
Modelling of Multiscale Structures in Flow Simulations for Petroleum Reservoirs

Jørg Aarnes¹, Vegard Kippe¹, Knut–Andreas Lie¹, and Alf Birger Rustad²

¹ SINTEF ICT, Dept. of Applied Mathematics

² Statoil Research Centre, Trondheim

Summary. Flow in petroleum reservoirs occurs on a wide variety of physical scales. This poses a continuing challenge to modelling and simulation of reservoirs since fine-scale effects often have a profound impact on flow patterns on larger scales. Resolving all pertinent scales and their interaction is therefore imperative to give reliable qualitative and quantitative simulation results. To overcome the problem of multiple scales it is customary to use some kind of upscaling or homogenisation procedure, in which the reservoir properties are represented by some kind of averaged properties and the flow is solved on a coarse grid. Unfortunately, most upscaling techniques give reliable results only for a limited range of flow scenarios. Increased demands for reservoir simulation studies have therefore led researchers to develop more rigorous multiscale methods that incorporate subscale effects more directly.

In the first part of the paper, we give an overview of some of the many scales that are considered important for flow simulations. Next, we present and discuss several upscaling approaches that have played a role in the history of reservoir simulation. In the final part, we present some more recent approaches for modelling scales in the flow simulations based upon the multiscale paradigm. We conclude with a discussion of benefits and disadvantages of using multiscale methods, rather than using traditional upscaling techniques, in reservoir simulation.

1 Introduction

Simulation of petroleum reservoirs started in the mid 1950's and has become an important tool for qualitative and quantitative prediction of the flow of fluid phases. Today, computer models have a widespread use as a complement or even a competitor to field observations, pilot field and laboratory tests, well testing and analytical models. However, even though reservoir simulation can be an invaluable tool to enhance oil-recovery, the demand for simulation studies depends on many factors. For instance, petroleum fields vary in size from small pockets of hydrocarbon that may be buried just a few meters beneath the surface of the earth, to huge reservoirs stretching out several square kilometres beneath remote and stormy seas. This illustrates that the

value of reservoir simulation studies depends on what kind of extra profit one can expect by increasing the oil-recovery from a given reservoir.

A typical North Sea reservoir is located some two kilometres under the sea floor and several hundred kilometres offshore with a sea depth between one hundred and three hundred meters. These reservoirs are often fairly large horizontally, but the vertical dimensions of the zones carrying hydrocarbon may be just a few meters. The rock structures are usually layered with different mixtures of rock types (called facies) in different layers forming an irregular pattern. In addition, geological activity has created faults and fractures that highly influence flow properties. On the other end of the length scale, the fluids in the reservoir are embedded in porous formations like water in a sponge, and flow through small channels formed by the void space between rock particles. The volume fraction of the porous media formations consisting of pores open to fluid flow (henceforth called porosity) is typically in the range 0.1–0.3. To produce petroleum from the reservoir, one drills wells into the rock formations carrying the petroleum resources. In simple terms, a well can be considered as a hole with a diameter of about ten inches extending several kilometres into the earth and through the reservoir. The hole is lined with a casing that is perforated at places where the well is to produce or inject fluids.

We shall let this type of reservoir provide a conceptual model of a petroleum reservoir on which we want to perform simulation studies. Flow processes in this type of reservoir involve a large gap in scales. On one end we have the kilometre scale of the reservoir. On the other end we have the micrometer scale of the pore channels. In between we find, for instance, the centimetre scale of the wells. To obtain a model of the reservoir, one builds geological models that attempt to reproduce the true geological heterogeneity in the reservoir rock. To create a geological model that properly reflects all pertinent scales that impact fluid flow is, however, impossible. Indeed, the size of a grid block in a typical geological grid-model is in the range 10–50 m in the horizontal direction and 0.1–1 m in the vertical direction. Thus, a geological model is clearly too coarse to resolve small scale features such as radiant flow around wells, and flows through narrow high permeable channels. Neither can we expect to resolve properly important geological features such as fractures and faults. Nevertheless, from a simulation point of view, geological models are too complex, i.e., they contain more information than we can exploit in simulation studies. Indeed, in simulation models we usually use a coarsened grid model, with the possible exception of regions in the proximity of wells, and variations in the geological model occurring at length scales below the simulation grid block scale are normally replaced with averaged or upscaled quantities.

Many fields of science face the problem of multiple scales and share the need for upscaling or homogenisation³. The main purpose of this chapter is to

³Homogenisation and upscaling are used interchangeably in the scientific literature to describe procedures for generating coarsened reservoir models. However,



Fig. 1. Layered geological structures as seen in these pictures typically occur on both large and small scales in petroleum reservoirs. Pictures are courtesy of Silje Søren Berg, University of Bergen.

discuss alternative strategies for modelling reservoir heterogeneity at the scale of the geological model in a coarsened model suitable for fluid flow simulation. We review and discuss some historically important upscaling techniques in Section 3. The basic motivation behind these upscaling regimes is to create simulation models that produce flow scenarios that are in close correspondence with the flow scenarios that one would obtain by running simulations directly on the geological models. Unfortunately, experience has shown that it is very difficult to design a robust upscaling scheme that gives somewhat reliable results for all kinds of flow scenarios. Therefore, due to the limitations of upscaling, a new type of methodology based on so-called *multiscale methods* has started to gain popularity. The multiscale paradigm generally refers to a class of methods that incorporate fine-scale information into a set of coarse-scale equations in a way that is consistent with the local properties of the differential operators. In Section 4 we present a recent class of methods that seek to incorporate subscale information more directly in the flow simulation as a means to avoid the use of upscaling. But before we discuss different upscaling and multiscale approaches for simulating flow in reservoirs, we will identify some scales of importance.

2 Models and Scales in Porous Media Flow

In petroleum reservoirs sedimentary structures are formed when sediments are deposited to form thin layers called laminae. Due to alternating layers of coarse and fine-grained material, laminae may exhibit large permeability contrasts on the mm-cm scale, but these effects are usually left unresolved

homogenisation is also the name of a specific mathematical theory (for asymptotic analysis of periodic structures) that has been applied, by coincidence, to upscaling geological models for reservoir simulation. Thus, to avoid confusion we will only use the term upscaling for this coarsening procedure.

in geological models. Laminae are stacked to form beds, which are the smallest stratigraphic units. The thickness of beds varies from millimetres to tens of meters, and different beds are separated by thin layers with significantly lower permeability. Beds are, in turn, grouped and stacked into parasequences or sequences (parallel layers that have undergone similar geologic history). Parasequences represent the deposition of marine sediments, during periods of high sea level, and they tend to be somewhere in the range from meters to tens of meters thick and have a horizontal extent of several kilometres.

The trends and heterogeneity of parasequences depend on the depositional environment. For instance, whereas shallow marine deposits may lead to rather smoothly varying permeability distributions with correlation lengths in the order of 10–100 meters, fluvial reservoirs may contain intertwined patterns of narrow high-flow channels on a background of significantly lower permeability. Fractures and faults, on the other hand, are created by stresses in the rock. A fault is a planar surface in the rock, across which the rocks have been displaced. A fracture is a crack or a surface of breakage, across which the rocks have not been displaced. These structures may have higher or lower permeability than the surrounding rocks. The reservoir geology can also consist of other structures like for instance shale layers, which consist of impermeable clays and are the most abundant sedimentary rock.

We have seen that geological structures in petroleum reservoirs occur at a range of length scales. Moreover, it is known that geological structures at all scales can have a profound impact on fluid flow. Hence, ideally geological features that span across all types of length scales should be reflected in the geological reservoir model. However, since a geological model can only cope with a certain range of scales, we need to discretise with respect to scale, try to identify the most important scales, and develop different models when studying different phenomena. Choosing scales is often done by intuition and experience, and it is hard to give very general guidelines, but an important concept in choosing model scales is the notion of representative elementary volumes (REVs). This concept is based on the idea that petrophysical flow properties (typically porosity and permeability) are constant on some intervals of scale, see Figure 2. REVs, if they exist, mark transitions between scales of heterogeneity, and present natural length scales for modelling.

In order to identify a range of length scales where REVs exist for porosity (or permeability), we move along the length-scale axis from the micrometer scale of the pores toward the kilometre scale of the reservoir. At the pore scale level the porosity is a rapidly oscillating function equal to zero (in solid rock) and one (in the pores). Hence, obviously no REVs can exist at this scale. At the next characteristic length scale, the core scale level, we find laminae deposits. Since the laminae consist of alternating layers of coarse and fine grained material, we cannot expect to find a common porosity value for the different rock structures. Moving further along the length scale axis we may find long thin layers, perhaps extending throughout the entire horizontal length of the reservoirs. Each of these individual layers may be nearly homogeneous since

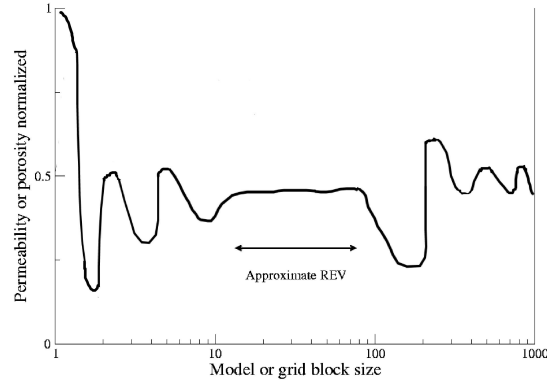


Fig. 2. Flow properties may be approximately constant on scale intervals.




Fig. 3. Various scales in a reservoir.

they are created by the same geological process, and probably contain approximately the same rock types. Hence, at this level it sounds reasonable to speak of REV's. If we move to the high end of the length scale axis we start to group more and more layers into families of layers with different sedimentary structures, and porosity REV's will probably not exist.

The previous discussion gives some grounds to claim that reservoir rock structures contain scales where REV's may exist. However, from a general point of view the existence of REV's in porous media is highly disputable. For instance a faulted reservoir with faults distributed continuously both in length and size throughout the reservoir, will typically have no REV. Moreover, no two reservoirs are identical so it is difficult to capitalise from previous experience. Indeed, porous formations in reservoirs may vary greatly, also in terms of scales. Nevertheless, the concept of REV's can serve as a guideline when deciding what scales to model. In the following we will discuss four different model classes, as shown in Figure 3: pore scale models, core scale models, geological models and simulation models.




Fig. 4. Graphical illustration of a pore network.

2.1 Pore-Scale Model

Modelling flow at the pore scale is quite different from modelling flow at the reservoir scale. On the reservoir scale, the fundamental equations are continuity of fluid phases and Darcy’s law, as we will see below. Darcy’s law basically assumes that the primary forces driving the flow are the pressure gradient and gravity. Flow at the pore scale, on the other hand, is mainly dominated by capillary forces, although gravitational forces can still be important. We therefore need a different set of equations to model flow at the pore scale. Since it would be off-track to give a thorough presentation of pore-scale flow modelling, we give only a very simplified overview.

A pore-scale model (as illustrated in Figure 3) may be about the size of a sugar cube. At the pore scale the porous medium is usually represented by a graph (see e.g., [52]). A graph is a pair (V, E) , where V is a set whose elements are called vertices (or nodes), and E is a subset of $V \times V$ whose elements are called edges. The vertices are taken to represent pores, and the edges represent pore-throats (i.e., connections between pores), see Figure 4.

The flow process where one fluid invades the void space filled by another fluid is generally described as an invasion–percolation process. Invasion is the process where a phase can invade a pore only if a neighbouring pore is already invaded. For each pore there is an entry pressure (i.e., the threshold pressure needed for the invading phase to enter the pore) depending on the size and shape of pores, the size of pore throats, as well as other rock properties. Among those pores neighbouring invaded pores, the invading phase will first invade

the pore with lowest threshold pressure. This gives a way of updating the set of pores neighbouring invaded ones. Repeating the process establishes a recursive algorithm, determining the flow pattern of the invading phase. In the invasion process we are interested in whether a phase has a path through the model (i.e., percolates) or not, and the time variable is often not modelled at all. For pore networks this is misleading, because we are also interested in modelling the flow after the first path through the model has been established. After a pore has been invaded, the saturations in the pore will vary with pressures and saturations in the neighbouring pores (as well as in the pore itself). New pores may also be invaded after the first path is formed, so that we may get several paths through the model where the invading phase can flow. Once the invading phase percolates (i.e., has a path through the model) we can start estimating flow properties. As the simulation progresses we will have an increasing saturation for the invading phase, which can be used to estimate flow properties at different saturation compositions in the model.

As mentioned initially this overview was simplified. To elaborate on the underlying complexity we need to mention wettability. When two immiscible fluids (such as oil and water) contact a solid surface (such as the rock), one of them tends to spread on the surface more than the other. The fluid in a porous medium that preferentially contacts the rock is called the wetting fluid. Note that wettability conditions are usually changing throughout a reservoir. The flow process where the invading fluid is non-wetting is called drainage and is typically modelled with invasion–percolation. The flow process where the wetting fluid displaces the non-wetting fluid is called imbibition, and is more complex, involving effects termed film flow and snap-off. A further presentation is beyond the scope here, but the interested reader is encouraged to see [52] and references therein.

From an analytical point of view, pore-scale modelling is very important as it represents flow at the fundamental scale (or more loosely, where the flow really takes place), and hence provides the proper framework for understanding the fundamentals of porous media flow. From a practical point of view, pore-scale modelling has a huge potential. Modelling flow at all other scales may be seen as averaging of flow at the pore scale, and properties describing the flow at larger scales are usually a mixture of pore-scale properties. At larger scales, the complexity of flow modelling is often overwhelming, with large uncertainties in determining flow parameters. Hence being able to single out and estimate the various factors determining flow parameters is invaluable, and pore-scale models can be instrumental in this respect. However, to extrapolate properties from the pore scale to an entire reservoir is very challenging, even if the entire pore space of the reservoir was known (of course, in real life you will not be anywhere close to knowing the entire pore space of a reservoir).



Fig. 5. Three core plugs with diameter of one and a half inches, and a height of five centimetres.

2.2 Core-Scale Model

When drilling a well it is common to bring up parts of the rock from the well trajectory. Three such rock samples are shown in Figure 5. These rock samples are called cores or core plugs, and are necessarily confined (in dimension) by the radius of the well, although they lengthwise are only confined by the length of the well. Cores taken from a well can give detailed information at scales restricted only by human apparatus. For instance, using thin slices of the rock one may study the pore structure with an electron microscope with micrometer resolution. However, it is common to do also flow experiments on cores (e.g., flood them with water in a laboratory), thereby obtaining flow properties for the core, e.g., relative permeability curves and capillary pressure curves.

Properties from core-scale flow experiments are often used as input for the geological model, or directly for the simulation model. Cores should therefore ideally be representative for the heterogeneous structures that one may find in a typical grid block in the geological model. However, flow experiments are usually performed on relatively homogeneous cores that rarely exceed one meter in length. Cores can therefore seldom be classified as representative elementary volumes. For instance, cores may contain a shale barrier that blocks flow inside the core, but which does not extend much outside the well bore region. Thus, if the core was slightly wider, then there would be a passage past the shale barrier. Flow at the core scale level is also more influenced by capillary forces than flow on a reservoir field scale.

This discussion shows that the problem of extrapolating information from cores to build a geological model is largely under-determined. Supplementary pieces of information are also needed, and the process of gathering geological data from other sources is described next.

2.3 Geological Model

The two key characteristics that affect the fluid flow in a reservoir rock are its porosity and permeability. Porosity is the fraction of space not occupied by the rock, i.e., the volume fraction occupied by the phases (oil, water and gas). Permeability is a measure of how easy fluid flows through the rock. The main purpose of the geological model is therefore to provide the distribution of these petrophysical parameters, besides giving location and geometry of the reservoir. Thus, a geological model is a conceptual, three-dimensional representation of a reservoir and consists of a set of grid cells that each has a constant porosity and a corresponding constant permeability tensor. The tensor is represented by a matrix where the diagonal terms represent direct flow (i.e., flow in one direction caused by a pressure drop in the same direction), and the off-diagonal terms represent cross-flow (flow caused by pressure drop in directions perpendicular to the flow). Thus, a full tensor is needed to model local flow in directions at an angle to the coordinate axes. For example, in a layered system the dominant direction of flow will generally be along the layers. Thus, if the layers form an angle to the coordinate axes, then a pressure drop in, say, the x -coordinate direction will typically produce flow at an angle to x -direction. This type of flow can be modelled correctly only with a permeability tensor with nonzero off-diagonal terms.

