UNIVERSITY
OF OSLO

Master's thesis

An Extended Fully-Implicit Hybrid Model for Geological CO_2 Storage

With applications to FluidFlower

Vetle Nevland

Computational Geosciences
60 ECTS study points

Department of Geosciences
Faculty of Mathematics and Natural Sciences

Spring 2023

Vetle Nevland

An Extended Fully-Implicit Hybrid Model for Geological CO_2 Storage

With applications to FluidFlower

Supervisors:
Odd Andersen
Knut-Andreas Lie
Helge Hellevang

Contents

Acknowledgements

There are several people that deserve praise for my opportunity to write this thesis. First and foremost, I would like to thank my external supervisors at SINTEF Digital, Knut-Andreas Lie and Odd Andersen for excellent supervision and patiency, especially for assistance in understanding the inner workings of MRST that helped me develop the product of this thesis. Despite late proposal, I am very grateful for the opportunity to contribute to the highly relevant 11th SPE CSP. This certainly made the work of my thesis more valuable. Jan Martin Nordbotten and his team at the University of Bergen deserve a big thanks for their realization of the FluidFlower concept. I would also like to thank my internal supervisor at UiO, Helge Hellevang, for your flexibility and formal advices. Last, but not least, thanks to my family and friends for the continuous support they have shown through the difficult, the good and the rewarding times of my Master's degree.

Abstract

Geological CO2 storage is an important initiative to mitigate climate change. Numerical models are valuable for estimating storage capacity and for risk assessment. Due to large scales involved, simplified models are desired. Vertical equilibrium (VE) models have shown orders-of-magnitude reduced runtime, but are challenged by formation heterogeneities. Herein, we extend a fully-implicit hybrid model developed in MRST with residual saturation and capillary exclusion. We relax the VE assumption (RVE) to model diffuse leakage of CO2 through sealing barriers. For comparisons, we extend a conventional trap analysis to be compatible with hybrid models.

Our extended hybrid model agrees well with full-dimensional models, with runtimes being shorter. From the RVE assumption, it manages to trace the history of individual plumes spread among sealing barriers, but misses lateral spreading of CO2 beyond these. Stochastic simulations indicate that differences in depth of the CO2 plume between the models decrease for an increased number of sealing barriers, while differences in net CO2 trapping show an increasing trend for more sealing barriers.

As a second purpose of the thesis, we develop a novel, semi-structured grid in MRST that conforms precisely to the facies of the FluidFlower rig, as part of the 11th SPE Comparative Solution Project. The grid is also compatible with hybrid modeling. Results from a full-dimensional compositional model correspond well with physical measurements, particularly adjacent to anticlines and faults, but also show grid-orientation effects that are subject to improvement.

The hybrid model shows promising results on the 11th SPE CSP, despite convergence issues. Whether the FluidFlower rig is a good candidate for hybrid modeling remains an open question. The purpose herein is to demonstrate the versatility of our hybrid model for simulations on heterogeneous formations. We believe it is a promising candidate for history-matching and greater insight into numerical challenges for CO2 storage simulation in general, and for the 11th SPE CSP in particular.

Chapter 1

Introduction

In the past decade, consequences of accumulated greenhouse gas emissions have materialized globally, and the severity of the climate crisis is indisputable. A promising initiative for mitigating the ongoing climate change is carbon capture and storage, a technology that has been active and developed over several decades. To safely store CO2 in the subsurface, it is crucial to make reliable estimates of where the CO2 ends up in the long term. Numerical simulations are invaluable for this purpose, and for ultimately performing safe and efficient injection operations.

Injected CO2 should preferentially be trapped inside the reservoir indefinitely. The leakage potential is affected by many factors including reservoir geometries, petrophysical properties, and operational management such as injection rates and bottom-hole pressure. Finding an optimal injection site that maximizes CO2 storage and minimizes leakage (risk) is not trivial, and numerical modeling is inevitable to answer important questions to this end, particularly for subsurface reservoirs lacking detailed petrophysical data. Although seismicity has provided petrophysical data of the Sleipner field, where injection has been ongoing since 1996 [2], the level of detail is limited. Full-scale models have been built based on this data, but how predictive they are remains elusive. This stresses the need for long-term numerical simulations.

One challenge is that traditional reservoir simulators are built for problem settings whose spatial and temporal scales are significantly smaller than the scales involved in CO2 storage. Injected CO2 tends to accumulate below impermeable layers and migrate horizontally over long distances. Thus, CO2 storage requires models that work with lateral dimensions of several hundred kilometers. Moreover, injected CO2 should be safely stored not only for decades but for millennia. Although the large spatial and temporal scales become manageable by choosing low enough spatial and temporal resolutions, it is not a viable option because CO2 storage involves physical processes at the same scale as traditional reservoir simulation [49]. In other words, high spatial and temporal resolutions are necessary for reliable simulations. Compared to conventional 3D reservoir simulators, this poses an almost intractable computational challenge [52, 48]. On top of this, the governing equations are highly nonlinear, requiring expensive iterative numerical schemes [52, 42].

An ongoing challenge with CO2 storage simulation is to simulate all *significant* physical processes with sufficient resolution [17]. To this end, several assumptions and simplifications have been applied to conventional simulators, achieving great accuracy [49]. A successful approach is the assumption of vertical equilibrium (VE). It is based on the fact that CO2 and brine tend to segregate rapidly in the vertical direction. If the time of buoyant segregation is shorter than the time steps in the simulation, it is reasonable to assume that segregation of CO2 and brine happens instantaneously [8]. In this case, the vertical distribution of fluids can be computed analytically and the spatial dimensions are reduced from three to two, facilitating faster computations.

Although VE models achieve great accuracy, often comparable to full 3D simulators, their application is challenged by heterogeneous formations. In regions of complex flow behavior, e.g., in low-permeable layers, the assumption of VE is usually not justified [18]. To circumvent this issue, hybrid models have been developed that increase accuracy significantly, with (in many cases) an acceptable sacrifice in computational requirements. The idea is to combine the respective advantages of VE and full-dimensional models by applying VE models in regions where the assumption of immediate vertical segregation is justified, and 3D models in regions with significant vertical flow [44]. Compared to pure VE, hybrid models have been shown to provide results closer to conventional 3D simulators [44, 43, 10]. The convenience of hybrid models is substantiated by the fact that there are large uncertainties in the petrophysical parameters of deep saline aquifers. Multiple realizations of the reservoir are thus a frequent demand, but this is costly for full-dimensional models. Fortunately, the significant reduction in runtime makes multiple realizations for hybrid models practical [64].

Hybrid models serve as a great tradeoff between VE models and fine-scale 3D models, but some modeling aspects remain elusive, for example, diffuse leakage. Møyner, Nilsen and Andersen [44, 43] touched upon this effect, but flow behavior internally in low-permeable barriers was not resolved. There is also a desire to find correlations between the accuracy of hybrid models and the distribution of low-permeable barriers in a multilayered system. The applicability of hybrid models is ultimately dependent on the relevant time scale for the problem, the spatial and temporal balance of forces, as well as the degree of reservoir heterogeneity [26].

Multiple variants of hybrid models exist in the scientific literature, the differences being primarily the numerical schemes, physical effects incorporated, and representation of the reduced dimensions. Møyner et al. [44] studied long-term CO2 storage in multilayered systems by coupling a full-dimensional 3D model with stacked VE models in a fully-implicit hybrid framework. The model was later extended with compositional effects [43]. Other studies [26, 69, 27] have developed hybrid models by combining a full-dimensional model with a *relaxed* VE model in which vertical flow is resolved as isolated one-dimensional problems. Current studies based on a fully-implicit solution scheme [44, 43] assume no residual saturation and negligible capillary entry pressure between layers. It would be valuable to further develop the fully-implicit hybrid model

by implementing these physical effects.

Another area of improvement is appropriate modeling of flow through nearly impermeable barriers, henceforth referred to as *semi-permeable*, *low-permeable* or *sealing barriers*. The three terms will be used interchangeably throughout this thesis. The hybrid models in [44, 43] represented semi-permeable barriers as lower-dimensional faces in the grid, adjusting their transmissibility by multipliers to account for diffuse leakage of brine. However, internal flow is lost when using a lower-dimensional representation. An alternative is to represent the barriers as full-dimensional cells. This has the benefit of resolving any flow internally in the semi-permeable barriers by explicitly modeling pressure gradients that drive the flow. Full-dimensional representations are computationally more demanding, but the semi-permeable barriers often constitute a minor part of the reservoir. Thus, hybrid models can still be significantly less complex than full-dimensional models. In return, results of hybrid models are in excellent agreement with full-dimensional models for semi-complex heterogeneous formations while running up to orders of magnitude faster [44, 10].

The open-source software MRST [61] provides a module `mrst-autodiff/hybrid-ve`, which is specific for the application of hybrid models, based on a fully-implicit numerical solution scheme. An essential part of this thesis is devoted to extending hybrid models with currently missing physical effects such as residual saturation and capillary exclusion, which may have a great impact on the migrating plume and long-term trapping [23]. Also, the possibility of improved modeling of diffuse leakage will be addressed by considering multilayered reservoir systems. Our novel hybrid framework is presented through studies on synthetic, two-dimensional, heterogeneous formations. Qualitatively and quantitatively, we compare it with a full-dimensional model by extending a conventional trap analysis [46] to be compatible with a hybrid framework.

Thereafter, we proceed with a study on data from a real reservoir. The intention was to model a section of the Utsira formation, an excellent candidate for modeling diffuse leakage due to its multilayered characteristics [16]. But over the time course of the thesis, a more relevant modeling problem emerged, namely the eleventh SPE Comparative Solution Project [54]. The benchmark study involves a multilayered reservoir generated in the FluidFlower laboratory rig at the University of Bergen [21, 56]. Quantifying the accuracy of numerical models for simulating flow in real reservoirs is challenging due to inadequate geometrical and petrophysical data. To this end, the FluidFlower benchmark study is a great initiative to provide robust comparisons of model predictions with empirical observations [30]. Availability of local data and plume shape facilitates history-matching; an invaluable technique that can be used to tune model parameters and improve the accuracy of our numerical model [56, 22, 20, 48]. The benchmark will also provide valuable insight into modern challenges involved when simulating CO2 storage for realistic geological settings [54].

