transmissibilities. The upscaling approach presented here, while suitable for CCAR grids, can just as easily be applied to other types of grids. Future work will include generalization of our algorithm to other gridding strategies such as patched refinements, unstructured grids, or curvilinear grids.

6. Conclusions. The following main conclusions can be drawn from this study:
- For highly heterogeneous (e.g., channelized) systems, the proposed integration of grid adaptivity and upscaling is shown to consistently provide more accurate coarse-scale models for global flow, relative to reference fine-scale results, than do existing upscaling techniques using TPFA discretizations, including local-global upscaling, applied to uniform grids of similar densities.
- Adaptation allows the use of a coarse-scale model that contains fewer grid points than would be required when using a Cartesian grid to achieve the same accuracy during simulation.
- Apart from more accurate global flow results, the proposed method also computes improved local flow results. Connected high permeability flow paths are preserved, and local velocity directions and magnitude in these important paths are computed with good accuracy.
- We developed a refinement indicator based on local fluxes that guides the grid

adaptation process during upscaling and further reduces solver and upscaling errors.

- Process dependency is strongly reduced, that is, the approach computes accurate global flow results also for flows driven by boundary conditions different from the generic boundary conditions used to compute the upscaled parameters.
- Upscaling accuracy can be significantly improved and process dependency reduced with the use of a gradient-based weighting of transmissibilities computed from multiple local flow problems in the local-global method.

REFERENCES

[1] Aavatsmark I, Barkve T, Mannseth T. Discretization on Non-Orthogonal, Quadrilateral Grids for Inhomogeneous, Anisotropic Media. Journal of Computational Physics, 127:2-14, 1996.

[2] Aavatsmark I, Eigestad GT and Nordbotten JM. A compact MPFA method with improved robustness. Proceedings of the 10th European Conference on the Mathematics of Oil Recovery, 2006.

[3] Berger MJ, Colella P. Local adaptive mesh refinement for shock hydrodynamics. Journal of Computational Physics, 82:64-84, 1989.

[4] Berger MJ, Oliger JE. Adaptive Mesh Refinement for Hyperbolic Partial Differential Equations. Journal of Computational Physics, 53:484-512, 1984.

[5] Bourgeat A. Homogenized Behavior of Two-Phase Flows in Naturally Fractured Reservoirs with Uniform Fractures Distribution. Computational Methods in Applied Mechanics and Engineering, 47:205-16, 1984.

[6] Chen Y, Durlofsky LJ. Adaptive Local–Global Upscaling for General Flow Scenarios in Heterogeneous Formations. Transport in Porous Media 62:157-185, 2006.

[7] Chen Y, Durlofsky LJ, Gerritsen MG, Wen XH. A coupled local-global upscaling approach for simulating flow in highly heterogeneous formations. Advances in Water Resources, 26:1041-1060, 2003

[8] Christie MA, Blunt MJ. Tenth SPE comparative solution project: a comparison of upscaling techniques. SPE Reservoir Evaluat Eng 2001;4:308-17.

[9] Deutsch C, Journel AG. GSLIB: Geostatistical Software Library and User's Guide, 2nd. Edition, Oxford Press, 1998.

[10] Durlofsky LJ. Numerical Calculation of Equivalent Grid Block Permeability Tensors for Heterogeneous Porous Media. Water Resources Research, 27:699-708, 1991.

[11] Edwards MG. Elimination of Adaptive Grid Interface Errors in the Discrete Cell Centered Pressure Equation. Journal of Computational Physics, 126:356-372, 1996.

[12] Edwards MG, Rogers CF. Finite Volume Discretization with Imposed Flux Continuity for the General Tensor Pressure Equation. Computational Geosciences, 2:259-290, 1998.

[13] Gautier Y, Blunt MJ, Christie MA. Nested Gridding and Streamline-Based Simulation for Fast Reservoir Performance Prediction. SPE Reservoir Simulation Symposium, 1999, SPE 51931.

[14] Gerritsen MG, Jessen K, Mallison BT, Lambers JV. A Fully Adaptive Streamline Framework for the Challenging Simulation of Gas-Injection Processes. SPE ATC, 2005, SPE 97270.

[15] Gerritsen MG, Lambers JV. Solving Elliptic Equations in Heterogeneous Media using Cartesian Grid Methods with Anisotropic Adaptation. To be submitted to Journal of Computational Physics, 2007.

[16] Ham FE, Lien FS, Strong, AB. A Cartesian Grid Method with Transient Anisotropic Adaptation. Journal of Computational Physics, 179:469-494, 2002.

[17] He C. Structured Flow-Based Gridding and Upscaling for Reservoir Simulation. PhD Thesis, Department of Petroleum Engineering, Stanford University.

[18] Henson VE, Yang UM. BoomerAMG: a Parallel Algebraic Multigrid Solver and Preconditioner. Applied Numerical Mathematics, 41:155-77, 2002.

[19] Holden L, Nielsen BF. Global Upscaling of Permeability in Heterogeneous Reservoirs: the Output Least Squares (OTL) Method. Transport in Porous Media, 40:115-43, 2000.

[20] Hornung R, Trangenstein J. Adaptive Mesh Refinement and Multilevel Iteration for Flow in Porous Media. Journal of Computational Physics, 136:522-545, 1997.

[21] Kippe V, Aarnes JE, Lie KA. A Comparison of Multiscale Methods for Elliptic Problems in Porous Media Flow. To be published in Computational Geosciences, Special issue on Multiscale Methods, 2007.

[22] Lee SH, Tchelepi HA, Jenny P, DeChant LJ. Implementation of a Flux-Continuous Finite-Difference Method for Stratigraphic, Hexahedron Grids. SPE Journal, 7:267:277, 2002.

[23] Lambers JV, Gerritsen MG, Mallison BT. Accurate Local Upscaling with Variable Compact Multi-point Transmissibility Calculations. To be published in Computational Geosciences, Special issue on Multiscale Methods, 2007.

[24] Mittal R, Iaccarino, G. Immersed Boundary Methods. Annual Review of Fluid Mechanics, 37:239-61, 2005.

[25] Nilsson J, Gerritsen MG, Younis RM. A Novel Adaptive Anisotropic Grid Framework for Efficient Reservoir Simulation. Proc. of the SPE Reservoir Simulation Symposium, 2005, SPE 93243.

[26] Nordbotten JM, Aavatsmark I, Eigestad GT. Monotonicity of Control Volume Methods. Submitted to Numerische Mathematik, 2006.

[27] Pickup GE, Ringrose PS, Jensen JL, Sorbie KS. Permeability Tensors for Sedimentary Structures. Mathematical Geology, 26:227-50, 1994.

[28] Sammon, P H, "Dynamic Grid Refinement and Amalgamation for Compositional Simulation", SPE RSS, 2003, SPE 79683.

[29] Trangenstein J, Bi Z. Large Multi-Scale Iterative Techniques and Adaptive Mesh Refinement for Miscible Displacement Simulation. SPE/DOE Improved Oil Recovery Symposium, 2002, SPE 75232.

[30] Watson DF. Contouring: A guide to the analysis and display of spacial data. Pergamon, 1994.

[31] Wen XH, Durlofsky LJ, Edwards MG. Use of Border Regions for Improved Permeability Upscaling. Mathematical Geology, 35:521-547, 2003.

[32] Younis RM, Caers J. A Method for Static-based Upgridding. Proc. of the 8th European Conference on the Mathematics of Oil Recovery, Sept. 2002