Geological models are built using a combination of stratigraphy (the study of rock layers and layering), sedimentology (study of sedimentary rocks), and interpretation of measured data. Unfortunately, building a geological model for a reservoir is like finishing a puzzle where most of the pieces are missing. Ideally, all available information about the reservoir is utilised, but the amount of data available is limited due to costs of acquiring them. Seismic surveys give a sort of X-ray image of the reservoir, but they are both expensive and time consuming, and can only give limited resolution (you cannot expect to see structures thinner than ten meters from seismic data). Wells give invaluable information, but the results are restricted to the vicinity of the well. While seismic has (at best) a resolution of ten meters, information on a finer scale are available from well-logs. Well-logs are basically data from various measuring tools lowered into the well to gather information, e.g., radiating the reservoir and measuring the response. Even well-logs give quite limited resolution, rarely down to centimetre scale. Detailed information is available from cores taken from wells, where resolution is only limited by the apparatus at hand. The industry uses X-ray, CT-scan as well as electron microscopes to gather high resolution information from the cores. However, information from cores and well-logs are from the well or near the well, and extrapolating this information to the rest of the reservoir is subject to great uncertainty. Moreover, due to costs, the amount of data acquisitions made is limited. You cannot expect well-logs and cores to be taken from every well. All these techniques give separately small contributions that can help build a geological model. However, in the end we still have very limited information

available considering that a petroleum reservoir can have complex geological features that span across all types of length scales from a few millimetres to several kilometres.

In summary, the process of making a geological model is generally strongly under-determined. It is therefore customary to use *geostatistics* to estimate the subsurface characteristics between the wells. Using geostatistical techniques one builds petrophysical realisations in the form of grid models that both honour measured data and satisfy petrophysical trends and heterogeneity. Since trends and heterogeneity in petrophysical properties depend strongly on the structure of sedimentary deposits, high-resolution petrophysical realisations are not built directly. Instead, one starts by building a facies model. A facies is the characteristics of a rock unit that reflects its origin and separates it from surrounding rock units. By supplying knowledge of the depositional environment (fluvial, shallow marine, deep marine, etc) and conditioning to observed data, one can determine the geometry of the facies and how they are mixed. In the second step, the facies are populated with petrophysical data and stochastic simulation techniques are used to simulate multiple realisations of the geological model in terms of high-resolution grid models for petrophysical properties. Each grid model has a plausible heterogeneity and can contain from a hundred thousand to a hundred million cells. The collection of all realisations gives a measure of the uncertainty involved in the modelling. Hence, if the sample of realizations (and the upscaling procedure that converts the geological models into simulation models) is unbiased, then it is possible to supply predicted production characteristics, such as the cumulative oil production, obtained from simulation studies with a measure of uncertainty.

2.4 From Geomodel to Simulation Model

For full sized reservoirs the traditional approach has been to model geological structures with a geological model, and fluid flow with a coarser simulation model. Indeed, core models and pore models are designed only to give input to the geological model and perhaps to derive flow parameters for the simulation model. As we have seen, geological models are designed to reproduce the true geological heterogeneity in the reservoir rock and possibly incorporate a measure of inherent uncertainty. The grid cells in the geomodels are considered as a representative volume having certain geometry and constant petrophysical properties like porosity and permeability. The petrophysical properties are parameters in the equations governing fluid flow, and the size of the grid cells are typically of order 10–50 m in the horizontal directions and 10 cm to 10 m in the vertical direction, allowing the geomodel to capture the most important variations and characteristics in the petrophysical parameters. However, rocks are heterogeneous at all levels and the current choice of scales means that geomodels are too coarse to capture small-scale geological structures like tidal depositional layers that occur on a centimetre or even smaller scale.

Geological structures on a finer scale than the resolution of the geomodel may have a strong impact on fluid flow. On the other hand, geomodels are typically too detailed for routine flow simulations. Indeed, flow simulations are usually performed on coarser models where the petrophysical properties in the geomodel are replaced with an average of some kind over regions that correspond to a grid block in the simulation model. There are several reasons for this. First of all, the global flow in the reservoir may be dominated by large-scale phenomena (occurring on a scale of 10-1000 m) and a coarser model may therefore be sufficient to predict the reservoir production with the required accuracy. Secondly, since coarse models contain fewer parameters, they are easier to history-match⁴ than fine-grid models. Similarly, the results of very detailed flow simulations may be masked by the large uncertainty involved in reservoir modelling. Thus, by running flow simulations directly on geomodels we get more detailed flow scenarios, but they might not produce more accurate production characteristics.

The primary reason for not running simulations on geomodels, however, is that geomodels typically contain far too many grid-blocks from a numerical point of view and cannot be used directly because the memory and the computational time required are outside the limits of current computers. It is perhaps tempting to believe that with future increase in computing power, one will soon be able to simulate the flow in a full field reservoir model on a grid fine enough to capture rock properties at the smallest significant scale. The development so far, indicates that such an idea is wrong. The trend is rather to increase the number of scales and the sizes of models whenever more computing power or new numerical techniques become available. Indeed, the sizes and complexity of geomodels have increased continuously with the increase in computer memory and processing power. Hence, the need for handling different scales will most likely persist in the foreseeable future and there will be a continuing need for rescaling techniques in order to upscale and downscale models.

Upscaling techniques have primarily been used to go from geomodels to simulation models, but similar techniques are also used to go from e.g., a pore-scale model to a geomodel. By upscaling detailed geomodels one generates reduced simulation models that have fewer grid cells and possibly less irregular grid structures. The effective petrophysical properties are calculated in each cell of the simulation grids based on properties of the underlying geomodels. In this process, the aim is to preserve the small-scale effects in the large-scale computations (as well as possible). Systematic small-scale variations in permeability and porosity can have a significant effect on a larger scale, and this should be captured in the upscaled model. Upscaling techniques range from simple averaging techniques to relatively complex flow analyses. Although the latter type involves solving numerous local flow problems that

⁴History matching is the process of calibrating parameters (here porosity and permeability) to observations.

require a substantial computational effort, they do not always produce reliable results. Indeed, if the heterogeneity in the reservoir rock does not contain well-separated length-scales, this type of upscaling may lead to highly erroneous results, even for single-phase flow.

Downscaling techniques are typically used when going from a simulation model to a geomodel. For instance, one can be interested in refining coarse-scale modifications in petrophysical properties obtained during a history match on a simulation model in order to update the underlying geomodel. In this process, the aim is to preserve both the coarse-scale trends and the fine-scale heterogeneity structures.

2.5 Reservoir Simulation Model

The fundamental equations that model flow in porous media are the continuity equations. These equations ensure conservation of mass for the different fluids phases, and take the following form

$$\frac{\partial}{\partial t}(\phi\rho_i s_i) + \nabla \cdot (\rho_i v_i) = q_i. \quad (1)$$

Here each fluid phase is modelled by its density ρ_i , phase velocity v_i , and saturation s_i . The saturations are the volume fractions occupied by each phase and satisfy the closure relation $\sum_{i=1}^n s_i = 1$. Production and injection wells are modelled by the source term q_i . The phase velocities v_i are related to the phase pressure p_i through an empirical relation called Darcy’s law

$$v_i = -\frac{\mathbf{K}k_{ri}}{\mu_i}(\nabla p_i - \rho_i G). \quad (2)$$

Here \mathbf{K} is the rock permeability given from the geological model; μ_i is the viscosity of phase i ; $k_{ri} = k_{ri}(s_1, \dots, s_n)$ is the relative permeability of phase i , i.e., reduced permeability due to the presence of other phases; p_i is phase pressure; and $G = gn_z$ where g is the gravitational constant and n_z is the unit vector pointing in the downward vertical direction. It is also customary to introduce the phase mobility: $\lambda_i = k_{ri}(s_i)/\mu_i$.

The primary unknowns in simulation models are the phase saturations and the phase pressures. Darcy’s law together with the continuity equations (1) gives us one equation for each phase. The closure relation $\sum_{i=1}^n s_i = 1$ for the phase saturations gives us one additional equation. This gives us a complete model for single-phase flow. The simplest model for single-phase flow is obtained by assuming incompressibility (constant density). In this case, the equations (1)–(2) simplify to an elliptic equation for the fluid pressure

$$-\nabla \cdot \left[\frac{\mathbf{K}}{\mu} (\nabla p - \rho G) \right] = \frac{q}{\rho}. \quad (3)$$

For multi-phase flow, however, we can eliminate only one of the saturations and extra closure relations in terms of capillary pressure functions are added.

That is, one assumes that the phase pressures are related through a function describing the difference between the phase pressures across fluid interfaces. These functions are assumed to depend solely on the phase saturations, and are obtained from experiments and simulations on core scale models, or from tables describing capillary pressure functions for different rock types with known properties. We should also mention that pore scale-modelling can provide capillary pressure curves, and provide an appropriate framework for understanding the capillary forces.

For two-phase flow we can eliminate one of the transport equations. We then get a coupled system consisting of pressure equation and an equation for the fluid transport. Moreover, if we assume that the two fluids, say, water (w) and oil (o) are incompressible, and introduce a so-called global pressure p (see e.g., [3] in this book) and a total velocity $v = v_o + v_w$, then the corresponding system reads

$$\nabla \cdot v = q, \quad \text{where} \quad v = -\mathbf{K}[\lambda \nabla p - (\lambda_w \rho_w + \lambda_o \rho_o)G], \quad (4)$$

$$\phi \frac{\partial s_w}{\partial t} + \nabla \cdot [f_w(v + \mathbf{K} \lambda_o \nabla p_{cow} + \mathbf{K} \lambda_o (\rho_w - \rho_o)G)] = \frac{q_w}{\rho_w}. \quad (5)$$

Here we have introduced the total mobility $\lambda = \lambda_w + \lambda_o$, the fractional flow of water $f_w = \lambda_w / \lambda$, the capillary pressure $p_{cow} = p_o - p_w$, and the accumulated contribution from the wells $q = q_w / \rho_w + q_o / \rho_o$.

The transport equation (5) contains three different terms that stem from three different types of forces: viscous forces give rise to the term $q_v = f_w v$, capillary forces give rise to $q_c = f_w \mathbf{K} \lambda_o \nabla p_{cow}$, and gravity forces give rise to $q_g = f_w \mathbf{K} \lambda_o (\rho_w - \rho_o)g$. The range of these different forces is influenced by the flow process and model. In particular, model size, flow-rate, heterogeneities and fluid system are all important ingredients in determining the balance of forces. There are no absolute rules as to what problems are dominated by the different forces, although some rules of thumb apply. Gas-oil systems are generally more influenced by gravity than oil-water systems due to larger density differences, and horizontally layered reservoirs with poor vertical communication are less influenced by gravity than reservoirs with layers that are tilted with respect to the vertical axis or have good vertical communication. Capillary effects are mainly small-scale phenomena, and are therefore more apparent in small-scale models, but capillary forces can play a dominant role also in large reservoir regions with strong heterogeneous structures. However, as a general rule we can say that viscous forces tend to be dominant on a reservoir (or field) scale, whereas capillary forces typically are dominant in core scale models. In fact, capillary forces are rarely included in simulation models, implying that small scale effects of capillary forces need to be adjusted for in the relative permeability curves.

In the introduction we discussed why reservoir simulation models are usually posed on a different length scale than the geomodel and how this introduces an upscaling of the grid models. Another reason for operating with different grid-models for the fluid simulation comes from the fact that geological

grid models tend to follow the spatial structure of the geology. For (classical) numerical methods to behave properly, one would generally prefer that grid-blocks are as regular as possible. However, reservoir dimensions impose very different length scales in the vertical and horizontal direction. Moreover, layers, fractures and faults all give rise to complex physical geometries and grid-blocks that represent the underlying geology are rarely orthogonal, but tend to be skewed, irregular, and even have non-neighbour connections. These geometrical problems are also related to the multiscale structure of the reservoirs and lack of local grid resolution, but will not be discussed at all in the current paper. Instead, we will focus on the multiscale structures in permeability and discuss approaches for meeting the challenges they impose on the simulation. To ease the presentation, we will therefore henceforth assume that our reservoirs have a shoe-box geometry and consist of Cartesian grid-blocks.

In the oil industry, it is common to discretise the pressure equation with a *finite-volume method*. These methods consider the grid cells V_i as control volumes, and invoke the divergence theorem to yield

$$\int_{V_i} q \, dx = - \int_{\partial V_i} \mathbf{K} [\lambda \nabla p - (\lambda_w \rho_w + \lambda_o \rho_o) G] \cdot n \, d\nu.$$

Here n is the outward unit normal on ∂V_i . A fully discrete scheme is obtained by expressing the flux across the cell interfaces in terms of the pressure at the cell centres. For instance, the most basic finite volume method, the two-point flux approximation (TPFA) scheme, takes the following form:

$$\int_{V_i} q \, dx = \sum_j T_{ij} (p_i - p_j). \quad (6)$$

Here the sum is taken over all non-degenerate interfaces, i.e., over all j such that $\partial V_j \cap \partial V_i$ has a positive measure. The coupling factors T_{ij} are called transmissibilities and we will return to these later.

Similarly, the transport equation (5) is usually discretised with a finite volume method where upstream weighting is used to discretise the term representing viscous forces. Invoking the divergence theorem once again we have

$$\int_{V_i} \phi \frac{\partial S_w}{\partial t} + \int_{\partial V_i} f_w (v + \mathbf{K} \lambda_o \nabla p_{cow} + \mathbf{K} \lambda_o (\rho_w - \rho_o) G) \, d\nu = \int_{V_i} \frac{q_w}{\rho_w} \, dx.$$

By upstream weighting, it is meant that $f_w v$ is computed using the saturation on the upstream side of the cell interfaces. For a discussion of how to discretise the terms representing gravity and capillary forces, we refer the reader to the previous chapter in this book [3], where a more thorough introduction to the governing equations and corresponding numerical methods is given.

3 Upscaling for Reservoir Simulation

In the introduction we saw how reservoir modelling often involves an upscaling phase, where high-resolution geological models containing several million

cells are turned into coarser grid models that are more suitable for fluid flow simulation. This process leads to many fundamental questions. For instance, do the partial differential equations modelling porous media flow at the coarse scale take the same form as the equations modelling flow at the subgrid level? And, if so, how do we honour the fine-scale heterogeneities at the coarse scale?

Even though upscaling has been a standard procedure in reservoir simulation for nearly four decades, nobody has answered these questions rigorously, except for cases with special heterogeneous formations such as periodic or stratified media. Most upscaling techniques are based on some kind of local averaging procedure in which effective properties in each grid-block are calculated solely from properties within the grid block. As such, the averaging procedure does not account for coupling with neighbouring coarse grid-blocks. This implies in particular that the upscaled quantities do not reflect the impact of global boundary conditions and large scale flow patterns in the reservoir. Because different flow patterns may call for different upscaling procedures, it is generally acknowledged that global effects must also be taken into consideration in order to obtain robust coarse-scale simulation models.

A related problem to upscaling is that of gridding. That is, what kind of grid do we want to use to represent our porous medium, what resolution do we need, and how do we orient the grid. To design robust simulation models, the grid should be designed so that the grid blocks capture heterogeneities on the scale of the block. This often implies that we need significantly more cells in a regular grid than if we use corner point grids in which the grid block corners are on straight, but possibly tilted lines, that extend from top-to-bottom of the reservoir. It is therefore much more common to use corner-point grids in reservoir simulation. In the vicinity of wells one might use a different type of grid, for instance some kind of radial grid, to account for a radial flow pattern in the near well region. The importance of obtaining a good grid representation of the porous medium should not be underestimated, but gridding issues will not be discussed any further here. We assume henceforth that we have a fixed coarse grid suitable for numerical simulation and focus on upscaling techniques for the pressure equation (4) and the saturation equation (5).

The literature on upscaling techniques is extensive, ranging from simple averaging techniques, e.g., [44], via local simulation techniques [10, 20], to multiscale methods [2, 17, 39, 41] and rigorous homogenisation techniques for periodic structures [11, 38, 43]. Some attempts have been made to analyse the upscaling process, e.g., [7, 62], but so far there is generally no theory or framework for assessing the quality of an upscaling technique. In fact, upscaling techniques are seldom rigorously quantified with mathematical error estimates. Instead, the quality of upscaling techniques is usually assessed by comparing upscaled production characteristics with those obtained from a reference solution computed on an underlying fine grid.