We believe the FluidFlower benchmark case is a relevant candidate for assessing our novel hybrid model on real data mimicking conditions of subsurface reservoirs. The

system is subject to atmospheric conditions, so CO2 will appear in a gaseous state with significantly lower density than the supercritical state applicable at subsurface conditions. As a result, the effect of buoyancy is enhanced and one expects vertical segregation of fluids to occur more rapidly than in the subsurface. In addition, the simulation times are too short for a capillary fringe of significant thickness to develop [51]. Thus, our assumption of a sharp interface with capillary exclusion is reasonable for the FluidFlower system. Despite high-permeable sand layers, the very small time steps required by numerical solvers challenges the VE assumption [56]. This raises the question of the applicability of our hybrid model on the FluidFlower. Nonetheless, the benchmark serves as a great demonstration of its versatility.

The trajectory of the thesis has deviated slightly from the intended, exclusive focus of developing a fully-implicit hybrid model to a parallel path where we also aim to generate a detailed, numerical grid in MRST of the geology of the FluidFlower rig. To this end, significant effort has been put into making the grid conform to the facies, including faults and pinch-outs, as well as adopting a semi-structured topology that facilitates hybrid modeling. Due to these unanticipated changes to the thesis, extensions of the hybrid model have been limited to implementation of capillary exclusion and residual saturation. Empirical studies show that dissolution is an important physical effect for the FluidFlower [56, 22]. Thus, dissolution should be implemented before comparing our hybrid simulations on the FluidFlower with other simulators that account for this effect.

The FluidFlower is projected to gain large scientific attention in the years to come and is an important initiative for improving numerical models for long-term CO2 storage. The work of this thesis contributes to the research on the FluidFlower by generating a flexible grid in MRST that supports conventional, compositional simulators as well as hybrid modeling.

Chapter 2

Theory

CO2 storage in the subsurface is a comprehensive study of multiphase flow involving numerous processes operating on a multitude of scales [52]. To generate a feasible model to simulate CO2 injection and migration, it is essential to explain and comprehend the elements that constitute the CO2 storage system. This section describes the significant physical processes involved in light of the relevant dimensional scales.

2.1 Physical description

Similar to traditional reservoir simulation, modeling of CO2 storage involves the study of flow in porous media. In this regard, we are considering the joint flow of CO2 and brine, labeled multiphase flow. The purpose of this section is to provide a concise description of the physics behind multiphase flow, in particular the properties of fluids and rocks, and the governing equations.

2.1.1 Force balance

Fluids in the subsurface are subject to different forces, whose balance governs the distribution and movement of the fluids. For multiphase flow, particularly a two-phase CO2–brine system, there are primarily three dominating forces: buoyant, capillary and viscous forces [67]. Their balance depends on multiple factors, including spatial and temporal scales, formation geometry, flow rates, and petrophysical and fluid properties. For high flow rates, the effect of viscous forces is dominating and the system is considered to be in the Darcy flow regime, attributable to significant pressure gradients [29]. In contrast, for low flow rates, viscous forces weaken and the effect of density is emphasized. The system approaches the capillary flow regime, attributable to distinct migration paths and capillary forces at the fluid–fluid interface. By addressing the force balance, we can distinguish regions of different flow regimes, a characterization that is important when it comes to numerical modeling, because certain regimes justify certain modeling simplifications [52, 29].

Viscous forces tend to dominate early in the injection stage when local pressure gradients are large, while capillary forces are diffusive and typically become significant

in the post-injection stage. For CO2 storage, gravity is the long-term dominating force dictating the spatial distribution of injected CO2 [67]. Capillary and buoyant forces have been shown to dominate viscous forces for CO2 migration in multiple reservoir storages, for example the Utsira formation as part of the Sleipner field on the Norwegian Continental Shelf [16]. This justifies the applicability of simplified models, which will be described in details in Chapters 5 and 6.

2.1.2 Fluid properties

To address how fluids flow in a reservoir system, it is essential to consider the unique properties of fluids. **Density**, ρ, is an important property that is responsible for buoyant forces that dominate on the long term, and it dictates the vertical distribution of fluids at hydrostatic equilibrium. Hydrostatic instabilities initiate buoyant forces that displace lighter fluids over heavier fluids.

Another important property is **dynamic viscosity**, μ, a measure of a fluid's resistance to shear forces. It contributes to the force balance by means of viscous forces. Density and dynamic viscosity are generally not fixed quantities, but are related to thermodynamic conditions through an equation of state [4]. Depending on the system under consideration, such equations can become quite complex, for instance if including phase mixing. For the models developed in this work, we assume phases remain distinct and do not account for mixing of components between phases. As a result, the equation of state simplifies significantly.

Fluids are also characterized by **compressibility**, describing the ability of an entity of fluid to change its volume subject to external forces. Compressibility is explicitly related to density by

$$c_f = \frac{1}{\rho}\frac{d\rho}{dp}, \tag{2.1}$$

where c_f is the compressibility factor of the fluid. Herein, it is assumed that density is only a function of pressure, p, but in general it depends on multiple properties including temperature, pressure and composition [52].

2.1.3 Petrophysical properties

It is not only inherent properties of the fluids that dictate migration paths of CO2. The surrounding rock matrix greatly influences the potential for flow, which explains why well placement is a crucial factor for optimizing long-term CO2 storage. **Permeability** is a property describing the ability of a particular rock to transmit fluids through its pores. In the case of multiphase flow, each phase has an effective permeability describing its ability to flow by the presence of other phases. The **absolute permeability**, $\boldsymbol{K}$, is the ability to transmit a single phase in the pore space. If multiple phases are present, phase α will experience more resistance to flow than if it occupied the pore space alone. Thus, the effective permeability of a phase is always less than the absolute permeability, whose ratio is referred to as **relative permeability**, $k_{r\alpha}$. Whereas absolute permeability is an

intrinsic property of the rock, relative permeability depends on the saturation of phase α in the pore space, usually as a nonlinear dependence that exhibits hysteresis [47].

Another essential property of the rock is **porosity**, ϕ, expressing the fraction of rock volume consisting of pores for fluids to flow through. A lower porosity implies less volume to store and transmit fluids. To maximize injection capacity while keeping the pressure buildup at a minimum, it is vital to locate formations with high porosity. Generally, porosity is a pressure-dependent parameter, implying that rock is a compressible material with compressibility c_r defined as

$$c_r = \frac{1}{\phi}\frac{d\phi}{dp}. \tag{2.2}$$

2.1.4 Capillary pressure

In a multiphase system, a fluid is subject to adhesive forces with the surrounding rock in addition to surface tension at the interface between the fluids present. This force balance is responsible for a pressure difference between the fluids known as capillary pressure. In a two-phase system, the capillary pressure is the pressure difference between a non-wetting phase n and a wetting phase w in hydrostatic equilibrium, expressed as

$$p_c = p_n - p_w, \tag{2.3}$$

where *wetting* expresses how strongly the fluid phase adheres to the solid rock. In a two-phase system of CO2 and brine, brine experiences the strongest attraction and thus represents the wetting phase while CO2 is the non-wetting phase. Depending on the dimensional scales of the system, capillary forces may have a significant impact on the distribution of phases, even in the context of geological CO2 storage [52]. Capillary pressure is often modeled through an inverse relationship of the saturation of the wetting phase, expressed as

$$p_c(s_w) = f^{-1}(s_w), \tag{2.4}$$

where f is assumed to be monotonic, in which case the expression can be inverted to obtain saturation s_w as a function of capillary pressure. As the water saturation approaches its residual value, the capillary pressure often diverges to infinity, as shown in Figure 2.1 for a typical $p_c - s_w$ relationship. In other words, an infinite amount of pressure is required for the non-wetting phase to force the residual part of the wetting phase out of the pore space. The invading non-wetting phase will never be able to displace this residual content. When the capillary pressure is in equilibrium with the gravitational force, the transition zone between the wetting phase and the non-wetting phase is a fully developed **capillary fringe** described by (2.4).

Capillary pressure functions vary according to the wettability of the fluid phases and the pore structure of the ambient rock. The entry pressure is defined as the minimum pressure required for the non-wetting phase to displace the wetting phase. The capillary pressure determines to what extent a non-wetting fluid is able to displace a wetting fluid

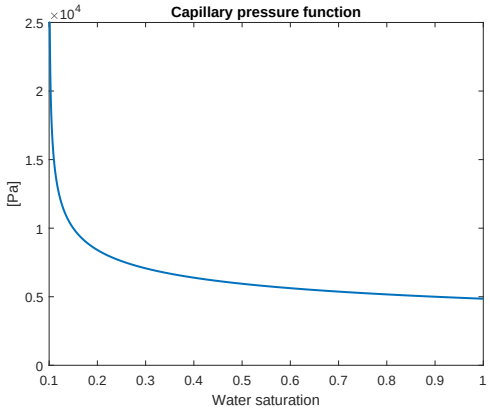

Figure 2.1: Typical relationship of capillary pressure and water saturation for primary drainage of a water–gas system with residual water saturation $s_{wr} = 0.1$. The abrupt jump from zero to 0.5 Pascal (Pa) at unit water saturation marks the entry pressure required for gas to enter pore space initially filled with water.

occupying some pore space, depending on whether the entry pressure in this region is exceeded or not. In the case of cylindrical pores, the capillary pressure is expressed analytically through the Young–Laplace equation [52],

$$p_c = \frac{2\gamma \cos(\theta)}{r}, \tag{2.5}$$

where γ is the surface tension between the two fluids and θ the contact angle between them. Tighter pore networks, i.e., pore throats with small radius r, give a larger capillary pressure. For long-term CO2 storage though, modeling is done on the reservoir scale which is too large to resolve pore size distribution on the microscale. Instead, it is customary to apply an averaged pore size so that the capillary pressure is essentially homogeneous for distinct formations [52].

2.1.5 Residual saturation

Residual saturation is the saturation for which a fluid in some pore space becomes immobile. That is, as saturation decreases and reaches the residual saturation, the fluid goes from being mobile to become residually trapped by capillary forces in isolated pockets surrounded by ambient fluid. Residual saturation influences the relative permeability curves. Instead of the relative permeability reaching zero at zero saturation, it reaches zero at the residual value. Moreover, if multiple phases are present in some pore space, the maximum relative permeability is not achieved for unit saturation but for one minus the residual saturation of the other phases.