It is not within our scope to give a complete overview over the many upscaling techniques that have been applied in reservoir simulation. Instead, we refer the reader to the many review papers that have been devoted to this

topic, e.g., [9, 19, 55, 60]. However, to understand why some upscaling methods work well for some flow scenarios, and not for others, whereas other upscaling techniques may work well for a completely different set of flow scenarios, we do need to discuss some basic concepts. To this end, we start by giving a brief introduction to upscaling rock permeability for the pressure equation (3) or (4) in Subsection 3.1. In Section 3.2 we discuss techniques for generating pseudo functions for upscaling the saturation equation (5). In Subsection 3.3 we present a different type of upscaling technique for the saturation equation that splits the variables into volume averaged and fluctuating components. In this approach various moments of the fluctuating components are used to derive a coarse scale equation for grid block saturations. Finally, in Subsection 3.4 we describe a class of upscaling techniques that assume steady state flow, i.e., that the time derivative term in the continuity equations can be neglected in the upscaling process.

Before we proceed with details, we should mention that upscaling is also applied to other quantities, but then mostly using very simple techniques. For instance, quantities such as saturation and porosity are upscaled using simple volume averaging. That is, for a grid block V we obtain porosity and saturations, denoted ϕ^* and s_i^* , respectively, by

$$\phi^* = \frac{1}{|V|} \int_V \phi(x) dx, \quad s_i^* = \frac{1}{\phi^* |V|} \int_V s_i(x) \phi(x) dx.$$

Fluxes are upscaled similarly. Moreover, if a grid block is known to contain a certain fraction of different rock types (also often termed flow-units), then a so-called majority vote can be used to approximate the effective properties. In this method one identifies which rock type occupies the largest volume fraction of the grid block and simply assigns the properties associated with this rock type to the entire grid block. Note however, that such a simple approach is not robust, see, e.g., [63]. Actually, the problem of upscaling rock-types is closely related to that of upscaling permeability, but has not received much attention in the literature.

3.1 Single-Phase Upscaling

The process of upscaling permeability for the pressure equation (3) or (4) is often termed single-phase upscaling. Most single-phase upscaling techniques seek homogeneous block permeabilities that reproduce the same total flow through each coarse grid-block as one would get if the pressure equation was solved on the underlying fine grid with the correct fine-scale heterogeneous structures. However, to design upscaling techniques that preserve averaged fine-scale flow-rates is in general nontrivial because the heterogeneities at all scales have a significant effect on the large-scale flow pattern. A proper coarse-scale reservoir model must therefore capture the impact of heterogeneous structures at all scales that are not resolved by the coarse grid.

To illustrate the concept behind single-phase upscaling, let p be the solution that we obtain by solving

$$-\nabla \cdot \mathbf{K} \nabla p = q, \quad \text{in } \Omega \quad (7)$$

on a fine grid with a suitable numerical method, e.g., a TPFA scheme of the form (6). To reproduce the same total flow through a grid-block V we have to find a homogenised tensor \mathbf{K}_V^* such that

$$\int_V \mathbf{K} \nabla p \, dx = \mathbf{K}_V^* \int_V \nabla p \, dx. \quad (8)$$

This equation states that the net flow-rate v_V through V is related to the average pressure gradient $\nabla_V p$ in V through the upscaled Darcy law $v_V = -\mathbf{K}^* \nabla_V p$.

In the discrete case, the choice of an appropriate upscaling technique also depends on the underlying numerical method. For instance, if the pressure equation is discretised by a TPFA finite-volume scheme of the form (6), then grid-block permeabilities are used only to compute interface transmissibilities at the coarse scale. As a result, one needs correct coarse-scale transmissibilities to reproduce a fine-scale flow field in an averaged sense. Hence, instead of upscaled block-homogenised tensors \mathbf{K}^* , one should seek block transmissibilities T_{ij}^* satisfying

$$Q_{ij} = T_{ij}^* \left(\frac{1}{|V_i|} \int_{V_i} p \, dx - \frac{1}{|V_j|} \int_{V_j} p \, dx \right), \quad (9)$$

where $Q_{ij} = -\int_{\partial V_i \cap \partial V_j} (\mathbf{K} \nabla p) \cdot n \, d\nu$ is the total Darcy flux across $\partial V_i \cap \partial V_j$.

If all transmissibilities T_{ij}^* are positive, then the TPFA scheme defined by

$$\sum_{j: \partial V_i \cap \partial V_j \neq \emptyset} T_{ij}^* (p_i - p_j) = \int_{V_i} q \, dx,$$

will reproduce the net grid block pressures $p_i = \frac{1}{|V_i|} \int_{V_i} p \, dx$, and hence also the coarse grid fluxes Q_{ij} . Unfortunately, there is no guarantee that the transmissibilities defined by (9) are positive, or even exist. Neither can we guarantee that the upscaled permeability tensors defined by (8) are positive definite. Positive transmissibilities and positive definite permeability tensors ensure stability of TPFA finite-volume and finite-element methods, respectively. The possible absence of these properties illustrates that the fundamental problem of single-phase upscaling is ill-posed. However, this ill-posedness is usually not a subject of debate since upscaling methods are usually devised such that the occurrence of unphysical upscaled quantities is avoided. For a more thorough discussion of existence and uniqueness, we refer the reader to [62]. In the next subsections we review some standard single-phase upscaling methods.

Averaging Methods

The simplest form of upscaling permeability is to take some kind of average of the permeabilities at the subgrid level. A general class of averaging techniques is defined by the power average,

$$\mathbf{K}_V^{*,p} = \left(\frac{1}{|V|} \int_V \mathbf{K}(x)^p dx \right)^{1/p}, \quad -1 \leq p \leq 1.$$

Note that $p = 1$ and $p = -1$ correspond to the arithmetic and harmonic means respectively, while the geometric mean is obtained in the limit $p \rightarrow 0$.

The use of power averaging can be motivated by the so-called Wiener-bounds [61], which state that for a statistically homogeneous medium, the correct upscaled permeability will be bounded above and below by the arithmetic and harmonic mean, respectively. This result has also a more intuitive explanation. To see this, consider the one dimensional pressure equation:

$$-\partial_x(\mathbf{K}(x)\partial_x p) = 0 \quad \text{in } (0, 1), \quad p(0) = p_0, \quad p(1) = p_1.$$

It is easy to see that the solution of this equation induces constant Darcy velocity v . This implies that $\partial_x p$ must scale proportional to the inverse of $\mathbf{K}(x)$. Hence, we derive

$$\partial_x p = \frac{p_1 - p_0}{\mathbf{K}(x)} \left[\int_0^1 \frac{dx}{\mathbf{K}(x)} \right]^{-1} = \frac{p_1 - p_0}{\mathbf{K}(x)} \mathbf{K}_V^{*, -1}.$$

If we insert this expression into equation (8) we find that the correct upscaled permeability \mathbf{K}_V^* is identical to the harmonic mean $\mathbf{K}_V^{*, -1}$.

This result, which shows that the harmonic mean gives the correct upscaled permeability, is generally valid only in one dimension, but it applies to a special case in higher dimensions as well. Indeed, consider the following problem:

$$\begin{aligned} -\nabla \cdot \mathbf{K} \nabla p &= 0 & \text{in } V = (0, 1)^3, \\ p(0, y, z) &= p_0, & p(1, y, z) = p_1, \\ (-\mathbf{K} \nabla p) \cdot n &= 0 & \text{for } y, z \in \{0, 1\}, \end{aligned} \tag{10}$$

where n is the outward unit normal on $\partial\Omega$.

Now, if \mathbf{K} models a perfectly stratified isotropic medium with layers perpendicular to the x -axis (so that $\mathbf{K}(x, \cdot, \cdot)$ is constant for all x), then the solution models uniform flow in the x -coordinate direction. This means that for each pair $(y, z) \in (0, 1)^2$ the solution $p_{y,z} = p(\cdot, y, z)$ solves

$$-\nabla \cdot \mathbf{K} \nabla p_{y,z}(x) = 0 \quad \text{in } (0, 1), \quad p_{y,z}(0) = p_0, \quad p_{y,z}(1) = p_1.$$

It follows that

$$-K(x) \nabla p = -(K(x) \partial_x p_{y,z}, 0, 0)^T = -\mathbf{K}_V^{*, -1} (p_1 - p_0, 0, 0)^T.$$

Hence, the correct upscaled permeability is equal to the harmonic mean.

Similarly, if \mathbf{K} models instead a stratified isotropic medium with layers perpendicular to the y - or z -axis, then the correct upscaled permeability would be equal to the arithmetic mean. These examples show that averaging techniques can give correct upscaling in special cases, also in three dimensional space. However, if we consider the model problem (10) with a less idealised heterogeneous structures, or with the same heterogeneous structures but with other boundary conditions, then both the arithmetic and harmonic average will generally give wrong net flow-rates. Indeed, these averages give correct upscaled permeability only for cases with essentially one-dimensional flow.

To try to model flow in more than one direction, one could generate a diagonal permeability tensor with the following diagonal components:

$$k_{xx} = \mu_a^z(\mu_a^y(\mu_h^x)), \quad k_{yy} = \mu_a^z(\mu_a^x(\mu_h^y)), \quad k_{zz} = \mu_a^x(\mu_a^y(\mu_h^z)).$$

Here μ_a^ξ and μ_h^ξ represent respectively the arithmetic and harmonic means in the ξ -coordinate direction. Thus, in this method one starts by taking a harmonic average along grid cells that are aligned in one coordinate-direction. One then computes the corresponding diagonal by taking the arithmetic mean of all “one dimensional” harmonic means. This average is sometimes called the harmonic-arithmetic average and may give good results if, for instance, the reservoir is layered and the primary direction of flow is along the layers.

Despite the fact that averaging techniques can give correct upscaling in special cases, they tend to perform poorly in practise since the averages do not reflect the structure or orientation of the heterogeneous structures. It is also difficult to decide which averaging technique to use since the best average depends both on the heterogeneities in the media and the flow process we want to model (flow direction, boundary conditions, etc). To illustrate the dependence on the flow process we consider an example.

Example 1. Consider a reservoir in the unit cube $[0, 1]^3$ with two different geomodels that each consist of a $8 \times 8 \times 8$ uniform grid-blocks and permeability distribution as depicted in Figure 6. We consider three different upscaling methods: harmonic average; arithmetic average; and harmonic-arithmetic average. The geomodels are upscaled to a single coarse grid-block, which is then subjected to three different boundary conditions:

BC1: $p = 1$ at $(x, y, 0)$, $p = 0$ at $(x, y, 1)$, no-flow elsewhere.

BC2: $p = 1$ at $(0, 0, z)$, $p = 0$ at $(1, 1, z)$, no-flow elsewhere.

BC3: $p = 1$ at $(0, 0, 0)$, $p = 0$ at $(1, 1, 1)$, no-flow elsewhere.

Table 1 compares the observed coarse-block rates with the flow-rate obtained by direct simulation on the $8 \times 8 \times 8$ grid. For the layered model, harmonic and harmonic-arithmetic averaging correctly reproduces the vertical flow normal to the layers for BC1. Arithmetic and harmonic-arithmetic averaging correctly reproduces the flow along the layers for BC2. The harmonic arithmetic averaging method also performs reasonably well for corner-to-corner flow (BC3).




Fig. 6. Plot of the logarithm of two permeability distributions: (left) a layered medium, and (right) a cube extracted from the lower part of the fluvial Upper Ness formation from the 10th SPE test case [18].

Table 1. Flow-rates relative to the reference rate Q_R on the fine grid.

	Model 1			Model 2		
	BC1	BC2	BC3	BC1	BC2	BC3
Q_H/Q_R	1	2.31e-04	5.52e-02	1.10e-02	3.82e-06	9.94e-04
Q_A/Q_R	4.33e+03	1	2.39e+02	2.33e+04	8.22	2.13e+03
Q_{HA}/Q_R	1	1	1.14	8.14e-02	1.00	1.55e-01

For model two, however, most of the methods produce some significant errors, and none of the methods are able to produce an accurate flow-rate for the test cases with boundary conditions specified by BC1 and BC3.

Averaging techniques can also be used to upscale transmissibility. One such method is the half-block method [45]. In this method each grid block is divided into two pieces (using the average surface between the two opposing faces), obtaining six upscaled permeabilities for each block. The transmissibility T_{ij} for the coarse grid interface, which essentially models the conductivity across the corresponding interface in the direction of the associated coordinate unit normal n_{ij} , is then obtained by taking the harmonic average of the half-block arithmetic averages on both sides of the interface. This method has the advantage of improving the model resolution at a low computational cost.

Numerical Pressure Computation Techniques

To approximate the combined effect of heterogeneous structures and imposed boundary conditions, the idea of inverting local pressure computations was suggested by Begg et al. [10]. In this approach one solves, for each grid block V , a set of homogeneous pressure equations on the form

$$-\nabla \cdot \mathbf{K} \nabla p = 0 \quad \text{in } V,$$

with prescribed boundary conditions. Thus, this method raises the immediate question: what kind of boundary conditions should be imposed?




Fig. 7. A schematic of the pressure solver method with $p=1$ and $p=0$ along the inflow and outflow boundaries respectively, and no-flow boundary conditions elsewhere.

One alternative is to impose three different sets of boundary conditions for each block to create a pressure drop across the coarse cell in each of the three coordinate directions (see Figure 7). This gives us a set of flow-rates for each grid block that can be used to compute an effective diagonal permeability tensor with components

$$k_{xx} = -Q_x L_x / \Delta P_x, \quad k_{yy} = -Q_y L_y / \Delta P_y, \quad k_{zz} = -Q_z L_z / \Delta P_z.$$

Here Q_ξ , L_ξ and ΔP_ξ are respectively the net flow, the length between opposite sides, and the pressure drop in the ξ -direction inside V .

Another popular option is to choose periodic boundary conditions. That is, one assumes that each grid block is a periodic cell in a periodic medium and imposes full correspondence between the pressures and velocities at opposite sides of the block, e.g., to compute k_{xx} we impose the following boundary conditions:

$$\begin{aligned} p(1, y) &= p(0, y) - \Delta p, & p(x, 1) &= p(x, 0), \\ v(1, y) &= v(0, y), & v(x, 1) &= v(x, 0) \end{aligned}$$

This approach yields a symmetric and positive definite tensor [20], and is usually more robust than the directional flow boundary conditions.

Improved accuracy of the upscaled permeability can be obtained if one uses an oversampling technique, in which the domain of the local flow problem is enlarged with a border region surrounding each grid block. In this approach we let the coarse grid-block V be embedded in a macro-block V' and solve the elliptic pressure equation (7) in V' with some prescribed boundary conditions on $\partial V'$. The flow is then solved in the whole enlarged region, but the averaging of the permeability tensor is performed only over the (original) coarse block. The motivation for using such a technique is to better account for permeability trends that are not aligned with the grid directions and possible large-scale connectivity in the permeability fields.

Example 2. We revisit the test-cases considered in Example 1, but now we compare harmonic-arithmetic averaging (HA) with the pressure computation




Fig. 8. A schematic of the pressure solver method for upscaling transmissibility.

techniques using; (D) the boundary conditions depicted in Figure 7; and (P) the periodic boundary conditions with a directional pressure drop. The latter method gives rise to full permeability tensors, but for the cases considered here the off-diagonal terms in the upscaled permeability tensors are small, and are therefore neglected for simplicity.

Table 2. Flow-rates relative to the reference rate Q_R on the fine grid.

	Model 1			Model 2		
	BC1	BC2	BC3	BC1	BC2	BC3
Q_{HA}/Q_R	1	1	1.143	0.081	1.003	0.155
Q_D/Q_R	1	1	1.143	1	1.375	1.893
Q_P/Q_R	1	1	1.143	0.986	1.321	1.867

Table 2 compares the observed coarse-block rates with the flow-rate obtained by direct simulation on the $8 \times 8 \times 8$ grid. For the layered model, all methods give rise to the same diagonal permeability tensor, and hence give exactly the same results. For model 2 we see that the numerical pressure computation methods give significantly better results than the HA averaging technique. Indeed, the worst results for the pressure computation methods, which were obtained for corner-to-corner flow, is within a factor 2, whereas the HA averaging technique underestimates the flow rates for BC1 and BC3 by almost an order of magnitude.

A similar approach can be used to derive upscaled transmissibilities. A simple way of doing this is illustrated in Figure 8. Here we create a flow across the interface between V_i and V_j by imposing a pressure drop between the grid block faces opposite to $\partial V_i \cap \partial V_j$. Thus, by solving the corresponding pressure solution in the two-block domain $V_i \cup V_j$ numerically, we can compute the average pressures p_i and p_j in V_i and V_j respectively, and derive an upscaled transmissibility by requiring that equation (9) holds.

Global and Local-Global Upscaling Techniques

The methods described so far have all been local in nature; the averaging techniques derive the upscaled permeability solely from local heterogeneous structures, whereas the pressure computation techniques try to account for flow responses by solving local problems with prescribed boundary conditions. The numerical examples show, however, that all the local methods may fail to capture the correct flow behaviour. For instance, none of the local methods were able to capture the correct flow behaviour for model 2 with boundary conditions prescribed by BC3. The main problem for the pressure computation techniques is, of course, that we do not know a priori the precise flow that will occur in a given region of the reservoir. Thus, it is generally not possible to specify the appropriate boundary conditions for the local flow problems in a unique manner (unless we already have solved the flow problem).

In order to account for global flow patterns, Holden and Nielsen [37] proposed to solve the pressure equation once on a fine grid, and extract information about the fine-grid flow pattern to compute correct upscaled transmissibilities. This approach may seem to contradict the purpose of upscaling, but it can be justified for the numerical treatment of two-phase flows. Indeed, for two-phase flow the pressure equation has to be solved several times throughout the simulation, and the cost of solving the pressure once (or even a few times) on a fine grid will typically be negligible compared with the total computational cost of the full multi-phase flow simulation. Another question is of course the robustness of such an upscaling with respect to changes in the well configuration and global boundary conditions. For some cases the upscaled transmissibilities may prove useful, while for other cases they may not. A similar technique avoiding the solution of the fine-grid problem was presented by Nielsen and Tveito [51].

Another way to avoid the full fine-grid computation was proposed by Chen et al. [15, 16] in the form of a local-global technique that uses global coarse-scale calculations to determine boundary conditions for local calculations that determine upscaled transmissibilities. Since the initial coarse-grid calculation may still give quite poor boundary conditions for the subgrid problems, they use an iteration procedure to ensure consistency between the local and global calculations.

A problem encountered in the global and local-global upscaling procedures is the occurrence of negative or anomalously large transmissibilities. Holden and Nielsen [37] avoid negative and very large transmissibilities by perturbing the transmissibilities through an iterative process, involving an optimisation problem, until all the transmissibilities are contained in some interval $[a, b] \subset \mathcal{R}^+$. Chen et al. [15, 16] observed that unphysical transmissibilities occur mainly in low-flow regions. Therefore, instead of using an optimisation procedure that alters all transmissibilities in the reservoir, they use a thresholding procedure, where negative and very large transmissibilities are replaced with transmissibilities obtained by a local pressure computation

technique. Since the transmissibilities are altered only in low-flow regions, the perturbation will have limited impact on the total flow through the reservoir.

Another class of methods that account for global boundary conditions are multiscale methods. One can think of these methods as a way of upscaling the flow field instead of computing coarse scale reservoir properties (e.g., porosity and permeability) while trying to preserve important trends in the fine-scale flow pattern. Two examples of such methods will be discussed in Section 4.

Other Techniques

There are many upscaling techniques that we have not discussed. For instance, we have not said anything about a class of methods that are based on the homogenisation theory (see e.g., [11, 38, 43]). Homogenisation is a rigorous mathematical theory for asymptotic analysis in periodic structures. A relevant result states that for a periodic medium with permeability $\mathbf{K}(\frac{x}{\varepsilon})$, there exists a constant symmetric and positive definite tensor \mathbf{K}_0 such that the solutions p_ε and $v_\varepsilon = -\mathbf{K}(\frac{x}{\varepsilon})\nabla p_\varepsilon$ to the following elliptic problem

$$-\nabla \cdot \mathbf{K}(\frac{x}{\varepsilon})\nabla p_\varepsilon = q$$

converges uniformly as $\varepsilon \rightarrow 0$ to the corresponding solutions of a homogenised equation of the form

$$-\nabla \cdot \mathbf{K}_0 \nabla p_0 = q.$$

Thus, if we consider each coarse grid-block as a cell in an infinite periodic medium, then homogenisation theory can be used to derive corresponding homogenised tensors for each grid-block. The use of homogenisation theory to upscale grid-block permeability for porous media flow has, however, been a subject of debate. A common question is why one should study a periodic medium when no natural medium is periodic? A natural response to this question is that the homogenisation theory provides the mathematical tools capable of proving the existence and uniqueness of the solution, and gives some verification that the governing equations at the macroscopic level take the same form as (7), which governs porous media flow at the continuous level.

Among other techniques for single-phase upscaling that have not been discussed above are the fast, but less robust, real-space renormalisation techniques [46], and the use of geostatistical methods and Monte-Carlo analysis to generate reservoir descriptions at the reservoir simulation scale directly [33]. Neither have we mentioned how the definition of the coarse-grid geometry can be linked to upscaling through non-uniform coarsening approaches [23], or the use of elastic grids [31]. For a deeper discussion of these, and other related, issues, we ask the reader to consult [19, 55, 60] and the references therein.

3.2 Pseudo Functions for Upscaling the Saturation Equation

Single-phase upscaling alone is often not sufficient to capture large-scale heterogeneity effects in a multi-phase system; for instance, when the flow follows

narrow high-flow channels that penetrate the coarse grid blocks. In such cases, only a small portion of a grid block may be subjected to flow, and the fluid displacement cannot be characterised by single-phase upscaling, even in an averaged sense. Similar problems arise in the presence of long and thin shale barriers or high-permeable layers. Thus, in addition to capturing the macroscopic effect of the rock permeability, we need to incorporate the large-scale effects of the relative permeabilities.

Whereas the problem of upscaling permeability for solving the pressure equation on a coarse scale is fairly well understood, derivation of upscaling techniques for the saturation equation has only been moderately successful. In fact, a robust methodology that can be used to upscale the saturation equation reliably for general heterogeneous formations still does not exist. This is obviously a very difficult task, since we need to trace sharp saturation fronts and at the same time model high-flow channels. Indeed, consider the saturation equation in its simplest form (neglecting gravity and capillary forces)

$$\phi \frac{\partial s}{\partial t} + \nabla \cdot f_w(s)v = 0, \quad (11)$$

and suppose that it is discretised with the following upstream-weighted finite-volume method:

$$\bar{\phi} \frac{s_i^{n+1} - s_i^n}{\Delta t} + \sum_j F^{us}(s_i, s_j) Q_{ij} = 0 \quad \forall V_i \subset \Omega. \quad (12)$$

Here s_i^n is the average saturation in grid-block V_i at time t_n , $\bar{\phi}$ denotes the arithmetic mean of the porosity and Q_{ij} is the total flux across the interface between grid-blocks V_i and V_j obtained from the coarse grid solution of the pressure equation. The numerical fractional flux function F^{us} is given as

$$F^{us}(s_i, s_j) = \begin{cases} f_w(s_i), & \text{if } Q_{ij} \geq 0, \\ f_w(s_j), & \text{if } Q_{ij} < 0. \end{cases}$$

In this scheme, the only way to properly account for the discrepancy between the scales of the coarse grid and the fine grid, is to define special fractional flow functions F_{ij} (so-called pseudo functions) that represent the averaged flux over each coarse grid interface. An intrinsic feature with this approach to two-phase flow upscaling is that pseudo functions depend, not only on the heterogeneous porous media structures, but also on the saturation distribution within the grid block and on the reservoir flow history, which in turn depends on well locations and global boundary conditions. All these effects must be accounted for in the macroscopic reservoir characterisation to achieve robust coarse-scale models.

Pseudo functions can generally be divided in two categories: (vertical) equilibrium pseudo functions and dynamic pseudo functions. Equilibrium pseudo functions have traditionally been used to reduce the dimension of the system

under static flow conditions, e.g., from three to two spatial dimensions and are reminiscent of the steady-state methods discussed below. Dynamic pseudo functions, on the other hand, aim at reproducing subscale effects under dynamic flow conditions by upscaling relative permeabilities or fractional flow functions. Below we briefly review some of the most important techniques that have been proposed for generating dynamic coarse-grid pseudo functions.

Dynamic Pseudo Functions for Two-Phase Flow

The main idea behind dynamic pseudo functions is to do simplified fine-scale flow simulations that mimic the flow patterns of the case that we want to model. An intrinsic feature with this approach to two-phase flow upscaling is that pseudo functions depend, not only on the heterogeneous porous media structures, but also on the saturation distribution within the grid block and on the reservoir flow history, which in turn depends on well locations and global boundary conditions. All these effects must be accounted for in the macroscopic simulation model to achieve reliable results. It is therefore generally not sufficient to perform local, or even extended local, flow simulations [36]. This makes generating reliable pseudo functions without too much computational effort a very challenging task.

The use of dynamic pseudo functions for two-phase flow simulation goes back to Jacks et al. [40], and Kyte and Berry [49]. Their objective was to design relative permeability functions in coarse grid-blocks that account for subgrid flow patterns and scale discrepancies to reduce numerical dispersion. Since then, many different methods have been developed for generating dynamic pseudo functions. Here we will describe two basic types: pseudo functions from individual phase flow-rates and pseudo functions based upon averaged total mobility.

The individual phase flow-rate methods include the methods of Jack et al. [40], the Kyte and Berry method [49], the flux-weighted potential method [34], and the pore-volume weighted method. All methods are based upon the upscaled Darcy’s law and vary only in the way the fine-grid simulations are factored into averaged quantities. Although, these methods have received some criticism for being unreliable (and for being limited to cases where capillary or gravity equilibrium can be assumed at the coarse scale) [9], the individual phase flow-rate methods have been widely used for upscaling of two-phase flow in reservoir simulation. To describe the essential ingredients in these methods, assume that we have a coarse grid overlaying a fine grid, as illustrated in Figure 9. Moreover, we assume that the flow has been computed on the fine grid such that the saturation and pressure history is known in each fine-grid cell along with derived quantities like relative permeability and flow-rates. Based on this, the aim is to define pseudo functions that produce a coarse-grid solution equal to the averaged fine-grid solution. The functions are dynamic in the sense that they are saturation (and pressure) dependent.

The method of Jacks et al. [40] (which is sometimes called the weighted relative permeability method) defines an upscaled relative permeability function for each coarse grid interface by taking a transmissibility-weighted average of relative permeabilities in cells on the upstream side. To be precise, assume that $Q_{ij} > 0$, so that there is a net flow from V_i to V_j and denote by $T_{i,kl}$ the transmissibility for a fine grid interface that lies on $\partial V_i \cap \partial V_j$, i.e., between two fine-grid cells $U_k \subset V_i$ and $U_l \subset V_j$. The upscaled relative permeability associated with $\partial V_i \cap \partial V_j$ is now defined as follows:

$$\bar{k}_{rw,ij} = \frac{\sum_{kl} T_{i,kl} k_{rw,k}(s_k)}{\sum_{kl} T_{i,kl}}.$$

Thus, the sum is taken over all interfaces $(\partial U_k \cap \partial U_l) \subset (\partial V_i \cap \partial V_j)$. The suffix k in $k_{rw,k}$ indicates that there may be different relative permeabilities in each layer (for different rock types). If $Q_{ij} < 0$ then saturations from the corresponding grid cells in V_j are used.

Similar methods to the method of Jacks et al. [40] first average the phase pressure p_α or the phase potential $\Phi_\alpha = p_\alpha - \rho_\alpha g D$ as will be described below and then calculate upscaled pseudo-relative permeabilities by inverting Darcy's law for each phase. For water,

$$\bar{k}_{rw} = -\frac{\bar{\mu}_w \bar{q}_w \Delta x}{\bar{K}(\Delta \bar{p}_w - g \bar{\rho}_w \Delta D)}.$$

Here D is the average depth of the coarse block, \bar{q}_w is the spatial average of the fine-grid water fluxes, and \bar{K} is an upscaled permeability obtained by using a single-phase upscaling technique. The coarse-grid viscosity and density are computed as pore-volume weighted or flow-rate weighted averages.

In the Kyte and Berry method the pseudo water phase-pressures \bar{p}_w are obtained by computing the weighted average along the centre column of the coarse grid block,

$$\bar{p}_w = \frac{\sum_{j=J_1}^{J_2} [K k_{rw} \delta y (p_w - g \rho_w (d - D))]_j}{\sum_{j=J_1}^{J_2} [k_{rw} K \delta y]_j}. \quad (13)$$

In the equation, values with subscript j refer to midpoint (or cell-average) values on the underlying fine-grid along centre column in the coarse block (i.e., $i = \frac{1}{2}(I_1 + I_2)$, see Figure 9). The weighting factor is the phase permeability multiplied by the thickness of each layer. The dynamic behaviour of the pseudo functions follows by tabulating each averaged quantity as a function of the corresponding average saturation. The model can easily be extended to three dimensions, see [59]. By using the procedure on each cell-interface on the coarse grid, one should, in principle, obtain ideal pseudo functions that reproduce the underlying fine-scale simulation in an averaged sense.

Let us now analyse the method in some detail. Notice first that since distinct weighting factors are used for different phases in (13), nonzero capillary




Fig. 9. Schematic of the coarse and the fine grid used in calculation of dynamic pseudo functions.

pressures may be introduced on the coarse grid even if the phase pressures are equal on the fine grid. Moreover, the weighing factors are directionally dependent, leading to pseudo functions that depend on the flow direction (and thus the history of the reservoir). More alarming, however, is the fact that pseudo-relative permeabilities may be negative or have arbitrary large magnitudes. Negative \bar{k}_{rw} values occur if the average flux \bar{q}_w has the same sign as the increment in the phase potential $\Delta\bar{p}_w - g\rho_w g\Delta D$, and if this increment tends to zero, then $|\bar{k}_{rw}| \rightarrow \infty$. The Kyte and Berry method can therefore be very sensitive to small perturbations in the flow conditions.

Alternative methods can be derived by using different weighting factors in to compute the phase pressures. The pore-volume weighted method reads

$$\bar{p}_w = \frac{\sum_{j=J_1}^{J_2} \sum_{i=I_1}^{I_2} [\phi\delta x\delta y(p_w - g\rho_w(d - D))]_{j,i}}{\sum_{j=J_1}^{J_2} \sum_{i=I_1}^{I_2} [\phi\delta x\delta y]_{j,i}},$$

and similarly, the flux-weighted potential method [34]

$$\bar{k}_{rw} = -\frac{\mu\bar{q}_w\Delta x}{K_x\Delta\bar{\Phi}_w}, \quad \bar{\Phi}_w = \frac{\sum_{j=J_1}^{J_2} [q_w\bar{\Phi}_w]_j}{\sum_{j=J_1}^{J_2} [q_w]_j},$$

where the sums are taken over the fine-grid column in the centre of the coarse block.

To avoid some of the problems associated with the individual phase flow-rate methods, Stone [58] suggested upscaling the fractional-flow function $f_o(S)$ based on averaging the total mobilities λ at the fine scale,

$$\bar{f}_o = \frac{\sum_{j=J_1}^{J_2} [q_t f_o]_j}{\sum_{j=J_1}^{J_2} [q_t]_j}, \quad \bar{\lambda} = \frac{\sum_{j=J_1}^{J_2} [T_x \lambda]_j}{\sum_{j=J_1}^{J_2} [T_x]_j}.$$

Here the sum is taken over the fine-grid column on the upstream side of the coarse-grid interface. Stone's formula assumes constant pressure drop in the layers and neglects gravity and capillary pressure. His method for is therefore usually inadequate in cases with significant gravity or capillary pressure effects, or in cases with large local variations in total mobility.