Residual saturation is closely related to a concept known as **hysteresis**. As a non-wetting fluid enters new pore space, it displaces the current ambient wetting fluid in what is referred to as a displacement process. The relative permeability curve for primary displacement starts increasing at zero saturation. Once the wetting fluid imbibes the regions occupied by the non-wetting fluid, the relative permeability of the non-wetting

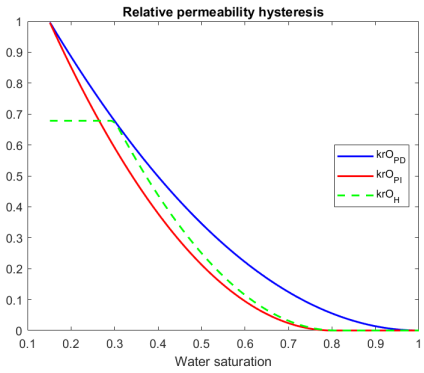

Figure 2.2: Relative permeability curves for a non-wetting phase (CO2) for a hysteretic system with non-zero residual saturation using the Killough model [34]. Blue curve is primary drainage curve, red curve is primary imbibition curve and green dashed curve shows imbibition curve for a reversal point $s_w > s_{wr} = 0.15$.

fluid decreases until it reaches zero at the residual saturation. In other words, the particular relative permeability curve followed depends on the history of the system [34]. For a two-phase system with wetting phase w and non-wetting phase n, the saturation domain in the relative permeability curves of the non-wetting phase for drainage and imbibition are given by, respectively,

$$\begin{aligned} D_{dr}(k_{rn}) &: \{s_n \mid s_n \in [0,\ 1 - s_{wr}]\}, \\ D_{im}(k_{rn}) &: \{s_n \mid s_n \in [s_{nr},\ 1 - s_{wr}]\}. \end{aligned} \tag{2.6}$$

That is, during injection, CO2 is mobile even for saturations below the residual value. Relative permeability curves for a CO2–brine system with residual saturation $s_{nr} = 0.2$ and $s_{wr} = 0.15$ including hysteretic effects are shown in Figure 2.2.

2.1.6 Governing equations

Darcy's law is an empirical relationship derived from the principle of conservation of momentum that describes flow through porous media. It states that the volumetric rate (flux) of a fluid through a pore network is proportional to the difference in hydraulic head $\delta h = \delta p/(\rho g) + \delta z$ at two prescribed locations. The constant of proportionality depends on the permeability. In the case of multiphase flow, relative permeability becomes a factor and the multiphase extension of Darcy's law becomes

$$\boldsymbol{u}_\alpha = -\frac{\boldsymbol{K} k_{r,\alpha}}{\mu_\alpha}(\nabla p_\alpha - \rho_\alpha \boldsymbol{g}), \tag{2.7}$$

where subscript α denotes a particular fluid phase, $\boldsymbol{g} = g\nabla z$, and the vertical component is assumed to be positive downwards, as appropriate for a subsurface system. Henceforth, vectors and tensors are represented in bold, and the distinction between them should be clear from context. For geological CO2 storage, (2.7) describes the fluxes of a two-phase system $\alpha = \{w, n\}$ where brine is the wetting phase w and CO2 is the non-wetting phase n.

The governing equations for multiphase flow incorporates conservation of mass with Darcy's law. The derivation is short and elegant. Consider an arbitrary closed volume Ω with boundaries denoted by $\partial\Omega$. To impose mass conservation, we require that any change of fluid mass M inside the volume is associated by a mass flux $\boldsymbol{f}$ into or out of the volume and/or external sources or sinks q. Accounting for contributions over the entire volume this becomes an integral expression,

$$\int_\Omega \frac{\partial M}{\partial t}\, dV = -\oint_{\partial\Omega} \boldsymbol{f}\cdot\boldsymbol{n}\, dA + \int_\Omega q\, dV. \tag{2.8}$$

The mass of the fluid per unit volume is given by $M = \rho\phi$ and the mass flux is $\boldsymbol{f} = \rho\boldsymbol{u}$, where $\boldsymbol{u}$ is the Darcy flux. By using the divergence theorem, the surface integral can be transformed to a volume integral,

$$\int_\Omega \Big(\frac{\partial \rho\phi}{\partial t} + \nabla\cdot(\rho\boldsymbol{u}) - q\Big) dV = 0. \tag{2.9}$$

Based on the assumption that the volume Ω is arbitrary, it must be the case that the integrand itself is zero, which provides the equation of mass conservation for a single fluid

$$\frac{\partial(\rho\phi)}{\partial t} + \nabla\cdot(\rho\boldsymbol{u}) = q. \tag{2.10}$$

Multiple extensions exist depending on the system under consideration. Density of CO2 varies as a function of pressure and temperature, as described by an equation of state, particularly near the critical point [52]. In the following, we assume that the fluids are compressible, and there is no mass transfer of components between phases. We refer to [38] for a derivation including mass components.

A multiphase extension yields one mass conservation equation per phase α, where the mass term is adjusted to account for pore space shared by multiple phases. The mass conserving equation for multiphase flow can be formulated as [38]

$$\frac{\partial}{\partial t}(\rho_\alpha\phi s_\alpha) + \nabla\cdot(\rho_\alpha\boldsymbol{u}_\alpha) = q_\alpha, \tag{2.11}$$

or, when incorporating (2.7),

$$\frac{\partial}{\partial t}(\rho_\alpha\phi s_\alpha) - \nabla\cdot\Big(\rho_\alpha\frac{\boldsymbol{K}k_{r\alpha}}{\mu_\alpha}(\nabla p_\alpha - \rho_\alpha g\nabla z)\Big) = q_\alpha, \tag{2.12}$$

where s_α and q_α are the saturation and source/sink term of phase α, $k_{r\alpha}$ and μ_α the relative permeability and viscosity of phase α, and $\boldsymbol{K}$ the intrinsic permeability of the rock.

The system (2.7) and (2.11) constitutes two equations but there are four unknowns, p_α and s_α for $\alpha = \{w, n\}$. To generate a closed system we need two additional closure equations. The first equation states that all available pore space is occupied by the present phases, i.e.,

$$s_w + s_n = 1. \tag{2.13}$$

The second one relates the difference in phase pressures to the capillary pressure curve, typically expressed in terms of water saturation

$$p_c(s_w) = p_n - p_w. \tag{2.14}$$

If one phase pressure and one saturation are chosen as primary variables, the other phase pressure is obtained by this capillary pressure-saturation relation.

Because the multiphase flow system consists of multiple equations with nonlinear dependencies, there are numerous ways to rewrite and solve them [52]. A convenient method for an incompressible system is to rewrite the governing equations as an elliptic pressure equation and a hyperbolic transport equation to decouple pressure and saturation. Thence, the pressure equation is not affected by the stability restrictions of the transport equation, allowing us to use as large time step as desired for the former [38]. Unfortunately, this splitting strategy is not as convenient for compressible systems because additional nonlinearities amplifies the coupling between the two equations.

Applying the product rule on the first term in (2.11) gives [52]

$$\frac{\partial}{\partial t}(\rho_\alpha \phi s_\alpha) = \rho_\alpha \phi \frac{\partial s_\alpha}{\partial t} + s_\alpha \rho_\alpha c_{t,\alpha} \frac{\partial p_\alpha}{\partial t}, \tag{2.15}$$

where $c_{t,\alpha} = c_\phi + \phi c_{f,\alpha}$ is the total compressibility. Plugging this into (2.11), dividing by ρ_α and summing over phases we end up with the following

$$\phi \sum_\alpha \frac{\partial s_\alpha}{\partial t} + \sum_\alpha s_\alpha c_{t,\alpha} \frac{\partial p_\alpha}{\partial t} - \sum_\alpha \frac{1}{\rho_\alpha} \nabla \cdot (\rho_\alpha \boldsymbol{u}_\alpha) = \sum_\alpha \frac{q_\alpha}{\rho_\alpha}. \tag{2.16}$$

The first term evaluates to zero because $\frac{\partial}{\partial t}(s_w + s_n) = 0$ and we get the pressure equation

$$\sum_\alpha s_\alpha c_{t,\alpha} \frac{\partial p_\alpha}{\partial t} - \frac{1}{\rho_\alpha} \nabla \cdot (\rho_\alpha \boldsymbol{u}_\alpha) = \sum_\alpha \frac{q_\alpha}{\rho_\alpha}, \tag{2.17}$$

to be solved for one of the phase pressures. By introducing the total Darcy velocity $\boldsymbol{u} = \boldsymbol{u_w} + \boldsymbol{u_n}$ and the fractional flow function $f_\alpha = \frac{\lambda_\alpha}{\lambda_w + \lambda_n}$, [38] shows how (2.11) can be developed to yield the transport equation

$$\frac{\partial}{\partial t}(\rho_w \phi s_w) + \nabla \cdot \Big[\rho_w f_w (\boldsymbol{u} + \lambda_n \boldsymbol{K}(\Delta \rho g \nabla z + \nabla p_c \nabla s_w))\Big] = q_w, \tag{2.18}$$

to be solved for the wetting phase saturation.

The elliptic pressure equation and hyperbolic transport equation possess different physical characteristics. By splitting the governing equations we acquire greater insight into the mathematical nature of the system by separating the parts that are weakly coupled [38]. This is leveraged in the IMPES solution scheme that is used in many works on VE models from the literature, e.g., [28, 27, 10]. In this thesis, we use a fully-implicit scheme with a black-oil formulation of the governing equations instead of operator splitting.

2.1.7 Trapping inventory

CO2 injected into a reservoir system is subject to different trapping mechanisms whose importance depends on fluid and petrophysical properties such as residual saturation, permeability and caprock geometry [31]. To address how much of the injected volume of CO2 has potential to migrate out of the modeling area or leak back to the surface, it is crucial to categorize the injected volumes by a trapping inventory.

The **free plume** characterizes parts of the injected CO2 that are mobile and have potential to move out of the domain and/or leak up through the caprock. After water has reimbibed pore space occupied by CO2, a residual amount of CO2 will be left immobilized in the pores trapped by capillary forces. This quantity has become **residually trapped** and cannot leak due to its immobility. Residual trapping also benefits long-term storage of CO2 implicitly by retarding the spread of the mobile plume [31]. Accumulated CO2 below impermeable layers is **structurally trapped** due to its inability to penetrate through the confining layer, as long as this stays intact. There may be additional traps whose scale is smaller than the grid resolution. The trapping capacity of these *subscale traps* is inevitably lost and will induce an error to the net trapping capacity. While there exists methods to account for subtraps [46], this will not be accounted for herein. The free plume comprises the main concern within CO2 storage due to its leakage potential. Therefore, a central goal is to maximize the amount residually trapped, obtained by maximizing the spread of CO2.

Dissolution is a third trapping mechanism, describing the mass transport of components between phases. CO2 in gas phase can be permanently trapped by dissolving into the brine phase; this is also sometimes referred to as **solubility trapping**. The time-scale at which solubility becomes significant is sensitive to the reservoir system and may span several orders of magnitude [31]. For the Sleipner field, dissolution has shown to be notable first after hundred years of migration, becoming the dominant trapping mechanism in late-stage migration [2]. For laboratory-scale reservoirs, e.g., the FluidFlower rig, dissolution is apparent in early stages, even during injection [56]. In this thesis, mixing of components between phases is not accounted for and solubility trapping is thus omitted.