Several authors [8, 35, 47] have proposed methods similar to e.g., Stone’s method, with “better” averaging of the total mobility. One possibility is to do local potential calculations with e.g., periodic boundary conditions and invert the solution with respect to the phase mobilities. These total mobility methods are more robust than the Kyte and Berry method in the sense that infinite values can be avoided and negative values occur less frequently (but can still occur if the net flows of the two phases are in opposite directions). In summary, both the Kyte and Berry method and the method proposed by Stone (and variants of the respective methods) have their drawbacks. The Kyte and Berry method is computationally expensive, and may be highly sensitive to small perturbations in the flow conditions. Total mobility methods like Stone’s method are also computationally expensive, but at least these methods tend to be more robust than the Kyte and Berry method.

3.3 Volume Averaged Equations

As an alternative to using pseudo functions, a more general framework based on using higher moments to develop volume-averaged equations (VAE) has been proposed by several authors, e.g., [22, 27, 50]. In this approach the basic idea is to express the unknown quantities in terms of average and fluctuating components. To illustrate, consider the simplified saturation equation (11) with unit porosity $\phi = 1$. Furthermore, for any given variable $\nu(x)$, write

$$\nu(x) = \bar{\nu} + \tilde{\nu}(x),$$

where $\bar{\nu}$ denotes a volume-averaged quantity (constant within the averaging region), and $\tilde{\nu}(x)$ denotes a spatially fluctuating quantity with zero mean in the averaging region. Substituting now the volume-average expansions for s , v , and f_w into (11), and then averaging the resulting equations gives

$$\bar{s}_t + \tilde{s}_t + \bar{v} \cdot \nabla \overline{f_w} + \bar{v} \cdot \nabla \tilde{f}_w + \tilde{v} \cdot \nabla \overline{f_w} + \tilde{v} \cdot \nabla \tilde{f}_w = 0. \quad (14)$$

Averaging each term in this equation, noting that the average of terms with only one fluctuating component vanishes, we obtain the following volume averaged equation for \bar{s} :

$$\bar{s}_t + \bar{v} \cdot \nabla \overline{f_w} + \overline{\tilde{v} \cdot \nabla \tilde{f}_w} = 0 \quad (15)$$

By subtracting (15) from (14) one obtains a corresponding equation for the fluctuating part \tilde{s} :

$$\tilde{s}_t + \bar{v} \cdot \nabla \tilde{f}_w + \tilde{v} \cdot \nabla \overline{f_w} + \tilde{v} \cdot \nabla \tilde{f}_w = \overline{\tilde{v} \cdot \nabla \tilde{f}_w}. \quad (16)$$

The equation for the fluctuating component can be used to generate equations for various moments of the fine-scale equations, see e.g., [21].

Assuming a unit mobility ratio, Efendiev et al. [25, 27] derive a single coarse-grid equation from (15) on the form

$$\frac{\partial \bar{s}}{\partial t} + \nabla \cdot G(x, \bar{s}) = \nabla \cdot D(x, t) \nabla \bar{s}, \quad (17)$$

where $G(x, \bar{s}) = \bar{v} f(\bar{s})$ and $D(x, t)$ is a history dependent function that models the cross term $\tilde{v} \cdot \nabla \tilde{f}_w$. If we compare this equation with (11) we observe that (17) contains an extra diffusion term. The general form of the coarse-scale equation (17) has been obtained by many authors, e.g., [21, 22, 29, 50], and the history dependence in the diffusion term has been rigorously verified using homogenisation theory, see e.g., [29, 43]. This shows that the correct form of the coarse-scale saturation equation does not, unlike the coarse-scale pressure equation, take the same form as the equation that models the phase transport at the continuous level. This observation gives some justification for developing coarse-scale models that attempt to incorporate the history dependence.

Unfortunately, the implementation of VAEs into a simulator can be a bit complicated due to the nonlocal memory effect. As a remedy for this difficulty, Efendiev and Durlafsky [26] introduced a generalised convection-diffusion model for modelling subgrid transport in porous media. Here they observe that the nonlocal dispersivity in (17) imparts a somewhat convective character to the model. Therefore, to eliminate the history-dependence in (17), they propose a generalised convection-diffusion model where they introduce a convective correction term so that the convective term in (17) is on the form

$$G(x, \bar{s}) = \bar{v} f(\bar{s}) + m(x, \bar{s}).$$

The convective correction term $m(x, \bar{s})$ is determined from local subgrid computations, as is the diffusion term $D(x, \bar{s})$.

Durlafsky has investigated the relation between upscaling methods based on the VAE methodology and existing upscaling techniques for two-phase flow, with an emphasis on pseudo-relative permeability generation [9, 49, 53, 58], the non-uniform coarsening approach⁵ [23], and the use of higher moments of the fine-scale variables to correlate upscaled relative permeabilities [21]. It is illustrated that some of the benefits and limitations with the various approaches can be understood and quantified with respect to volume averaged equations. Specifically, it is demonstrated that for flows without gravity and capillary forces, the VAE approach shows a higher degree of process independence than the traditional pseudo-relative permeability methods. However, to our knowledge, all published literature on VAEs and so-called generalised convection-diffusion models for two-phase flow neglect capillary and gravity forces. In addition, the methodology has for the most part been tested on cases with unit mobility ratio [22, 21, 24, 27]. In recent work, Efendiev and Durlafsky [25, 26] make an effort to generalise the VAE approach to two-phase

⁵The non-uniform coarsening method tries to generate grids that are coarsely gridded in low-flow regions and more finely gridded in high-flow regions. Thus, rather than modelling subgrid effects explicitly, the non-uniform coarsening methodology attempts to minimise errors caused by subgrid effects by introducing higher grid resolution in regions where subgrid terms would otherwise be important.

flow with general mobility ratios, and obtain promising results. However, further research is needed to validate the method and investigate if it can be extended to include capillary and gravity forces.

3.4 Steady-State Equilibrium Upscaling for Two-Phase Flow

Dynamic pseudo methods are numerical in nature and tend to ignore the underlying physics. The focus is on adjusting parameters to match a numerical solution, rather than basing the upscaling on effective physical properties. Steady-state upscaling (see e.g., [56]) is an alternative where the essential assumption is steady-state flow, i.e., that the time derivative in equation (1) is zero. Away from the well, this assumption leads us to a steady-state equation for each phase on the form on the following form:

$$\nabla \cdot \left[\frac{\mathbf{K}k_{ri}}{\mu_i} (\nabla p_i - \rho_i G) \right] = 0. \quad (18)$$

It may seem very restrictive to impose steady-state flow. However, away from saturation fronts and in regions with low flow we may have small saturation changes with respect to time, and small changes in fractional flow. Disregarding the temporal derivative in (1) can therefore be justified in these types of flow regions.

To derive effective relative permeabilities from equation (18) we first need to determine the saturation distribution inside each coarse grid block. This will be done from various physical principles, depending on the flow regime. An upscaled relative permeability that corresponds to the prescribed subgrid saturation distribution is then obtained by upscaling the phase permeability $\mathbf{K}k_{ri}$ using a pressure-computation single-phase upscaling technique, and dividing the result with the corresponding upscaled permeability tensor \mathbf{K} . Thus, steady-state relative permeability curves are defined by the following formula:

$$\bar{k}_{rw} = \bar{\mathbf{K}}^{-1} \overline{\mathbf{K}k_{rw}}. \quad (19)$$

This technique may thus be viewed as an extension to two-phase flows of the pressure computation technique for single-phase upscaling.

To model the subgrid saturation, additional assumptions are made. The precise assumptions will depend on the scale that we wish to model, the flow process (e.g., the magnitude of the flow velocity) as well as an assessment of which physical effects that are important. Assuming for the moment that we have a way of extracting an appropriate saturation field on the subgrid from the coarse scale solution, it is clear that also other saturation-dependent physical parameters, such as the capillary pressure, may be upscaled. Thus, as opposed to dynamic pseudo functions, we see that these methods do not seek numerical devices that reproduce subscale flow effects. Instead, steady-state upscaling techniques seek effective physical properties and therefore have a more physical motivation. On the other hand, steady-state methods do not

correct for grid effects and numerical dispersion like dynamic pseudo functions tend to.

Steady-state upscaling methods may be put into three categories that are linked to the balance of three forces: viscous forces, capillary forces, and gravitational forces which give rise to flows q_v , q_c , and q_g (see (5) and the subsequent discussion). Understanding the balance of these forces may indicate which method to warrant. It is often not clear how to quantify these three forces for a particular modelling problem, and they will usually vary both with time and space. However, we can try to investigate the role of the different forces by comparing the respective contributions to the saturation equation

$$\phi \partial_t s_w = -\nabla \cdot [q_v + q_c + q_g]. \quad (20)$$

For certain types of models it may be possible to choose representative values for q_v , q_c and q_g . Thus, if we normalise them with respect to their sum and place them in a ternary diagram, then the position in the diagram may serve as a guideline deciding on which steady-state method to apply, see, e.g., [57]. Another possibility is to measure the balance of these three forces using laboratory experiments. This type of experiment is usually performed on core samples with constant saturation fraction injected through an inlet boundary, and pressure controlled production at an outlet boundary, where the outlet boundary is at the opposite face of the inlet boundary. By comparing the observed flow-rates at the outlet boundary with corresponding flow-rates obtained from numerical simulation studies, possibly through some tuning of relevant parameters, it is possible to say something about the importance of including the various forces in the upscaled model. Note that care need to be exercised when using laboratory experiments. The balance of forces are sensitive to scale, and also to whether reservoir conditions are used during a core flooding experiment. We now present how to determine the subgrid saturations for the different steady-state methods.

Capillary Equilibrium Methods

In these steady-state upscaling methods one assumes, in addition to the steady-state assumption, that capillary pressure is constant in space and time. In practise, approximately constant capillary pressure occurs mainly in regions with very low flow-rates. This implies that viscous and gravity forces are dominated by capillary forces. Viscous forces, and sometimes also gravity forces, are therefore neglected in capillary equilibrium steady-state methods.

In capillary equilibrium steady-state methods, the saturation in each grid-cell within the block to be upscaled is determined by inverting the capillary pressure curve. Indeed, for a two-phase system the capillary pressure is modelled as a strictly monotone function of one of the phases, so it is invertible on its range. The capillary pressure function is often scaled by the so-called Leverett J-function. This function takes the following form:

$$J(s_w) = \frac{p_c \sqrt{\frac{\mathbf{K}}{\phi}}}{\sigma \cos(\theta)}, \quad (21)$$

where σ is the interfacial tension between the phases, and θ is the contact angle between the phases. The contact angle and interfacial tension are usually measured in a laboratory. The Leverett J-function is usually a known (tabulated) function of s_w . Now, since the capillary pressure is a known constant within each grid block and \mathbf{K} and ϕ are known subgrid quantities, we can invert (21) to obtain the subgrid saturation distribution

$$s_w = J^{-1} \left(\frac{\bar{p}_c \sqrt{\frac{\mathbf{K}}{\phi}}}{\sigma \cos(\theta)} \right). \quad (22)$$

With the J-function scaling, the capillary pressure curve for each grid block will be scaled according to the ratio between permeability and porosity, and tend to give large saturation contrasts in the model.

Once the subgrid saturation is known, we solve the steady-state equation (18) for each phase, and use (19) to derive effective relative permeabilities. Gravity forces can be included in the model by replacing p_c in (22) with $\varrho = p_c + (\rho_w - \rho_o)gz$, where z is vertical position in the model. This follows by observing that in absence of viscous forces we have $v_w = f_w \mathbf{K} \lambda_o \nabla \varrho$. Thus, conceptually the mapping defined by ϱ allows us to include gravity forces in the capillary pressure function. However this approach should be used only if capillary forces are dominant since the Leverett J-function does not model effects of gravity. Finally, we remark that the assumption that viscous forces are negligible can be relaxed, as is shown in [6].

Steady-State Upscaling Methods for Viscous Dominated Flows

The basic assumption in the capillary equilibrium methods was that the flow was dominated by capillary forces, or equivalently that the Darcy velocity v is so small that the viscous forces can be neglected. At the other end of the rate scale we have the *viscous limit*, which can be interpreted as the limit when capillary and gravitational forces tend to zero relative to the viscous forces. This can be the case in, for instance, nearly horizontal high permeable layers with restricted vertical extent. For a more in depth discussion on the applicability of viscous limit steady-state the reader may consult [30]. By convention, one often refers to this as the limit when the Darcy velocity goes to infinity.

In the viscous-limit approach to steady-state upscaling methods, one neglects gravity and capillary forces and assumes constant fractional flow. In addition, it is assumed that the pressure gradient across a grid-block is the same for both phases. If k_{ro} is nonzero, we now can divide the flow-rate of the phases with each other (given by Darcy's law) to derive

$$\frac{v_w}{v_o} = \frac{\frac{-Kk_{rw}}{\mu_w} \nabla p}{\frac{-Kk_{ro}}{\mu_o} \nabla p} = \frac{k_{rw}\mu_o}{k_{ro}\mu_w}.$$

The relative permeability curves are assumed to be known functions of the phase saturations and rock type (i.e., functions of x). With the additional equation that saturations add to one, solving for saturations is just a matter of solving two equations with two unknowns $s_w(x)$ and $s_o(x)$,

$$s_w + s_o = 1, \quad k_{rw}(s_w, x) = \left(\frac{v_w \mu_w}{v_o \mu_o} \right) k_{ro}(s_o, x).$$

Note that saturations obtained are unique when the relative pressure curves are strictly monotonic. Moreover, if we have a single set of relative permeability curves, the subgrid saturations will be constant within the coarse block. As in the capillary equilibrium methods, effective relative permeabilities are obtained from the subgrid saturation distribution by using (19) and an effective permeability tensor $\bar{\mathbf{K}}$ is obtained by solving the steady-state single-phase equation.

It should be noted that for numerical implementation a strategy presented in [48] should be preferred. Rather than upscaling each phase separately, it is suggested to upscale the total mobility λ_t , and at the final step calculate the individual relative permeabilities using fractional flow. This approach has two advantages, the first being that you only need to solve the steady-state single phase equation once, the other being that you reduce contrast in the model, softening the numerical properties of the equation.

General Steady-State Methods

The two steady-state upscaling methods described above apply to cases with either very low or very high flow-rates. This may apply to certain regions of an oil-reservoir, but one will always have regions with somewhat moderate flow-rates. Moreover, many grid blocks will have some cells with low permeability and some blocks with high permeability. Hence, one will often have both high and low flow-rates inside the same grid block. In these situations, neither the capillary equilibrium method nor the viscous limit method tend to perform well. Still, the basic steady-state upscaling methodology that involves upscaling steady-state equations of the form (18) and using (19) to derive effective relative permeabilities may still be used. However, this requires full flow simulations for each grid block and is therefore computationally expensive. Moreover, since we can not exploit physical principles to derive upscaled parameters, we need to determine an initial saturation distribution for the simulations, as well as fluid injection rates on the boundary. On the upside we have that no physical assumptions are required other than steady-state, so that effects from gravity, phase transfer, and compressibility can be included.

Final Remarks on Two-Phase Upscaling

As with all mathematical tools for handling simulation of flow, the value of upscaling techniques will ultimately be judged by its success in applications to practical problems. Upscaling two-phase flow reliably has proved to be a very difficult task. In fact, although pseudo methods are widely used, Barker and Thibeau [9] argue that existing pseudo methods can only be used to upscale two-phase flow reliably for cases where gravity equilibrium (gravity forces dominate capillary and viscous forces) or capillary equilibrium can be assumed at the coarse-grid scale. One of the reasons why it is so hard to upscale reservoir flow scenarios reliably is that different types of reservoir structures require different upscaling methodologies. Hence, in principle one should divide grid blocks into different categories that each represents a particular type of representative elementary volumes. This implies that upscaling should ideally be linked to gridding. Of the methods discussed in this section, only pseudo methods tend to correct for grid effects, meaning that simulation results differ if you change the grid resolution. Relative permeability curves obtained using pseudo methods may therefore be effective for reproducing the production history of a field, but may perform poorly if the production strategy of the field is changed (e.g., if a new well is drilled). Conclusively, pseudo and steady-state methods both have their drawbacks. Pseudo-functions have been used by the industry for quite some time, so their limitations are well-known. Methods based on the VAE methodology are still immature for real field applications.