Finally, dissolved CO2 can be immobilized by chemical interactions with brine and the surrounding rock. The chemical processes involved are in many practical situations extremely slow and **mineral trapping** can safely be ignored in models of CO2 storage relevant for the FluidFlower and the Norwegian Continental Shelf [65]. Among the four trapping mechanisms, structural trapping is the only one not providing inherent security. Residually trapped CO2 is effectively immobilized, while solubility and mineral trapping facilitate removal of a distinct CO2 phase either by dissolution into the brine phase, where the larger density of brine transfer dissolved CO2 downwards, or by precipitation into a solid, immobile by nature. Structural trapping is mainly considered to facilitate storage security on a short to intermediate time scale, residual and solubility trapping at an intermediate to long time scale, and mineral trapping on a long time scale [41].

Structurally trapped CO2 is not trapped indefinitely as there might be imperfections and cracks in the sealing barrier that eventually permit diffuse leakage. Hence, it is common to separate the structurally trapped CO2 in a residual part and a mobile part. Parts of the free plume outside structural traps can be separated into a mobile part, whose fate is undetermined, and a residual part that is destined to be left immobilized and trapped indefinitely once brine reimbibes the pore space. In summary, at any time step during a simulation, one can separate the injected CO2 into six categories, as exemplified in Figure 2.3, where the mass of CO2 reaches a constant value once injection ceases. The categories are colored according to storage security, where greener and redder colors represent safer and more dangerous trapping mechanisms, respectively.

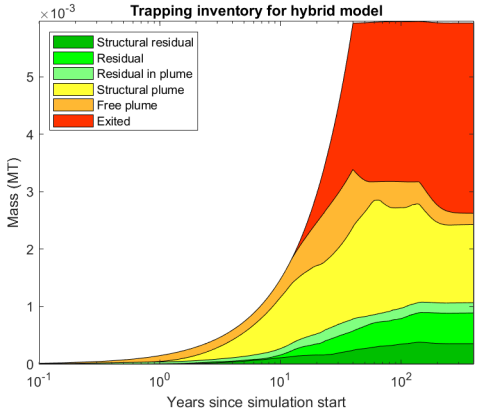

Figure 2.3: Example of a trapping inventory for a simulation with 40 years of injection and 400 years of migration, where mass of injected CO2 is separated into six categories.

To localize structural traps, we can use a technique called spill-point analysis. It is based on the assumption of an infinitesimally small trickle of CO2 and uses solely the geometry of sealing barriers of the formation to compute spill locations, spill paths and downstream traps, hierarchically described by a spill-tree [4]. The accuracy of the spill-point analysis is limited by grid resolution. If the grid cells are larger than the wavelengths of local geometry oscillations, the associated subscale traps are not resolved [59]. As a consequence, the analysis may predict erroneous spill points and spill paths that divert CO2 to other regions that are not apparent in the spill tree for the continuous setting.

2.2 Dimensional scales

As briefly outlined in Chapter 1, the dimensional scales for CO2 storage simulation are large and span multiple orders of magnitude. The relative importance of the scales vary according to the specific problem. Accounting for the physics occurring at every relevant scale and combine them into a comprehensive numerical framework is extremely demanding. Therefore, it is essential to discern the most important scales and apply simplifications to other less influential scales in order to make the computations manageable [17, 7]. In this section, we briefly outline the governing scales for geological

CO2 storage and present common simplifications.

2.2.1 Spatial scales

The spatial scales cover multiple orders of magnitude, characterized by different physical effects and contributions to the governing force balance. At the coarsest, one can distinguish between a microscopic scale, representing flow dynamics at the pore scale, and a macroscopic scale, representing flow within and between geological formations [52].

As for CO2 storage simulation, we are interested in long-term behavior at the macroscopic scale at which gravitational forces dominate. Still, fluid flow at the microscale influences dynamics at the macroscale [8]. This is particularly true near wells and other regions characterized by large pressure gradients where viscous forces are dominant. In addition, capillary forces is a pore-scale phenomena that in many cases have a notable impact on the distribution of fluids at the macroscale and may ultimately prevent CO2 from penetrating a confining caprock [4]. At gravity–capillary equilibrium, the functional dependence between saturation and capillary pressure dictates the spatial extent of the capillary fringe. If it is sufficiently small compared to the spatial scale of the formation, it can often be ignored and replaced by a sharp interface. Hence, the governing equations describe a system that is coupled through multiple spatial scales. To bridge the gap between the different scales, upscaling and multiscale modeling are essential components in long-term CO2 storage [52, 42].

2.2.2 Temporal scales

The temporal scales of CO2 storage also span multiple orders of magnitude, and the distinction closely resembles that of the spatial scales. Their relevance depends on operational demands and the particular storage site. It is customary to identify five temporal scales, sorted from smallest to largest; the nanoscale, microscale, mesoscale, macroscale, and megascale [52]. For long-term CO2 storage simulation considered in this thesis, physics at the nano- and microscale are much smaller than the long-term plume migration of interest and is omitted (i.e., considered to be in equilibrium at all times). The mesoscale defines the time used for fluids to segregate, which depends on the storage site. Herein, we will use simplified numerical models based on the assumption that CO2 and brine are vertically segregated, except internally in semi-permeable barriers for heterogeneous formations. Hence, the mesoscale is generally of minor importance apart from calculation of fluxes in the vicinity of low-permeable regions. Our concern is mainly the injection phase represented by the macroscale, and the post-injection plume migration represented by the megascale. Diffuse leakage through low-permeable layers is active on the megascale, but induces changes to the pressure field that impact the macroscale [52]. The macro- and megascale are the scales at which we want the governing equations to apply for. Thus, it is important to bridge the gap between these two scales by upscaling relevant properties from the mesoscale.

2.2.3 Upscaling

Pore-scale properties are inevitable for accurate simulation results, but representing every pore numerically would require innumerable cells and is computationally infeasible. This challenge is inherently related to the characteristic scales outlined in the previous section. Effects at the finer scales are too small to be accounted for directly in a numerical model tailored for larger scales. Still, it is possible to map fine-scale information to a coarser scale with the help of compression operators [52]. Generally, this is done by applying some sort of averaging of the properties. In practice, heterogeneities at the pore-scale are homogenized – or upscaled – to yield a single value representative at a coarser scale. VE models apply the principle of upscaling by agglomerating properties of vertically stacked fine-scale cells into a single upscaled variables for VE cells.

Permeability and porosity are rock properties defined for each grid cell and need to be upscaled. To this end, standard averaging techniques are appropriate where the layers in the formation are either parallel or perpendicular to the pressure drop [38]. For VE regions in a hybrid framework, the main flow direction is horizontal, parallel to the top surface, in which case permeability $\boldsymbol{K}_{\parallel}$ is upscaled as an arithmetic average. For flow not aligned with the axial directions, upscaling is done by solving local flow problems, a computationally more demanding approach [38].

There is also a desire to go the other way around, that is, retrieve fine-scale information from a coarse solution. Reconstructed fine-scale variables are obtained by applying reconstruction operators on coarse variables [52]. For particularly simple problems, the operator may be a closed-form expression, for instance in the sharp-interface models that will be presented in Section 5.4. In general though, reconstruction is non-trivial and will require numerical integration and nonlinear calculations [52, 47].

Finding an optimal upscaling of properties is a demanding task. Different methods for averaging properties yield different results in terms of accuracy. Absolute permeability is a tensorial quantity that complicates the search for an optimal upscaling method when the flow direction deviates from the principal directions of the layers. Despite the challenges, multiscale methods are an inevitable part of reservoir simulation to obtain a continuum description linking pore-scale behavior to the macroscale governing equations [39]. Discussion of appropriate upscaling techniques is beyond the scope of this thesis. For more details, we refer to chapter 15 of [38].

Chapter 3

Numerical methods

The multiphase flow equations presented in Section 2.1.6 include explicit and implicit nonlinear dependencies and are strongly coupled, leading to a system that is difficult to solve. Analytical solutions exist only for particularly simple problems involving homogenized reservoir properties and limited spatial and temporal scales [52]. Long-term migration of CO2 is usually a too complex problem and one has to rely on numerical computations to solve the two-phase flow system.

The governing equations (2.7) and (2.11) represent the state of the multiphase flow system in a continuous setting. A computer is not able to render the concept of a derivative, and rather has to describe the system in a discrete sense. The governing equations must correspondingly be converted to a set of algebraic equations that can be solved numerically as an approximation to the continuous counterpart.

To this end, the MATLAB Reservoir Simulation Toolbox (MRST) – an open-source software for reservoir simulations [61, 38, 39] – is used to develop the models studied in this thesis. This includes the full computational pipeline from grid setup, equation discretization, numerical solvers, and a wide range of visualization options for the results. The core of MRST is an object-oriented framework that utilizes automatic differentiation to efficiently calculate derivatives to machine precision and assemble the results into a system of linear equations that can be solved using efficient methods from numerical linear algebra.

The choice of numerical methods used to discretize a set of partial differential equations is not trivial. Depending on the equation system, the methods differ by convergence, stability, and computational efficiency and the tradeoff among these is an important consideration [55].

3.1 Two-point flux approximation

To derive a discretized version of the governing multiphase flow equations to be solved numerically, the physical properties need to take on discrete values in the spatial dimensions that constitute the computational domain. The principle of mass conservation governing the continuous system must also hold in a discrete sense to give

reliable solutions. Finite difference methods are mathematically simple to implement, but are restricted to regular geometry and mass conservation is only valid when the spatial resolution goes to zero [3]. To facilitate unstructured grids and conserve mass, it is appropriate to use finite-volume methods to discretize space that are based on conservation of mass over cells accompanied by continuity of fluxes across faces [38]. A widely used finite-volume method for multiphase systems is two-point flux approximation (TPFA), which will be derived in the following. The derivation herein is based on a two-phase flow system of a liquid and a gas, as relevant for CO2 storage simulation.