4 Multiscale Methods for Solving the Pressure Equation

Several types of multiscale methods have recently begun to emerge in various disciplines of science and engineering. By definition, multiscale methods is a label for techniques that model physical phenomena on coarse grids while honouring small-scale features that impact the coarse-grid solution in an appropriate way. Mathematical-oriented multiscale methods are often geared toward solving partial differential equations with rapidly varying coefficients. These methods try to incorporate subgrid information by utilising solutions of local flow problems to build a set of equations on a coarser scale. This is similar to the approach of flow-based upscaling methods, but the multiscale methods are different in the sense that they ensure consistency of the coarse-scale equations with the underlying differential operator. In practice this means that the solution of the coarse-scale system is consistent with the differential operator also on the finer scale on which the local flow problems are solved, and therefore gives an accurate solution on this scale as well.

To accomplish this, the multiscale methods may express the fine-scale solution as a linear combination of basis functions, where the basis functions are solutions to local flow problems. Here we describe and discuss two methods




Fig. 10. Schematic reservoir model with a coarse grid superimposed on the fine grid.

of this type that are designed for solving pressure equations arising in reservoir simulation. The methods we shall consider are the multiscale finite-volume method (MsFVM) [41, 42], and the multiscale mixed finite-element method (MsMFEM) [17, 2]. Both methods generate mass-conservative velocity fields on the subgrid scale, while solving only a coarse-scale system globally.

The starting point for both methods is a subdivision of the reservoir into polyhedral grid-blocks, with a coarse grid superimposed on the fine grid, e.g., as shown in Figure 10. Usually each coarse grid-block is a connected union of fine grid-blocks, but the coarse grid may be any connected non-overlapping partition of the reservoir, provided quantities may be properly mapped from the underlying fine grid. We shall assume that the reservoir data (\mathbf{K} and λ) are given on the fine grid, and since this grid must be suitable for numerical discretisation of (23) using standard numerical techniques, it can not be completely arbitrary. However, in the multiscale methods we use the fine grid only to solve local problems. The properties that enter the coarse-scale equations are extracted from these local solutions alone, hence the multiscale formulations are in a sense independent of coarse grid geometry.

As a consequence of this, we can allow very general coarse grids. For instance, we can allow grids that are locally structured, but globally unstructured. This convenient feature allows us to use the subgrid to model physical features such as fractures and faults, or to tune the grid geometry to advanced well architectures. Moreover, few restrictions are placed on the geometry of each coarse grid-block. This is particularly true for the MsMFEM, which has demonstrated accurate results with almost arbitrarily shaped grid-blocks [1]. Such a flexibility makes multiscale methods amenable for solving problems that demand unstructured grids, or grids with high grid-block aspect ratios.

Before describing each of the methods in more detail, we specify the model equations we will work with. For simplicity, we consider the elliptic pressure equation only, i.e., we assume incompressible flow. The methods can be extended to solve the parabolic equation arising from compressible flow using, for instance, an implicit Euler or the Crank-Nicholson method to discretise the time derivative. Moreover, we will disregard gravity effects, since these are straightforward to include, but complicate notation. After these simplifi-

cations, the pressure equation for two-phase flow (4) reads

$$v = -\mathbf{K}\lambda\nabla p, \quad \nabla \cdot v = q, \quad (23)$$

and we assume this is to be solved in a domain Ω subject to no-flow boundary conditions $v \cdot n = 0$ on $\partial\Omega$.

4.1 The Multiscale Finite-Volume Method

The MsFVM, introduced by Jenny et al. [41, 42], is essentially a finite-volume scheme that extracts a mass-conservative velocity field on a fine scale from a coarse-grid solution. The method employs numerical subgrid calculations (analogous to those in [39]) to derive a multi-point finite-volume stencil for solving (23) on a coarse grid. The method then proceeds and reconstructs a mass-conservative velocity field on a fine grid as a superposition of local subgrid solutions, where the weights are obtained from the coarse-grid solution.

Computation of the Coarse-Scale Solution

The derivation of the coarse-scale equations in the MsFVM is essentially an upscaling procedure for generating coarse-scale transmissibilities. The first step is to solve a set of homogeneous boundary-value problems of the form

$$-\nabla \cdot \mathbf{K}\lambda\nabla\phi_i^k = 0, \quad \text{in } R, \quad \phi_i^k = \nu_i^k, \quad \text{on } \partial R, \quad (24)$$

where R are so-called interaction regions as illustrated in Figure 11 and ν_i^k are boundary conditions to be specified below. The subscript i in ϕ_i^k denotes a corner point in the coarse grid (x_i in the figure) and the superscript k runs over all corner points of the interaction region (x^k in the figure). Thus, for each interaction region associated with e.g., a hexahedral grid in three dimensions, we have to solve a total of eight local boundary-value problems of the form (24). The idea behind the MsFVM is to express the global pressure as a superposition of these local pressure solutions ϕ_i^k . Thus, inside each interaction region R , one assumes that the pressure is a superposition of the local subgrid solutions $\{\phi_i^k\}$, where k ranges over all corner points in the interaction region (i.e., over the cell centres of the coarse grid-blocks).

We define now the boundary conditions ν_i^k in (24). At the corner points of the interaction region, the boundary condition ν_i^k satisfies $\nu_i^k(x^l) = \delta_{kl}$, where δ_{kl} is the Kronecker delta function. Assuming we are in three dimensions, the corner-point values $\nu_i^k(x^l)$ are extended to the edges of the interaction region by linear interpolation. To specify the boundary conditions on each face F of the interaction region, one solves a two-dimensional boundary problem of the form

$$-\nabla \cdot \mathbf{K}\lambda\nabla\nu_i^k = 0 \quad \text{in } F, \quad (25)$$

with the prescribed piecewise-linear boundary conditions on the boundary of F (the edges of the interaction region that connect the corner points). In two




Fig. 11. The shaded region represents the interaction region R for the MsFVM, where x_i denotes corner points and x^k the midpoints of the coarse grid-blocks K_k . The midpoints x^k are the corner points of the interaction region R .

dimensions it is not necessary to interpolate between corner points, and the boundary conditions $\nu_i^k(x)$ are obtained by solving (25) subject to the corner point values along the edges of the interaction region.

Having computed the local pressure solutions ϕ_i^k , we now show how to obtain a global solution as a superposition of the local subgrid solutions. To this end, we observe that the cell centres x^k constitute a corner point for eight interaction regions (in a regular hexahedral grid). Moreover, for all corner points x_i of the coarse grid, the corresponding boundary conditions ν_i^k for the different pressure equations coincide on the respective faces of the interaction regions that share the corner point x^k . This implies that the base function

$$\phi^k = \sum_i \phi_i^k \quad (26)$$

is continuous (in a discrete sense). In the following construction, the base functions defined in (26) will serve as building blocks that are used to construct a global “continuous” pressure solution.

Thus, define now the approximation space $V^{\text{ms}} = \text{span}\{\phi^k\}$ and observe that all base functions vanish at all but one of the grid-block centres x^k . This implies that, given a set of pressure values $\{p^k\}$, there exists a unique extension $\{p^k\} \rightarrow p \in V^{\text{ms}}$ with $p(x^k) = p^k$. This extension is defined by

$$p = \sum_k p^k \phi^k = \sum_{i,k} p^k \phi_i^k. \quad (27)$$

We derive now a multi-point finite-volume stencil by assembling the flux contribution across the grid-block boundaries from each base function. Thus, let

$$f_{k,l} = - \int_{\partial K_l} n \cdot \mathbf{K} \lambda \nabla \phi^k ds$$

be the local flux out of grid-block K_l induced by ϕ^k . The MsFVM for solving (23) then seeks constant grid-block pressures $\{p^k\}$ satisfying

$$\sum_k p^k f_{k,l} = \int_{K_l} q dx \quad \forall l. \quad (28)$$

We now explain how to reconstruct a mass-conservative velocity field on the underlying subgrid from the coarse-grid solution.

Reconstruction of a Conservative Fine-Scale Velocity Field

To reconstruct a mass-conservative velocity field on a fine scale, notice first that the expansion (27) produces a mass-conservative velocity field on the coarse scale. Unfortunately, this velocity field will not preserve mass across the boundaries of the interaction regions. Thus, to obtain a velocity field that also preserves mass inside the coarse grid-blocks, one solves

$$v_l = -\mathbf{K}\lambda\nabla p_l, \quad \nabla \cdot v_l = \frac{1}{|K_l|} \int_{K_l} q dx \quad \text{in } K_l, \quad (29)$$

with boundary conditions obtained from (27), i.e.,

$$v_l = -\mathbf{K}\lambda\nabla p \quad \text{on } \partial K_l, \quad (30)$$

where p is the expanded pressure defined by (27).

Hence, in this approach the subgrid fluxes across the coarse-grid interfaces that we obtain from p are used as boundary conditions for a new set of local subgrid problems. If these subgrid problems are solved with a conservative scheme, e.g., a suitable finite-volume method, then the global velocity field $v = \sum_{K_l} v_l$ will be mass conservative. Moreover, since the fluxes correspond to interfaces that lie in the interior of corresponding interaction regions, it follows that the boundary condition (30) is well defined. Note, however, that since the subgrid problems (29)–(30) are solved independently, we lose continuity of the global pressure solution, which is now defined by $p = \sum_{K_l} p_l$.

It is possible to write the new global fine-scale solutions p and v as linear superpositions of (a different set of) local base functions, see [41] for details. This may be convenient when simulating two-phase flows with low mobility ratios. This is because low mobility ratios lead to a slowly varying mobility field. As a result, it is sometimes possible to achieve high accuracy with a unique set of base functions. Thus, by representing p and v as a linear superposition of local base functions, solving the subgrid problems (24)–(25) and (29)–(30) becomes part of an initial preprocessing step only.

It is worth noting that the present form of the MsFVM, which was developed by Jenny et al. [41], does not model wells at the subgrid scale. Indeed, the source term in (29) is equally distributed within the grid-block. Thus, in order to use the induced velocity field to simulate the phase transport, one has to treat the wells as a uniform source within the entire well block. However, to get a more detailed representation of flow around wells one needs only to replace (29) in grid blocks containing a well with the following equation:

$$v_l = -\mathbf{K}\lambda\nabla p_l, \quad \nabla \cdot v_l = q \quad \text{in } K_l.$$

This completes the description of the MsFVM. The next section is devoted to a multiscale method with similar properties that is based on a so-called mixed formulation of the elliptic pressure equation (23); see [3].

4.2 A Multiscale Mixed Finite-Element Method

For finite-volume methods and finite-element methods the velocities are derived quantities of the unknown pressures, computed using some kind of numerical differentiation. In this section we consider a method that obtains the velocity directly. The underlying idea is to consider both the pressure and the velocity as unknowns and express them in terms of basis functions. The corresponding method is called a mixed finite-element method. In the lowest-order mixed method one uses piecewise-constant scalar basis functions for the pressure and piecewise-linear basis functions for the velocity (this is usually referred to as the Raviart–Thomas basis functions, see [54]).

In mixed finite-element discretisations of elliptic equations on the form (23), one seeks a solution (p, v) to the mixed equations

$$\int_{\Omega} u \cdot (\mathbf{K}\lambda)^{-1}v \, dx - \int_{\Omega} p \nabla \cdot u \, dx = 0, \quad (31)$$

$$\int_{\Omega} l \nabla \cdot v \, dx = \int_{\Omega} ql \, dx, \quad (32)$$

in a finite-dimensional product space $U \times V \subset L^2(\Omega) \times H_0^{1,\text{div}}(\Omega)$. If the subspaces $U \subset L^2(\Omega)$ and $V \subset H_0^{1,\text{div}}(\Omega)$ are properly balanced (see, e.g., [12, 13, 14]), then p and v are defined (up to an additive constant for p) by requiring that (31)–(32) holds for all $(l, u) \in U \times V$.

In MsMFEMs one constructs a special approximation space for the velocity v that reflects the important subgrid information. For instance, instead of seeking velocities in a simple approximation space spanned by base functions with linear components, as in the Raviart–Thomas method, one computes special multiscale base functions ψ in a manner analogous to the MsFVM, and defines a corresponding multiscale approximation space by $V^{\text{ms}} = \text{span}\{\psi\}$. An approximation space for the pressure p that reflects subgrid structures can be defined in a similar manner. However, there are reasons for not doing so.

First, the pressure equation (23) models incompressible flow. This means that the pressure is never used explicitly in the flow simulation, except possibly to determine well-rates through the use of an appropriate well-model. It is therefore often sufficient to model pressure on the coarse scale as long as we can obtain a detailed velocity field without knowing how the pressure behaves at the subgrid scale. This is one of the nice features with mixed finite-element methods. Indeed, by treating the pressure and velocities as separate decoupled variables, we gain some freedom in terms of choosing the resolution of the two approximation spaces U and V more independently. In other




Fig. 12. Schematic of the coarse and fine grid for the MsMFEM. The shaded region denotes the support of the velocity base function associated with the edge between the two cells K_i and K_j .

words, the computational effort can be spent where it is most needed. On the other hand, the approximation spaces can not be chosen arbitrarily. Indeed, the convergence theory for mixed finite element methods, the so-called *Ladyshenskaja–Babuška–Brezzi* theory (see [12, 13, 14]) states that the approximation spaces must satisfy a relation called the *inf-sup* condition, or the LBB condition. By defining a multiscale approximation space also for the pressure variable, it can be quite difficult to show that this relation holds. Therefore, in the following we assume, unless stated otherwise, that U is the space of piecewise constant functions, i.e.,

$$U = \{p \in L^2(\Omega) : p|_K \text{ is constant for all } K \in \mathcal{K}\}.$$

Observe that this space is spanned by the characteristic functions with respect to the coarse grid-blocks: $U = \text{span}\{\chi_m : K_m \in \mathcal{K}\}$ where

$$\chi_m(x) = \begin{cases} 1 & \text{if } x \in K_m, \\ 0 & \text{else.} \end{cases}$$

Next, we We define the approximation space V for the velocity.

Approximation Space for the Velocity

Consider a coarse grid that overlays a fine (sub)grid, for instance as illustrated in Figure 12. For the velocity we associate one vector of base functions with each non-degenerate interface γ_{ij} between two neighbouring grid-blocks K_i and K_j . To be precise, for each interface γ_{ij} we define a base function ψ_{ij} by

$$\psi_{ij} = -\mathbf{K}\lambda\nabla\phi_{ij}, \quad \text{in } K_i \cup K_j, \quad (33)$$

where ϕ_{ij} is determined by




Fig. 13. The figure depicts the x -component of two MsMFEM base functions for an interface between two rectangular (two-dimensional) grid-blocks that are not penetrated by a well. The base function to the left corresponds to homogeneous permeability, whereas the base function to the right corresponds to random coefficients.

$$(\nabla \cdot \psi_{ij})|_{K_i} = \ell(x) / \int_{K_i} \ell(x) dx, \quad (34)$$

$$(\nabla \cdot \psi_{ij})|_{K_j} = -\ell(x) / \int_{K_j} \ell(x) dx. \quad (35)$$

with no-flow boundary conditions along the edges $\partial K_i \cup \partial K_j \setminus \gamma_{ij}$.

The function $\ell(x)$ can be defined in various ways. One option that was proposed by Chen and Hou [17] is to simply choose $\ell(x) = 1$. This option produces mass-conservative velocity fields on the coarse scale, but the subgrid velocity field is mass conservative only if we treat the wells as a uniform source term within grid-blocks containing wells. To model flow around wells at the subgrid scale, one can for instance choose $\ell(x) = q(x)$ in coarse blocks penetrated by an injection or production well. Another possibility is to choose $\ell(x) = 1$ and then to reconstruct a velocity field that models near-well flow patterns from the coarse grid solution. This final option is analogous to the way one models flow around wells with the MsFVM.

We are here primarily interested in getting a velocity field that can be used to model flows in heterogeneous porous media on a subgrid scale. To this end we must solve the subgrid problems (33)–(35) with a conservative scheme, for instance the Raviart–Thomas mixed finite-element method. Moreover, it is generally acknowledged that it is important to model the flow around wells correctly to obtain reliable production forecasts from reservoir simulations. In the current MsMFEM we have the possibility to do this in a rigorous manner by incorporating near-well flow patterns in the base functions. We will therefore choose $\ell(x) = 1$ away from the wells and $\ell(x) = q(x)$ in grid blocks penetrated by wells. As discussed in [2, 5], it is not necessary to change

the approximation space for p even if we change $l(x)$ in blocks penetrated by wells.

Figure 13 shows the x -component of the velocity base functions for the MsMFEM obtained for two cases with homogeneous and random permeability, respectively. We see that the base function corresponding to random coefficients fluctuates rapidly and hence reflects the fine-scale heterogeneous structures. Note also that the base functions ψ_{ij} will generally be time dependent since they depend on λ , which is time dependent through $s_w(x, t)$. This indicates that one has to regenerate the base functions for each time step. However, for nearly incompressible fluids, the total mobility λ usually varies significantly only in the vicinity of propagating saturation fronts. It is therefore usually sufficient to regenerate a small portion of the base functions at each time step, see [2]. Similar observations have been made for the MsFVM, see [41].

Computing the Coarse Scale Solution

In the MsFVM the only information from the subgrid calculations that was explicitly used in the coarse grid equations was the flux contribution across the coarse grid interfaces from the respective base functions. This implies that a lot of potentially valuable information is disregarded in the coarse scale equations. For the MsMFEM, all information (apart from the subgrid pressure solutions) is exploited and enters directly into the coarse-scale system. Indeed, the MsMFEM seeks

$$p \in U, v \in V^{ms} \quad \text{such that (31)–(32) holds for } \forall l \in U, \forall u \in V^{ms}.$$

This leads to a linear system of the form

$$\begin{bmatrix} \mathbf{B} & -\mathbf{C}^T \\ \mathbf{C} & \mathbf{0} \end{bmatrix} \begin{bmatrix} \mathbf{v} \\ \mathbf{p} \end{bmatrix} = \begin{bmatrix} \mathbf{0} \\ \mathbf{f} \end{bmatrix}, \quad (36)$$

where

$$\begin{aligned} \mathbf{B} &= \left[\int_{\Omega} \psi_{ij} \cdot (\lambda \mathbf{K})^{-1} \psi_{kl} \, dx \right], \\ \mathbf{C} &= \left[\int_{K_m} \text{div}(\psi_{kl}) \, dx \right], \\ \mathbf{f} &= \left[\int_{K_m} f \, dx \right]. \end{aligned}$$

The linear system (36) is indefinite, which generally makes it harder to solve than the symmetric, positive definite systems arising from, e.g., the two-point flux approximation (6). This is not particular to the multiscale version of the mixed finite-element method, but a well-known characteristic of mixed

finite-element methods in general. Fortunately it is possible to reformulate the discrete form of the mixed equations (31)–(32) to eventually obtain a symmetric, positive definite linear system by applying a technique usually referred to as hybridisation [14].

The idea is to first remove the constraint that the normal velocity must be continuous across block interfaces and integrate (23) to get a weak form containing jump terms at block boundaries. Continuity of the normal component is then reintroduced by adding an extra set of equations, where the pressure at the interfaces plays the role of Lagrange multipliers. This procedure does not change the velocity solution, nor the grid cell pressures, but enables the recovery of pressure values at the cell interfaces, in addition to inducing the desired change in structure of the linear system.

In our multiscale setting, the mixed-hybrid problem is to find

$$(p, v, \pi) \in (U \times \tilde{V}^{ms} \times \Pi),$$

where

$$U \subset L^2(\Omega), \quad \tilde{V}^{ms} \subset (L^2(\Omega))^2, \quad \text{and } \Pi \subset L^2(\{\partial K\}),$$

such that

$$\int_{\Omega} u \cdot (\mathbf{K}\lambda)^{-1} v \, dx - \int_{\Omega} p \nabla \cdot u \, dx + \sum_{\kappa \in \{\partial K\}} \int_{\kappa} [u \cdot n] \pi \, ds = 0, \quad (37)$$

$$\int_{\Omega} l \nabla \cdot v \, dx = \int_{\Omega} ql \, dx, \quad (38)$$

$$\sum_{\kappa \in \{\partial K\}} \int_{\kappa} [v \cdot n] \mu \, ds = 0, \quad (39)$$

holds for all $(l, u, \mu) \in (U \times \tilde{V}^{ms} \times \Pi)$. The square brackets denote the jump in a discontinuous quantity. The new approximation space for the Darcy velocity is still defined by (34)–(35), but since \tilde{V}^{ms} is not required to be a subspace of $H_0^{1,\text{div}}$, the matrices \mathbf{B} and \mathbf{C} in the resulting linear system

$$\begin{bmatrix} \mathbf{B} & -\mathbf{C}^T & \mathbf{\Pi}^T \\ \mathbf{C} & \mathbf{0} & \mathbf{0} \\ \mathbf{\Pi} & \mathbf{0} & \mathbf{0} \end{bmatrix} \begin{bmatrix} \mathbf{v} \\ \mathbf{p} \\ \pi \end{bmatrix} = \begin{bmatrix} \mathbf{0} \\ \mathbf{f} \\ \mathbf{0} \end{bmatrix}, \quad (40)$$

will have a block-diagonal structure. The new matrix $\mathbf{\Pi}$ is defined by

$$\mathbf{\Pi} = \left[\int_{\partial K_m} \psi_{kl} \cdot n \, ds \right],$$

and block reduction of the system (40) yields

$$\mathbf{M}\pi = \tilde{\mathbf{f}}, \quad (41)$$

where

$$\begin{aligned}\mathbf{M} &= \mathbf{S}(\mathbf{C}^T \mathbf{D}^{-1} \mathbf{C} - \mathbf{B})\mathbf{S}^T, \\ \tilde{\mathbf{f}} &= \mathbf{S}\mathbf{C}^T \mathbf{D}^{-1} \mathbf{f}, \\ \mathbf{S} &= \mathbf{\Pi} \mathbf{B}^{-1}, \quad \text{and} \quad \mathbf{D} = \mathbf{C} \mathbf{B}^{-1} \mathbf{C}^T.\end{aligned}$$

Here \mathbf{B} is block-diagonal, so \mathbf{B}^{-1} can be computed block-by-block. Moreover, \mathbf{D} is diagonal⁶ and can therefore easily be inverted, and \mathbf{M} is symmetric positive definite (see, e.g., [14]). Hence, \mathbf{M} can be computed explicitly, and the linear system can be solved using one of the many efficient methods specialised for symmetric positive-definite linear systems.

4.3 Examples

Both the MsMFEM and the MsFVM solve a coarse-scale equation globally while trying to resolve fine-scale variations by using special multiscale basis functions. The basis functions are local solutions of the elliptic equation, and the therefore overall accuracy only depends weakly on the coarse grid size. We shall demonstrate this with an example.

Example 3. Consider a horizontal, two-dimensional reservoir with fluvial heterogeneity. An injection well is located in the centre and a producer is located in each of the four corners. The permeability on the 220×60 fine grid is taken from the bottom layer of Model 2 in the Tenth SPE Comparative Solution Project [18]. The dataset has many narrow and intertwined high-permeability channels that greatly impact the flow pattern. We solve the pressure equation using both the MsFVM and the MsMFEM with various coarse grid dimensions, and employ an upstream finite-volume method for the saturation equation on the underlying fine grid. A reference solution is computed using a two-point flux approximation scheme on a grid that is refined four times in each direction. Figure 14 shows the resulting saturation fields at dimensionless time $t = 0.3\text{PVI}$. The solutions are quite difficult to distinguish visually. We therefore measure errors in the saturation fields by

$$\delta(S) = \frac{\epsilon(S)}{\epsilon(S_{\text{ref}})}, \quad \epsilon(S) = \frac{\|S - \mathcal{I}(S_{\text{ref}}^{4\times})\|_{L^1}}{\|\mathcal{I}(S_{\text{ref}}^{4\times})\|_{L^1}}.$$

Here the operator $\mathcal{I}(\cdot)$ represents mapping from the refined to the original grid. The results displayed in Table 3 verify that the accuracy is indeed relatively insensitive to coarse grid size, although there is a slight degradation of solution quality as the grid is coarsened.

⁶This may be seen by observing that $\mathbf{X}\mathbf{Y}^T$ will be diagonal for any two matrices \mathbf{X} and \mathbf{Y} having the sparsity structure of \mathbf{C} , . The result then follows from the fact that multiplication by \mathbf{B}^{-1} does not change the sparsity structure of \mathbf{C} .

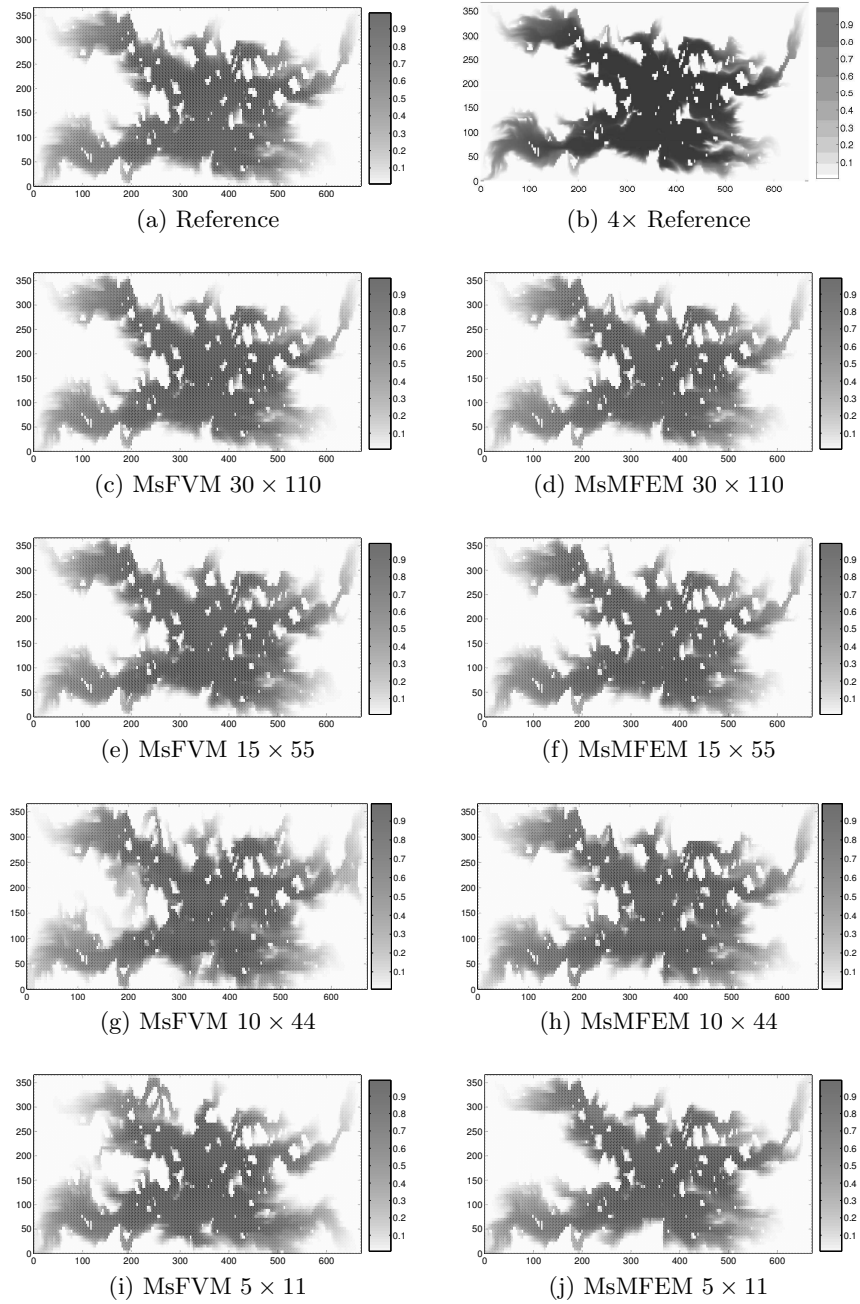


Fig. 14. MsMFEM and MsFVM solutions for various coarse grids on the bottom layer of the SPE 10 test case. The two figures at the top show the reference solution on the original fine and the $4\times$ refined grids.

Table 3. $\delta(S)$ for various coarse grids.

	30×110	15×55	10×44	5×11
MsMFEM	1.0916	1.2957	1.6415	1.9177
MsFVM	1.0287	1.6176	2.4224	3.0583

Boundary Conditions

The accuracy of the multiscale methods is largely determined by the boundary conditions for the local problems. In fact, it can be shown that by using properly scaled restrictions of a global fine-scale solution as boundary conditions for the basis functions, it is possible to replicate the fine-scale solution by solving the pressure equation on the coarse grid with the MsMFEM [2, 1] or the MsFVM [28]. For simulations with many time steps, computation of an initial fine-scale solution may be justified. Indeed, if the pressure field is reasonably stable throughout the simulation, it is possible to exploit the initial solution to obtain proper boundary conditions at subsequent time-steps [2, 42]. We usually refer to such boundary conditions as *global* boundary conditions, and they have been shown to give very accurate results, e.g., for the SPE10 model [4, 42].

Example 4. We now demonstrate in Table 4 that the multiscale methods with global boundary conditions are capable of replicating a fine-scale solution. Since we here use the TPFA to compute the fine-grid solution, we have $\delta(S) \equiv 1$ only for the MsFVM. If the lowest-order Raviart–Thomas mixed finite-element method had been used to compute the fine-grid solution and basis functions for the MsMFEM, we would get $\delta(S) \equiv 1$ for the MsMFEM instead.

Table 4. $\delta(S)$ for various coarse grids, boundary conditions determined using a fine-scale solution.

	30×110	15×55	10×44	5×11
MsMFEM	0.9936	0.9978	0.9938	0.9915
MsFVM	1.0000	1.0000	1.0000	1.0000

Multiscale Methods used as Upscaling Methods

Although the multiscale methods provide velocities on whatever we decide to use as the fine-grid, they can also be utilised as upscaling methods. Indeed, the velocities can easily be integrated along grid lines of an arbitrary coarse grid. The multiscale methods can therefore be viewed as seamless upscaling methods for doing fluid-transport simulations on user-defined coarse grids.

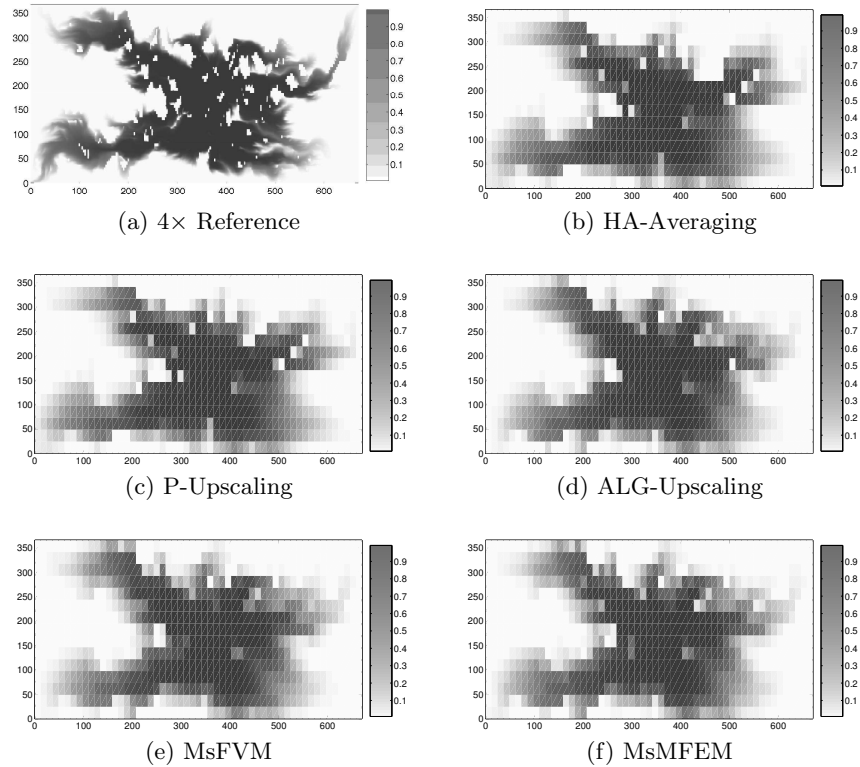


Fig. 15. Saturation profiles on a 15×55 grid obtained by the various upscaling and multiscale methods.

Next we show that the multiscale methods give comparable accuracy to standard upscaling techniques on fixed coarse grids if the saturation equation is solved on the coarse grid.