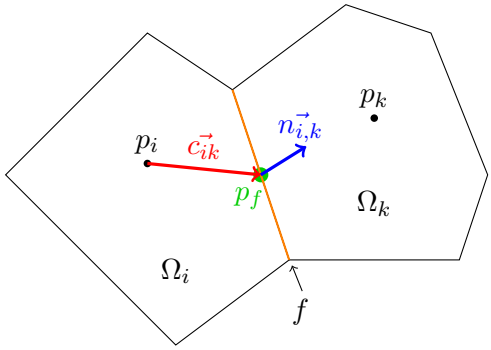

Figure 3.1: Sketch of flux of phase α across an interface $f = \Gamma_{i,k}$ shared by two cells Ω_i and Ω_k used to derive the Two-Point Flux Approximation (TPFA).

Consider the second term in (2.11). This includes the divergence of the flux function, defined over cells, as opposed to the phase flux itself defined over faces. The TPFA scheme used to discretize the phase fluxes is thus essentially a mapping from faces to cells (see Figure 3.1). To this end, we assume incompressible fluids and consider the pressure equation (2.17) now simplified to

$$\begin{aligned} \nabla \cdot \boldsymbol{u}_\alpha &= q_\alpha, \\ \boldsymbol{u}_\alpha &= -\frac{\boldsymbol{K} k_{r\alpha}}{\mu_\alpha}(\nabla p_\alpha - \rho_\alpha g \nabla z). \end{aligned} \tag{3.1}$$

Integrating over cell volume Ω_i and applying the divergence theorem gives

$$\int_{\partial\Omega_i} \boldsymbol{u}_\alpha \cdot \boldsymbol{n}\, dS = \int_{\Omega_i} q_\alpha\, dV. \tag{3.2}$$

The net flux from cell i to cell k across the interface $f = \Gamma_{i,k}$ with normal vector $\boldsymbol{n}_{i,k}$ and area $A_{i,k}$ is given by

$$\boldsymbol{u}_\alpha^{i,k} = \int_{\Gamma_{i,k}} \boldsymbol{u}_\alpha \cdot \boldsymbol{n}\, dS. \tag{3.3}$$

Assuming the interfaces between cells are continuous, each half-face has a complement with the same face area $A_{k,i} = A_{i,k}$ but opposite normal vector $\boldsymbol{n}_{k,i} = -\boldsymbol{n}_{i,k}$. Evaluating the flux at the centroid, $x_{i,k}$, of f, using the midpoint rule, we have

$$u_\alpha^{i,k} = A_{i,k}\boldsymbol{u}_\alpha(x_{i,k}) \cdot \boldsymbol{n}_{i,k} = -A_{i,k}\frac{\boldsymbol{K} k_{r\alpha,f}}{\mu_\alpha}(\nabla p_\alpha - \rho_\alpha g \nabla z)|_{x_{i,k}} \cdot \boldsymbol{n}_{i,k}, \tag{3.4}$$

where $u_\alpha^{i,k}$ is the volumetric flux directed from cell i to cell k.

The idea of the TPFA method is to approximate the flux across faces by two-point differences between adjacent cells. We assume linear pressure internally in each cell so that the pressure at the center of cell i equals the cell-averaged pressure $p_{\alpha,i}$, and the pressure at the center of face f is $p_{\alpha,f}$. The pressure gradient can then be approximated as $\nabla p_\alpha = (p_{\alpha,f} - p_{\alpha,i})/|\boldsymbol{c}_{ik}|$, and the flux from cell i across f becomes

$$\begin{aligned} u_\alpha^{i,k} &= -A_{i,k} \frac{\boldsymbol{K} k_{r\alpha,f}}{\mu_\alpha} \frac{\big((p_{\alpha,f} - p_{\alpha,i}) - \rho_\alpha g(z_f - z_i)\big)\boldsymbol{c}_{i,k}}{|\boldsymbol{c}_{ik}|^2} \cdot \boldsymbol{n}_{i,k} \\ &= T_{i,k} \lambda_{\alpha,f} \big((p_{\alpha,i} - p_{\alpha,f}) - \rho_\alpha g(z_i - z_f)\big), \end{aligned} \tag{3.5}$$

where $\lambda_{\alpha,f} = k_{r\alpha,f}/\mu_\alpha$, $\boldsymbol{c}_{i,k}$ is the vector pointing from the centroid of cell i to $x_{i,k}$ and $T_{i,k}$ is the half-face transmissibility for half-face $\Gamma_{i,k}$. The transmissibility is solely described by the cell geometry and the permeability of the rock and is therefore independent of phases α present.

Pressures are by default evaluated for cells, but the flux is evaluated at faces. Fortunately, the face pressure $p_{\alpha,f}$ can be eliminated by addressing the one-sided flux from the other side of the interface f in direction $\boldsymbol{n}_{k,i} = -\boldsymbol{n}_{i,k}$,

$$u_\alpha^{k,i} = T_{k,i} \lambda_{\alpha,f} \big((p_{\alpha,k} - p_{\alpha,f})\big) - \rho_\alpha g(z_k - z_f)\big). \tag{3.6}$$

For this to be valid, continuity of pressure across interfaces must be imposed, an inherent assumption of TPFA. To conserve mass, the one-sided fluxes across an interface must be equal but oppositely directed, $u_\alpha^{i,k} = -u_\alpha^{k,i} = u_\alpha^f$. Combining the expressions for the one-sided fluxes, $p_{\alpha,f}$ can be eliminated, and we arrive at the TPFA scheme giving the net flux across a face bounded by cells i and k,

$$\begin{aligned} u_\alpha^f &= (T_{i,k}^{-1} + T_{k,i}^{-1})^{-1} \lambda_{\alpha,f} \big((p_{\alpha,i} - p_{\alpha,k}) - \rho_\alpha g(z_i - z_k)\big) \\ &= T_{ik} \lambda_{\alpha,f} (\nabla p_\alpha - \rho_\alpha g \nabla z)_{ik}, \end{aligned} \tag{3.7}$$

where T_{ik} is the face transmissibility given as the harmonic average of the half-face transmissibilities $T_{i,k}$ and $T_{k,i}$.

The benefit of using the TPFA scheme is that it is simple to implement and strictly monotone if transmissibilities are non-zero [38]. On the other hand, the scheme is infamous for producing grid-orientation effects for anisotropic permeability and grids not aligned with the principal directions of the permeability tensor (non-K-orthogonal grids) [14], customary features for many grids used to model real formation with faults. To obtain a more accurate solution with mitigated grid-orientation effects, one generally has to rely on more advanced methods such as MPFA-O and mimetic methods [38, 39] or flow-weighted upwind interpolation [32]. In this thesis, the permeability tensor is isotropic, and grid-orientation effects for the TPFA scheme is only apparent when $\boldsymbol{c}_{i,k} \nparallel \boldsymbol{n}_{i,k}$.

3.2 Single-point upwinding

Just like the TPFA scheme uses the harmonic average to evaluate the transmissibility at an interface between two cells, phase saturations are also cell-wise properties that need to be evaluated for faces to yield well-defined phase mobility $\lambda_{\alpha,f}$ over face f. However, applying harmonic average to phase mobility will effectively immobilize a phase if one of a pair of neighboring cells has zero phase saturation [52]. Arithmetic averaging is another alternative, but this could lead to over- and undershooting of saturation [42]. Saturation is a property transported with finite speed by the hyperbolic flux term of (2.11). This suggests that phase mobilities should be upstream weighted [15]. The single-point upwind scheme works by choosing the saturation from the upwind cell, i.e., the cell which the flux across the face originates from, mathematically,

$$\lambda_{\alpha,f} = \begin{cases} \lambda_{\alpha,f}(s_{\alpha,i}) & \Theta_{\alpha,i} > \Theta_{\alpha,k}, \\ \lambda_{\alpha,f}(s_{\alpha,k}) & \Theta_{\alpha,i} < \Theta_{\alpha,k}, \end{cases} \tag{3.8}$$

where $\Theta_\alpha = p_\alpha - \rho_\alpha g z$ is the phase potential. Upwinding is simple to implement and has been proven to converge to the correct solution [15]. However, if the flux functions are not self-sharpening, the SPU-scheme suffers from numerical diffusion and better alternatives for the upwind operator are higher-order upwind schemes, high-resolution discontinuous Galerkin methods and WENO discretization [39, 42].

3.3 Discrete DIV, GRAD and UPW operators

To complete the spatial discretization of the multiphase flow system, we need discrete versions of the divergence, gradient, and upwind operators. The definition of these operators makes it possible to discretize the governing equations on a form that is independent of the topology of the grid and closely resembles the continuous operators. Moreover, the discrete operators are general in the sense that their definition is persistent between numerical methods, not only TPFA [38].

The continuous gradient operator evaluates changes in a property across a finite distance. The discrete version is formulated as the difference between two neighboring cell properties evaluated at the interface between them. In other words, the discrete `GRAD` maps from cells to faces. Let C denote the set of cells, F the set of faces, $\mathcal{R}^C$ the dimension for cells and $\mathcal{R}^F$ the dimension for faces. Then we have that the discrete gradient $\texttt{GRAD}: \mathcal{R}^C \to \mathcal{R}^F$ for a vector of scalar functions $\boldsymbol{p} \in \mathcal{R}^C$ is given by

$$\texttt{GRAD}(\boldsymbol{p})[f] = \boldsymbol{p}(c_d(f)) - \boldsymbol{p}(c_u(f)), \tag{3.9}$$

where c_d and c_u are the downstream and upstream cells for face f, respectively.

The discrete divergence operator maps from faces to cells to make the face-evaluated Darcy fluxes $\boldsymbol{u}_\alpha$ consistent with the cell-evaluated conservation equation (2.11). Hence,

we have that $\texttt{DIV}: \mathcal{R}^F \to \mathcal{R}^C$ for a vector field $\boldsymbol{v} \in \mathcal{R}^F$ is given by

$$\texttt{DIV}(\boldsymbol{v})[c] = \sum_{f \in F(c)} \boldsymbol{v}[f]\boldsymbol{I}(c = c_d) - \sum_{f \in F(c)} \boldsymbol{v}[f]\boldsymbol{I}(c = c_u), \tag{3.10}$$

where element c of vector $\boldsymbol{I}(c = c_d)$ or $\boldsymbol{I}(c = c_u)$ evaluates to one if c is downstream neighbor or upstream neighbor to face f, respectively.

The discrete version of the upwind operator maps from cells to faces and is given by

$$\texttt{UPW}(\boldsymbol{x}, \Theta)_{ik} = \begin{cases} x_i, & \texttt{GRAD}(\Theta)_{ik} < 0, \\ x_k & \texttt{GRAD}(\Theta)_{ik} > 0, \end{cases} \tag{3.11}$$

where $\texttt{GRAD}(\Theta)_{ik} = (p_k - p_i) - \rho g(z_k - z_i)$ is the difference in phase potential Θ for cell i and k.