Example 5. To validate the multiscale methods as upscaling methods, we again use the data from the bottom layer of the SPE10 model and solve the saturation equation on the same coarse grid that is used in the multiscale methods. For comparison, we also compute solutions obtained by using some of the upscaling methods described in Section 3. Figure 15 shows saturations for the 15×55 coarse grid. Table 5 shows the corresponding error measure $\delta_c(S)$, where the subscript c indicates that errors are computed on the coarse (upscaled) grid. We see that the multiscale methods perform slightly better than the upscaling methods. A more substantial improvement of saturation accuracy cannot be expected, since the listed upscaling methods for the pressure equation are quite robust, and generally perform well.

We now reverse the situation from Example 5 and compute saturations on the fine grid also for the upscaling methods. This requires that we downscale

Table 5. $\delta_c(S)$ for various coarse grids.

	30×110	15×55	10×44	5×11
MsMFEM	0.9891	0.9708	0.9556	0.9229
MsFVM	0.9932	0.9675	1.0351	1.1439
ALG-Upscaling	0.9813	0.9826	1.0223	1.0732
P-Upscaling	1.0313	1.0216	1.1328	1.2371
HA-Averaging	1.0371	1.0654	1.2300	1.8985

Table 6. $\delta_c(S)$ for various coarse grids, saturations computed on the original grid.

	30×110	15×55	10×44	5×11
MsMFEM	1.0775	1.2849	1.8564	2.9680
MsFVM	1.0190	1.7140	2.8463	5.9989
ALG-Upscaling	1.1403	1.5520	2.3834	4.4021
P-Upscaling	1.6311	2.6653	4.1277	6.5347
HA-Upscaling	1.6477	2.9656	4.8572	15.3181

the velocity field by reconstructing a conservative fine-scale velocity field from a coarse-scale solution. For this we employ the procedure described in [32]. After having solved the saturation equation on the fine grid, we map the saturations back to the upscaled grid and compute the errors there. From the results in Table 6 we see that the multiscale methods give much better results than simple averaging and purely local upscaling. The local-global method, on the other hand, gives good accuracy, but computationally it is also much more expensive than the multiscale methods.

4.4 Concluding Remarks

As we have seen, the multiscale methods are capable of producing accurate solutions on both the fine and coarse scales, and may therefore be utilised as either robust upscaling methods or efficient approximate fine-scale solution methods. Having fine-scale velocities available at a relatively low computational cost gives great flexibility in the choice of solution method for the saturation equation. In light of the limited success of upscaling methods for the saturation equation, it seems beneficial to perform the transport on the finest grid affordable, and a fine-scale velocity field may aid in choosing such a grid as well as ensure the accuracy of the resulting solution.

Although we have shown only very simple examples, we would like to emphasise the advantage multiscale methods have in their flexibility with respect to grids. In particular this is true for the MsMFEM, since it avoids the dual-grid concept. For the MsMFEM each coarse-grid block is simply a connected union of fine-grid blocks, which makes the method virtually grid-independent given a fine-scale solver. Therefore it is straightforward to perform multiscale




Fig. 16. The fine-grid is a corner-point grid, which is the industry standard for reservoir modeling and simulation. The coarse-grid blocks are given by different colours.




Fig. 17. The coarse grid is tuned to model a well and the near-well region. Away from the well regions the coarse-grid blocks are aligned with the physical layers.

simulations on models with rather complicated fine-grids, such as the one displayed in Figure 16.

Furthermore, since the only real constraint on the geometry of coarse-grid blocks seems to be that it must allow computation of the local solutions, we may tune the coarse-grid to model particular reservoir features or wells without worrying about grid-block shapes. Figure 17 shows an example where the coarse-grid blocks are aligned with the physical layers. A vertical well is modelled by a narrow block, and the near-well region is modelled by another block with a hole. Even such exotic grid-blocks are unproblematic for the MsMFEM, and the multiscale simulations yield accurate results.

Finally, it is worth mentioning that the multiscale methods are flexible not only with regard to grids, but also with regard to fine-grid solution methods.

The fine-scale numerical method may be chosen independently for each of the local problems, and thus a higher-order method or local grid-refinement, may be employed for higher accuracy in regions of interest. Such approaches will not significantly impact the efficiency of the methods, since multiscale efficiency mainly comes from adaptively updating only a small portion of the local solutions. The multiscale methods may therefore have the potential to combine accuracy and efficiency in an ideal way.

References

1. J. E. Aarnes, S. Krogstad, and K.-A. Lie. A hierarchical multiscale method for two-phase flow based upon mixed finite elements and nonuniform coarse grids. *Multiscale Model. Simul.*, to appear, 2006.
2. J.E. Aarnes. On the use of a mixed multiscale finite element method for greater flexibility and increased speed or improved accuracy in reservoir simulation. *Multiscale Model. Simul.*, 2(3):421–439, 2004.
3. J.E. Aarnes, T. Gimse, and K.-A. Lie. An introduction to the numerics of flow in porous media using Matlab. In G. Hasle, K.-A. Lie, and E. Quak, editors, *Geometrical Modelling, Numerical Simulation and Optimization: Industrial Mathematics at SINTEF*. Springer Verlag.
4. J.E. Aarnes, V. Kippe, and K.-A. Lie. Mixed multiscale finite elements and streamline methods for reservoir simulation of large geomodels. *Adv. Water Resour.*, 28(3):257–271, 2005.
5. J.E. Aarnes and K.-A. Lie. Toward reservoir simulation on geological grid models. In *Proceedings of the 9th European Conference on the Mathematics of Oil Recovery*. EAGE, Cannes, France, 2004.
6. O. J. Arntzen. Higher-order terms in the capillary limit approximation of viscous-capillary flow. *Transp. Porous Media*, 57:17–34, 2004.
7. J.W. Barker and P. Dupouy. An analysis of dynamic pseudo relative permeability methods. In *Proceedings of the 5th ECMOR conference, Leoben, Austria*. 1996.
8. J.W. Barker and F.J. Fayers. Transport coefficients for compositional simulations with coarse grids in heterogeneous media. *SPE Adv. Tech Series*, 2(2):103–112, 1994.
9. J.W. Barker and S. Thibeau. A critical review of the use of pseudorelative permeabilities for upscaling. *SPE Reservoir Eng.*, 12(2):138–143, 1997.
10. S.H. Begg, R.R. Carter, and P. Dranfield. Assigning effective values to simulator gridblock parameters for heterogeneous reservoirs. *SPE Reservoir Eng.*, 4(4):455–463, 1989.
11. A. Benoussan, J.-L. Lions, and G. Papanicolaou. *Asymptotic Analysis for Periodic Structures*. Elsevier Science Publishers, Amsterdam, 1978.
12. D. Braess. *Finite elements: Theory fast solvers and applications in solid mechanics*. Cambridge University Press, Cambridge, 1997.
13. S.C. Brenner and L.R. Scott. *The mathematical theory of finite element methods*. Springer–Verlag, New York, 1994.
14. F. Brezzi and M. Fortin. *Mixed and hybrid finite element methods*. Computational Mathematics. Springer–Verlag, New York, 1991.

15. Y. Chen and L.J. Durlofsky. Adaptive local-global upscaling for general flows in heterogeneous formations. *Transp. Porous Media*, submitted.
16. Y. Chen, L.J. Durlofsky, M. Gerritsen, and X.H. Wen. A coupled local-global upscaling approach for simulation flow in highly heterogeneous formations. *Adv. Water Resour.*, 26:1041–1060, 2003.
17. Z. Chen and T.Y. Hou. A mixed multiscale finite element method for elliptic problems with oscillating coefficients. *Math. Comp.*, 72:541–576, 2003.
18. M. A. Christie and M. J. Blunt. Tenth SPE comparative solution project: A comparison of upscaling techniques. *SPE Reservoir Eval. Eng.*, 4(4):308–317, 2001. url: www.spe.org/csp.
19. M.A. Christie. Upscaling for reservoir simulation. *JPT J. Pet. Tech.*, 48:1004–1010, 1996.
20. L.J. Durlofsky. Numerical calculations of equivalent gridblock permeability tensors for heterogeneous porous media. *Water Resour. Res.*, 27(5):699–708, 1991.
21. L.J. Durlofsky. Use of higher moments for the description of upscaled, process independent relative permeabilities. *SPE J.*, 2(4):474–484, 1997.
22. L.J. Durlofsky. Coarse scale models of two-phase flow in heterogeneous reservoirs: Volume averaged equations and their relation to existing upscaling techniques. *Comput. Geosci.*, 2:73–92, 1998.
23. L.J. Durlofsky, R.C. Jones, and W.J. Milliken. A nonuniform coarsening approach for the scale-up of displacement processes in heterogeneous porous media. *Adv. Water Resour.*, 20:335–347, 1997.
24. Y. Efendiev. *The Multiscale Finite Element Method and its Applications*. PhD thesis, California Institute of Technology, 1999.
25. Y. Efendiev and L.J. Durlofsky. Numerical modeling of subgrid heterogeneity in two phase flow simulations. *Water Resour. Res.*, 38:1128, 2002.
26. Y. Efendiev and L.J. Durlofsky. A generalized convection-diffusion model for subgrid transport in porous media. *Multiscale Model. Simul.*, 1:504–526, 2003.
27. Y. Efendiev, L.J. Durlofsky, and S.H. Lee. Modeling of subgrid effects in coarse scale simulations of transport in heterogeneous porous media. *Water Resour. Res.*, 36(8):2031–2041, 2000.
28. Y. Efendiev, V. Ginting, T. Hou, and R. Ewing. Accurate multiscale finite element methods for two-phase flow simulations. *J. Comp. Phys.*, submitted.
29. Y. Efendiev and B. Popov. On homogenization of nonlinear hyperbolic equations. *Commun. Pure and Appl. Anal.*, 4:297–311, 2005.
30. S. Ekrann and J. O. Aasen. Steady-state upscaling. *Transp. Porous Media*, 41(3):245–262, 2000.
31. M.H. Garcia, A.G. Journel, and K. Aziz. Automatic grid generation and adjustment method for modeling reservoir heterogeneities. Technical report, Stanford Center for Reservoir Forecasting, 1990.
32. Y. Gautier, M. J. Blunt, and M. A. Christie. Nested gridding and streamline-based simulation for fast reservoir performance prediction. *Comput. Geosci.*, 3:295–320, 1999.
33. J.J. Gómez-Hernández. *A Stochastic Approach to the Simulation of Block Conductivity Conditioned Upon Data Measured at a Smaller Scale*. PhD thesis, Stanford University, CA, 1991.
34. R.E. Guzman, D. Giordano, F.J. Fayers, A. Godi, and K. Aziz. Evaluation of dynamic pseudo functions for reservoir simulation. In *SPE Annual Technical Conference and Exhibition, 6-9 October, Denver, Colorado*, SPE 35157, 1996.

35. T.A. Hewett and R.A. Behrens. Scaling laws in reservoir simulation and their use in a hybrid finite difference/streamtube approach to simulating the effects of permeability heterogeneity. In L.W. Lake, H.B. Carroll, and T.C. Wesson, editors, *Reservoir Characterization II*. Academic Press, San Diego, Ca, 1991.
36. T.A. Hewett, K. Suzuki, and M.A. Christie. Analytical calculation of coarse-grid corrections for use in pseudofunctions. *SPE J.*, 3(3):293–304, 1998.
37. L. Holden and B.F. Nielsen. Global upscaling of permeability in heterogeneous reservoirs; the output least squares (ols) method. *Trans. Porous Media*, 40(2):115–143, 2000.
38. U. Hornung. *Homogenization and porous media*. Springer–Verlag, New York, 1997.
39. T.Y. Hou and X-H. Wu. A multiscale finite element method for elliptic problems in composite materials and porous media. *J. Comput. Phys.*, 134:169–189, 1997.
40. H. Jacks, O.J.E. Smith, and C.C. Mattax. The modeling of a three-dimensional reservoir with a two-dimensional reservoir simulator - the use of dynamic pseudo functions. *SPE J.*, pages 175–185, 1973.
41. P. Jenny, S. H. Lee, and H. A. Tchelepi. Multi-scale finite-volume method for elliptic problems in subsurface flow simulation. *J. Comput. Phys.*, 187:47–67, 2003.
42. P. Jenny, S. H. Lee, and H. A. Tchelepi. Adaptive multiscale finite-volume method for multiphase flow and transport in porous media. *Multiscale Model. Simul.*, 3(1):50–64, 2004.
43. V.V. Jikov, S.M. Kozlov, and O.A. Oleinik. *Homogenization of differential operators and integral functionals*. Springer–Verlag, New York, 1994.
44. A.G. Journel, C.V. Deutsch, and A.J. Desbarats. Power averaging for block effective permeability. In *SPE California Regional Meeting, 2-4 April, Oakland, California*, SPE 15128, 1986.
45. M.J. King, D.G. MacDonald, S.P. Todd, and H. Leung. Application of novel upscaling approaches to the Magnus and Andrew reservoirs. In *European Petroleum Conference, 20-22 October, The Hague, Netherlands*, SPE 50643, 1998.
46. P.R. King. The use of renormalization for calculating effective permeability. *Transp. Porous Media*, 4:37–58, 1989.
47. P.R. King, A.H. Muggeridge, and W.G. Price. Renormalization calculations of immiscible flow. *Transp. Porous Media*, 12:237–260, 1993.
48. A. T. Kumar and G. R. Jerauld. Impacts of scale-up on fluid flow from plug to gridblock scale in reservoir rock. In *SPE/DOE Improved Oil Recovery Symposium, 21-24 April, Tulsa, Oklahoma*, SPE 35452, 1996.
49. J.R. Kyte and D.W. Berry. New pseudofunctions to control numerical dispersion. *SPE J.*, pages 269–276, August 1975.
50. P. Langlo and Magne S. Espedal. Macrodispersion for two-phase, immiscible flow in porous media. *Adv. Water Resour.*, 17:291–316, 1994.
51. B. F. Nielsen and A. Tveito. An upscaling method for one-phase flow in heterogeneous reservoirs; a Weighted Output Least Squares (WOLS) approach. *Comput. Geosci.*, 2:92–123, 1998.
52. P.-E. Øren, S. Bakke, and O.J. Arntzen. Extending predictive capabilities to network models. *SPE J.*, 3(4):324–336, 1998.
53. G.E. Pickup and K.S. Sorbie. Development and application of a new two-phase scaleup method based on tensor permeabilities. *SPE J.*, pages 369–382, 1996.

54. P.A. Raviart and J.M. Thomas. A mixed finite element method for second order elliptic equations. In I. Galligani and E. Magenes, editors, *Mathematical Aspects of Finite Element Methods*, pages 292–315. Springer–Verlag, Berlin – Heidelberg – New York, 1977.
55. P. Renard and G. de Marsily. Calculating equivalent permeability. *Adv. Water Resour.*, 20:253–278, 1997.
56. E.H. Smith. The influence of small scale heterogeneity on average relative permeability. In L.W. Lake, H.B. Carroll, and T.C. Wesson, editors, *Reservoir Characterization II*. Academic Press, 1991.
57. K. D. Stephen, G. E. Pickup, and K. S. Sorbie. The local analysis of changing force balances in immiscible incompressible two-phase flow. *Transp. Porous Media*, 45:63–88, 2001.
58. H.L. Stone. Rigorous black-oil pseudo functions. In *SPE Symposium on Reservoir Simulation, 17-20 February, Anaheim, California*, SPE 21207, 1991.
59. G.W. Thomas. An extension of pseudofunction concepts. In *SPE Reservoir Simulation Symposium, 15-18 November, San Francisco, California*, SPE 12274, 1983.
60. X.-H. Wen and J.J. Gómez-Hernández. Upscaling hydraulic conductivities in heterogeneous media: An overview. *J. Hydrol.*, 183:ix–xxxii, 1996.
61. O. Wiener. *Abhandlungen der Mathematisch*. PhD thesis, Physischen Klasse der Königlichen Sächsischen Gesellschaft der Wissenschaften, 1912.
62. X.H. Wu, Y. Efendiev, and T.Y. Hou. Analysis of upscaling absolute permeability. *Discrete Contin. Dyn. Syst. Ser. B*, 2(2):185–204, 2002.
63. P. Zhang, G. E. Pickup, and M. A. Christie. A new upscaling approach for highly heterogenous reservoirs. In *SPE Reservoir Simulation Symposium, 31 January-2 February, The Woodlands, Texas*, SPE 93339, 2005.