3.4 Fully-implicit equation system

With the transport of phase properties accounted for, and discrete versions of gradient, divergence, and upwind operators, the TPFA scheme with single-point upwinding (TPFA-SPU) is generalized to

$$\begin{aligned} \texttt{DIV}(\boldsymbol{u}_\alpha)_{ik} &= q_\alpha \\ \boldsymbol{u}_{\alpha,ik} &= -\texttt{UPW}(\lambda_\alpha, \Theta_\alpha)_{ik} \cdot T_{ik} \cdot \texttt{GRAD}(\Theta_\alpha)_{ik} \\ \texttt{GRAD}(\Theta_\alpha)_{ik} &= \texttt{GRAD}(p)_{ik} - \bar{\rho}_{\alpha,ik} \cdot g \cdot \texttt{GRAD}(z)_{ik}, \end{aligned} \tag{3.12}$$

where $\bar{\rho}_{\alpha,ik}$ is the arithmetic average of phase density over cells i and k, which themselves depend on the cell pressures p_i and p_k for compressible fluids.

To complete the numerical scheme, the remaining terms in (2.11) need to be discretized. This necessitates a discretization in time, for which we choose a fully-implicit approach [42]. Combined with the discrete spatial operators, the two-phase flow system is discretized compactly as

$$\frac{(\rho_\alpha \phi s_\alpha)^{n+1} - (\rho_\alpha \phi s_\alpha)^n}{t^{n+1} - t^n} + \texttt{DIV}(\rho_\alpha \boldsymbol{u}_\alpha)^{n+1} = q_\alpha^{n+1}, \tag{3.13}$$

with $\boldsymbol{u}_\alpha$ as given in (3.12). In MRST, the AD-OO framework combined with the flexibility of so-called state functions provides a practical way to change the definition of the discrete operators `DIV`, `GRAD`, and `UPW` without modifying the governing equation system [39, pp. 152-198].

An alternative to the fully-implicit method is to adopt a sequential solution strategy where an elliptic pressure equation is decoupled from a hyperbolic transport equation. First, the former is solved implicitly for the pressure, and then the latter is solved explicitly for the saturation. Hence, the method is referred to as Implicit Pressure Explicit Saturation (IMPES). For long-term CO2 storage problems, the effect of gravity

at the sharp interface between CO2 and brine appears as a second-order term that amplifies the parabolic character of the transport equation, which usually yields a too strong coupling between the pressure and transport equations for the IMPES method to produce satisfactory results [52, 48]. Moreover, the IMPES method is generally not robust for systems approaching steady state or if saturations change significantly during a time step, as is customary for strongly buoyant CO2–brine systems [47, 11]. The fully implicit scheme, on the other hand, is unconditionally stable for arbitrary time step selection for decaying problems, but generally requires more costly computations than the IMPES method [42]. Implementation-wise, the fully-implicit scheme is more convenient in the sense that the original formulation retains when including additional physical effects [48], for instance, residual saturation, capillary pressure, and compressibility, as will be accounted for in this thesis. Despite the stability of the scheme, longer time steps induce larger errors and require more iterations to converge.

3.5 Newton's method

A fully-implicit discretization gives a system of nonlinear algebraic equations that generally cannot be solved for analytically. This is particularly true for multiphase systems where additional nonlinear dependencies arise in the relative permeability and capillary pressure functions. The equation system must be linearized and solved numerically by an iterative procedure. A frequently used approach is Newton's method which possesses quadratic convergence given that the initial solution guess $\boldsymbol{x}_0^{n+1}$ is sufficiently close to the true solution. The method iteratively finds better approximations to the zeros of a function $\boldsymbol{F}(\boldsymbol{x}_{k+1})$ by a linear Taylor expansion around the current best guess $\boldsymbol{x}_k$ of the solution. Let the abstract form of the equation system at time step $n+1$ be written as $\boldsymbol{F}(\boldsymbol{x}^{n+1}, \boldsymbol{x}^n) = 0$, where $\boldsymbol{x}^{n+1} = [\boldsymbol{p}^{n+1}, \boldsymbol{s}^{n+1}]$ is the vector of primary unknowns, chosen as the water pressures $\boldsymbol{p}^{n+1} = \boldsymbol{p}_w^{n+1}$ and the CO2 saturations $\boldsymbol{s}^{n+1} = \boldsymbol{s}_g^{n+1}$ for a CO2–brine system considered in this thesis. A first-order Taylor expansion of $\boldsymbol{F}$ around the solution guess $\boldsymbol{x}_k^{n+1}$ gives the linearized equation

$$\hat{\boldsymbol{F}}(\boldsymbol{x}_{k+1}^{n+1}) = \boldsymbol{F}(\boldsymbol{x}_k^{n+1}) + \boldsymbol{J}(\boldsymbol{x}_k^{n+1})(\boldsymbol{x}_{k+1}^{n+1} - \boldsymbol{x}_k^{n+1}) = 0, \tag{3.14}$$

where $\boldsymbol{J} = \frac{d\boldsymbol{F}}{d\boldsymbol{x}}$ is the Jacobian matrix containing derivatives with respect to the primary unknowns. A Newton iteration k for $k \in [0, K]$, where K is the imposed maximum number of iterations, at time step $n+1$ is given by

$$\boldsymbol{x}_{k+1}^{n+1} = \boldsymbol{x}_k^{n+1} - \omega \frac{\boldsymbol{F}(\boldsymbol{x}_k^{n+1}, \boldsymbol{x}_k^n)}{\boldsymbol{J}(\boldsymbol{x}_k^{n+1})}, \tag{3.15}$$

where $\omega \in [0, 1]$ is a relaxation parameter that controls the step size we take in the direction of the true solution and may facilitate convergence by avoiding overshooting [37].

For long-term CO2 storage simulations where long time intervals are customary, the solution procedure is sensitive to overshooting. Thus, despite increased simulation time, it is vital to ensure that the number of iterations K is sufficiently high and the relaxation ω sufficiently small to obtain a reliable solution. Also, it is convenient to have a **stopping criteria** that stops iterations once the absolute difference between the new and old solutions fall below a prescribed value ϵ, mathematically formulated as

$$|\boldsymbol{x}_{k+1}^{n+1} - \boldsymbol{x}_{k}^{n+1}| < \epsilon. \tag{3.16}$$

Computing the Jacobian matrix $\boldsymbol{J}$ is a computationally expensive step. In MRST, the derivatives are computed by automatic differentiation to machine precision. This avoids having to analytically derive $\boldsymbol{J}$, which streamlines the implementation. Once the residual equations and Jacobians are assembled into a linearized problem, the system can be solved by one of the many linearized solvers offered in MRST [38].

Chapter 4

The FluidFlower benchmark study

The Norwegian Continental Shelf contains several reservoir sites that serve as benchmark studies for simulating CO2 storage and migration [19]. Unfortunately, model validation is challenging due to lack of petrophysical and geometrical data from the fields, as well as empirical observations of plume shape. Recently, the new 11th SPE benchmark, labeled the FluidFlower concept, was announced and involves CO2 injection and migration inside a laboratory-generated reservoir [54]. Rather than applying our hybrid model on subsurface reservoirs, we consider the meter-scale FluidFlower rig [13]. After extracting geometrical and petrophysical data, we use MRST to set up a grid and numerical model that represents a digital twin of the physical rig. This facilitates better understanding of the physical system and development of numerical simulators [35].

The rig comprises a multilayered geological system of different sand types. An illustration of the geological layering is shown in Figure 4.1. The heterogeneous permeability, combined with a large density difference between water and gaseous CO2 and thin formation thickness makes the system a reasonable candidate for hybrid modeling. On the other hand, due to very short time-stepping required for convergence [56], the applicability of VE – and therefore the computational gain of a hybrid model – remains elusive. In this regard, the motivation for running simulations on the FluidFlower rig is twofold. It is of interest to generate a detailed grid in MRST that conforms to the geometry of the layers in the rig and will provide a basis for experimenting with different physical settings and numerical models. Despite the system not necessarily being a perfect application for hybrid modeling, it is still an excellent setting for demonstration and validation of hybrid models on realistic formations.

Despite the laboratory rig being at the meter scale, i.e., orders of magnitude smaller than the reservoir scale, it nonetheless mimics petrophysical features from the Norwegian Continental Shelf. Fluids and rocks are subject to atmospheric conditions, but dominant physical processes such as capillarity, dissolution and convective mixing are still sustained [30], and is therefore a sensible substitute for a subsurface reservoir [36]. The geometry of the unconsolidated sand layers is complex enough to be representative of geological structures found in reservoirs that are typical candidates for CO2 storage, yet sufficiently idealized to be accurately reproduced by numerical models [20].

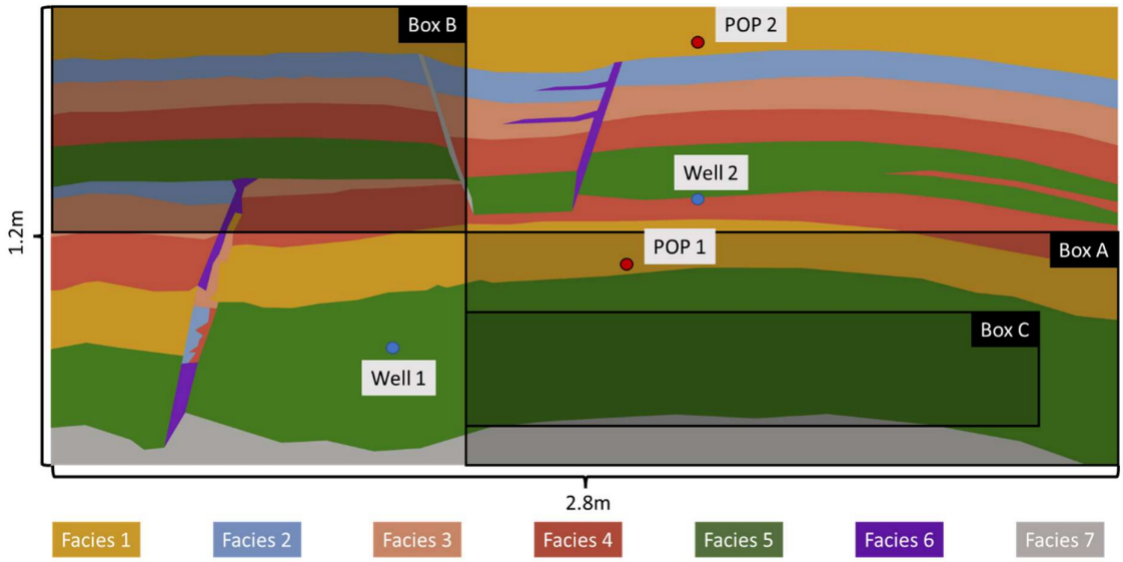

Figure 4.1: Geometry of the FluidFlower rig. The geology consist of seven unconsolidated sand facies of different petrophysical properties. Three deviated faults are apparent, where the lower left is heterogeneous, the upper left is impermeable and the upper right is high-permeable. Boxes A, B and C are for evaluating response quantities, but will be ignored herein. Image credit: [21].

Empirical observations of CO2 injected into the FluidFlower rig reveals that dissolution is an important mechanism governing the spreading of CO2, particularly in the form of convective mixing that gives a downward flux of CO2 from the buoyantly segregated plume [56, 68]. Despite its significance, the work of this thesis is devoted to implementation of residual saturation and capillary exclusion in a hybrid framework, of which dissolution has been omitted. Generation of a semi-structured grid tailored for the FluidFlower benchmark has been prioritized in this thesis. Hence, one should be aware that the results of our hybrid model will not be directly comparable to empirical observations.

4.1 Simulations using MRST

In the following, we provide a brief demonstration of using MRST to generate a semi-structured grid from the given geometrical data of the FluidFlower rig that is compatible with a hybrid framework. For a concise description of setting up and running a simulation in MRST, we refer to Section 7.1.

The dataset for the 11th SPE CSP is given as a `.geo` file providing the geometry of the rig, which is represented as a set of polygons. The data is made up of three geometrical entities, `Point`, `Line`, and `Physical Surface`. Zero-dimensional neighboring `Points` generate one-dimensional `Lines`, which are connected by a `Curve Loop` that forms a two-dimensional `Physical Surface`. Each `Physical Surface` represents a polygon uniquely describing the geometry of a facie of the rig. The reader is referred to [25] for a thorough documentation of `.geo` files.

Generating a grid that conforms to the geometrical data and is applicable for hybrid modeling is perhaps the most cumbersome step in preparing the FluidFlower case for simulation in MRST. Generally, the facies are nicely stacked vertically, but the presence of three faults complicates the geometry. Moreover, to generate and describe the connectivity between VE cells in a hybrid framework, the associated parts of the grid must be logically indexed [44]. One option is to use a logical corner-point grid to discretize the reservoir, but the inconsistent geometry of the reservoir poses geometrical challenges [12]. Instead, we adopt a composite VE–fine grid that facilitates hybrid modeling by letting the structured parts of facies satisfying (5.18), be discretized as VE cells, while remaining regions, including faults and pinch-outs, are discretized as full-dimensional cells.

To permit consistent calculation of fluxes in VE–to–fine transition regions, we introduce a set of full-dimensional, Cartesian buffer cells around faults. In other words, all transitions between VE and fine regions are composed of structured cells, as depicted in Figure 6.1c. Subgrids representing individual facies must be merged into a global grid representative of the entire reservoir. The glued counterpart is unstructured and does not have any trivial logical indexing. This must be derived manually from the logical indices of neighboring, structured cells in the parent grid which the current subgrid is glued to.

The collection of MATLAB scripts presented in the following is based on functionality from MRST. At the time of writing, it is located in the folder `grid` of this thesis' GitHub repository `Master-Thesis/FluidFlower` at https://github.com/vetlenev/Master-Thesis, with utility functions located in the folder `SPE11-utils`. The scripts are projected to be a part of MRST's contribution for the 11th SPE benchmark.

We start by prescribing horizontal and vertical dimensions N_x and N_z for a global, virtual, Cartesian background grid of size $L_x = 2.8$ meters and $L_z = 1.2$ meters, dictating the resolution of subgrids. Individual polygons, representing subgrids of given facies, are created as instances of a class `PolygonGrid` by the call

```
poly = PolygonGrid(all_polys, poly_num);
```

Here, `all_polys` is a cell array of data points for each polygon, and `poly_num` is the index of the polygon to create an instance for. The class contains functionality for generating subgrids, finding overlapping nodes, interpolating nodes, and gluing together grids. It also includes subclasses particular for edgecases, e.g., `PinchOuts` and `Faults` for representing complex, discountinuous geometries as unstructured cells. The stepwise procedure for generating the grid is as follows.

Step 1. We start by generating a subgrid for the topmost polygon extending horizontally over the entire global domain. From the polygonal points, we extract the top and bottom sides of the polygon and distribute $N_x + 1$ nodes uniformly between the leftmost and rightmost node. For a given `side` of the polygon, this is done by

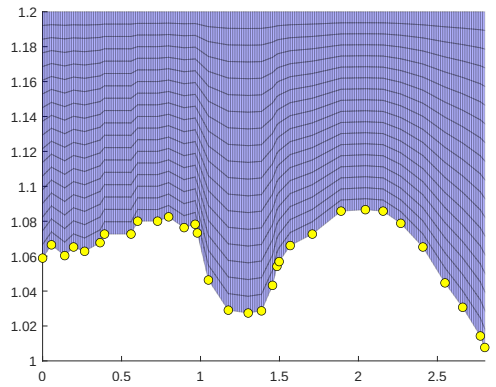

(a) Discretization of top layer (facie 1). Yellow points are adjusted nodes from polygonal data.

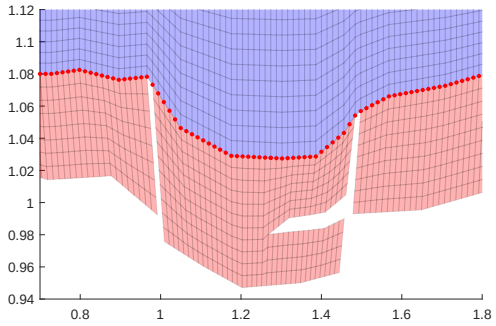

(b) Top layer glued together with its bottom neighbor layer (facie 2). Red points are nodes connecting the two subgrids.

Figure 4.2: Discretization of the topmost part of the FluidFlower rig. Nodes on black lateral curves are interpolated between top and bottom surface of associated subgrid.

the following class function acting on the polygon instance `poly`, with `boundary_mask` accessing the desired side, and `G_glob` being the global, virtual background grid.

```
[poly, closest_mask] = correct_dx(poly, G_glob, boundary_mask, side);
```

Next, we set the z-coordinate of nodes in the subgrid closest to each polygonal point, indexed by `closest_mask`, to the z-coordinate of this point. The remaining nodes on the top and bottom surfaces are interpolated between the closest adapted nodes.

```
poly = interpolateZ_remaining(poly, closest_mask, boundary_mask, side);
```

Finally, internal nodes are interpolated between associated nodes on the top and bottom surface by calling the class function

```
poly = interpolateInternal(poly);
```

The complete subgrid for the topmost polygon is shown in Figure 4.2a.

Step 2. We build the global grid by working downwards from the topmost polygon, sequentially gluing together neighboring polygons. Most of the structured subgrids are generated by the following generic algorithm.

1. Find polygonal points of the current subgrid overlapping with the bottom sides of upper neighboring polygons. Select all nodes from the bottom sides that are interpolated between the overlapping nodes.

   ```
   [nodes, pts] = findOverlappingNodes(poly, poly_other, side);
   ```

The call returns a list of overlapping polygonal points `pts` and the associated interpolated `nodes` from the upper neighbor `poly_other`. The returned nodes dictate the horizontal dimension and define the top side of the current subgrid. In our case, we are finding overlapping nodes with an upper neighbor, in which case we set `side` to *'top'*. The vertical dimension is set as the subgrid's fractional vertical size of the global grid.

2. Create a curvilinear subgrid based on the horizontal and vertical dimension, and allocate coordinates of overlapping nodes to the top side.

3. Coordinates of the bottom side are created in the same way as the topmost polygon.

4. West and east sides are handled by calling

```
poly = interpolateSide(poly);
```

The function runs a procedure similar to how nodes were assigned at the top and bottom sides of the subgrid, but now performed along a vertical side.

5. Interpolate internal nodes between associated top-bottom pair and west-east pair of nodes.

For a polygon with index `poly_num` with upper neighbors of indices `poly_num_upper`, the generic algorithm is executed by the following routine

```
[poly_list, nodes] = glueToUpperPolygon(all_polys, poly_list, poly_num,
                                        poly_num_upper, nodes, G_glob);
```

The routine generates a subgrid of the polygon with index `poly_num` by querying the associated polygonal points from the global list `all_polys` and assigns it to the list `poly_list` of glued polygons. New overlapping nodes in the glued grid are appended to the `nodes` structure. The complete subgrid of the top polygon (facie 1) glued together with its lower neighbor (facie 2) is shown in Figure 4.2b.

Step 3. There layers containing pinch-outs require separate treatment. First, we create a curvilinear background grid with interpolated nodes on the top and bottom surfaces as in step 2. Then, we locate the line segment between pairs of top and bottom nodes that lie closest to the tip of the pinch-out. The associated top and bottom nodes are used to separate the background grid into three parts; a region A *not* containing a pinch-out, a region B *above* the pinch-out and a region C *below* the pinch-out. The three subgrids are illustrated in Figure 4.3. Finally, we interpolate internal nodes between pairwise top-bottom and west-east nodes, for each subgrid. For a layer containing a pinch-out with index `poly_num_pinch`, the procedure is collectively performed by the routine

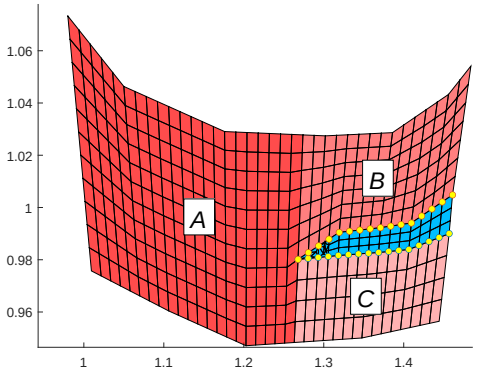

Figure 4.3: Discretization of the topmost pinch-out. The three neighbors A, B and C are structured Cartesian subgrids. Neighbor connections for the pinch-out are marked by yellow nodes.

```
[poly_list, nodes] = gluePinchOuts(all_polys, poly_list,poly_num,
                                   poly_num_upper, poly_num_pinch,
                                   nodes, G_glob);
```

Step 4. When we arrive at the lower tip of the two upper faults, the layer below overlaps with multiple upper subgrids of different horizontal dimensions. Here, we generate a parent subgrid that we split into multiple smaller subgrids, each overlapping with a single upper neighbor. The resulting subgrid has a nonuniform horizontal resolution, as shown in Figure 4.5 for the layer below the two upper faults.

Step 5. Faults are discretized as full-dimensional Cartesian quadrilaterals and Delaunay triangles. The relevant polygons are created as instances of `Faults`, with functionality for generating quadrilaterals and triangulated cells that conform to each other. The first step for generating a composite grid is to separate each fault into polygons that are either discretized as quadrilaterals or as triangles.

The purpose of the quadrilaterals is to make the cells inside a fault as regular as possible and minimize angles with the axial directions to reduce grid-orientation effects. We start by extracting the external boundaries of each polygon; that is, all faces on the boundary of the polygon that are shared with a structured subgrid outside the fault, using the function `findOverlappingNodes`. For the internal boundaries, i.e., faces shared with other polygons *inside* the fault, we interpolate nodes at a uniform distance between provided polygonal data points. Then, we create a Cartesian subgrid with nodes displaced to match the boundaries. If opposing boundaries contain different number of nodes, it is necessary to **collapse** some of the faces, as indicated by the yellow nodes in Figure 4.4.

The triangulated grid is hand-crafted to conform to the boundary of a fault, and also supports a quasi-random distribution of nodes inside the fault to give a finer triangulation. We use the built-in MATLAB function

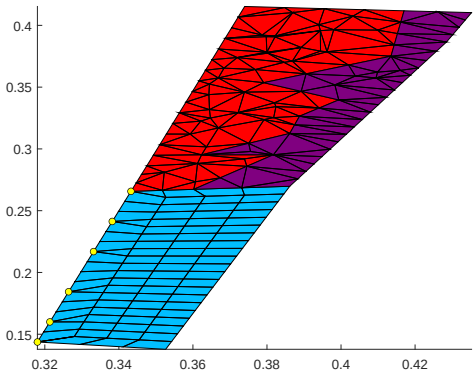

Figure 4.4: Discretization of a subset of the bottom, heterogeneous fault. The red and purple subgrids (facie 2 and facie 4) are triangulated to reasonably discretize the complex transition between the facies. The light blue subgrid (facie 6) is discretized as quadrilaterals where collapsed faces are indicated by yellow nodes.

```
dt = delaunayTriangulation(nodes, C);
```

to create a triangulation of the nodes given in `nodes` that constitute the mesh of the fault, and a closed loop of nodes `C` that constrains the triangulation. Since the function creates a triangulated grid based on the convex hull of `nodes`, we need to use the built-in method `isInterior` to only extract the triangles that are inside the polygon. The resulting triangulated grid is obtained as follows,

```
inpoly = isInterior(dt);
dt_inpoly = dt.ConnectivityList(inpoly,:);
G = triangleGrid(nodes, dt_inpoly);
```

where the second line extracts the connections between the triangles inside the polygon. To create more regularly shaped triangles, we supply additional points inside the original triangles by a quasi-random distribution where triangles of larger volumes are assigned more points, as described in [53]. Our implementation allows adjusting the local resolution at a particular fault or pinch-out. Because the additional points are inside the polygon, we retain the constraint `C` but expand our array of `nodes` with the added nodes `new_nodes`. The new triangulation is generated by

```
dt = delaunayTriangulation([nodes; new_nodes], C);
```

Grid generation and node adaption for polygons constituting a fault are collectively performed by the following class function,

```
poly = poly.distributeAndCollapseNodes(west, east, top, bottom, size, dim)
```

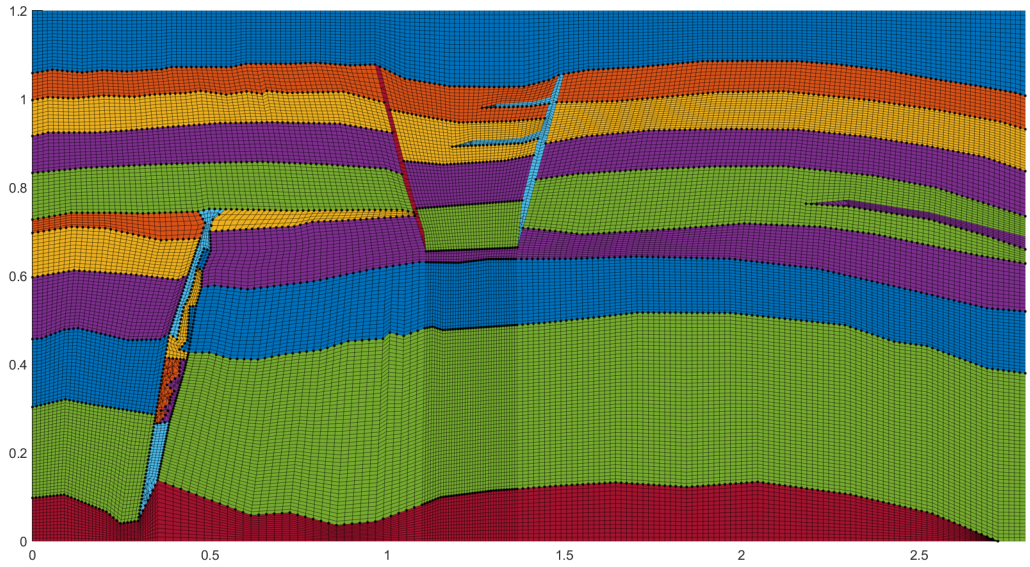

Figure 4.5: Semi-structured grid of FluidFlower generated in MRST, conforming perfectly to complex structures including faults, pinch-outs and anticlines.

where the first four arguments specify nodes on the boundaries of the polygon instance `poly`. The last two arguments give the sizes and dimensions of the global grid, used to distribute nodes for interpolation.

Once all subgrids are generated, they are glued together to form a new polygon structure `poly.glued` by sequentially calling the following function for each subgrid `G_sub`.

```
poly.glued.G = glue2DGrid_FF(poly.glued.G, G_sub);
```

The function returns an updated version of the global grid `poly.glued.G` with the subgrid `G_sub` glued to the grid structure. Nodes from one of the grids that are within a specified Euclidean distance from nodes in the other grid determines the intersections to be glued. The resulting semi-structured composite grid is shown in Figure 4.5. The desired resolution can be provided by the user, and the grid is compatible with hybrid modeling.

Once we have a computational grid for the FluidFlower rig, the most cumbersome work is done. The remaining steps are more generic for most type of simulations performed in MRST, and follows the workflow outlined in Section 7.1.

Chapter 5

Vertical equilibrium models

Vertical equilibrium (VE) models significantly reduce the computational complexity while retaining as much as possible of the conventional black-oil formulation [62]. It is based on the assumption that the reservoir fluids are rapidly segregated vertically after an initial distribution. This assumption is justified in regions where the fluids are subject to predominant buoyant forces caused by large density differences [8]. In a typical CO2 storage scenario the lateral scales are much larger than the vertical scales, implying that lateral flow will dominate vertical flow in the long term. Combined with significant density differences between CO2 and brine, even in supercritical state [41], the segregation process can in many cases be assumed to happen instantaneously, at least compared to the time scale of lateral up-dip migration [48]. Thus, the fluids are in hydrostatic equilibrium and their distribution in the vertical direction is given by an analytical expression. The vertical dimension is effectively removed by integrating the flow equations vertically between the bottom and top of the formation. Therefore, the VE models are based on a set of upscaled variables that only depend on lateral dimensions. Vertical pressure and saturation are trivial to reconstruct analytically from the hydrostatic assumption.

Because the VE columns are agglomerations of fine-scale cells, petrophysical and geometric fine-scale properties including absolute permeability, porosity, and transmissibility need to be averaged on the coarse scale [38]. Section 2.2.3 justified that arithmetic averaging of absolute permeability $\boldsymbol{K}$ is the correct upscaling method for VE cells. If β, τ and H define the bottom, top and thickness of the VE column, respectively, and $\boldsymbol{x}$ the location of fine-scale cells, the arithmetic average of $\boldsymbol{K}$ is given by

$$\bar{\boldsymbol{K}} = \frac{1}{H} \int_{\beta}^{\tau} \boldsymbol{K}(\boldsymbol{x}) \, d\vec{x}. \tag{5.1}$$

The discretized arithmetic average used in the numerical solver for a VE column made up of n fine cells of thicknesses h_i is

$$\bar{\boldsymbol{K}} = \frac{1}{\sum_{i=1}^{n} h_i} \sum_{i=1}^{n} h_i \boldsymbol{K_i}. \tag{5.2}$$

Assuming uniform spacing between the fine cells in the vertical direction, $h_i = H/n$, and the arithmetic average simplifies to

$$\bar{\boldsymbol{K}} = n \sum_{i=1}^{n} K_i. \tag{5.3}$$

Notice that for aquifers with homogeneous permeability we get $\bar{\boldsymbol{K}} = K_i = \boldsymbol{K}$, even for nonuniform grid spacing.

Porosity is a volumetric quantity that is upscaled by adding up pore volumes from fine cells with volume V_i and normalizing by total bulk volume of the VE column. The discretized average is

$$\bar{\phi} = \frac{1}{\sum_{i=1}^{n} V_i} \sum_{i=1}^{n} V_i \phi_i. \tag{5.4}$$

Because face areas of cells are incorporated in the definition of transmissibility, the upscaled transmissibility of a VE column is obtained by summation over m fine-scale transmissibilities for faces bounding the VE column, without scaling by formation thickness,

$$\bar{T} = \sum_{i=1}^{m} T_i. \tag{5.5}$$

Here, it is implicitly assumed that the fine-scale face transmissibilities T_i are computed as harmonic averages of associated half-faces

$$T_i = \Big(\frac{1}{T_{i,d}} + \frac{1}{T_{i,u}} \Big)^{-1}, \tag{5.6}$$

where $T_{i,d}$ and $T_{i,u}$ are half-face transmissibilities for face i from downstream cell d and upstream cell u, respectively.

When the fluid phases are in hydrostatic equilibrium, the vertical pressure distribution follows a hydrostatic profile dictated by the weight of the overlying column of fluid. The hydrostatic pressure distribution for phase α in an aquifer with slope angle θ (see Figure 5.1) is given by

$$p_\alpha(z) = P_\alpha + \rho_\alpha g (z - z_0) \cos(\theta), \tag{5.7}$$

where P_α is the pressure at a prescribed depth z_0.

The VE model is obtained by vertically integrating the governing equation for multiphase flow, (2.12). Retaining generality, we let x_1 and x_2 be the coordinates of the plane that defines the up-dip direction of the aquifer (lateral coordinates) and z the coordinate orthogonal to this plane (transversal coordinate). Hence, for a sloped aquifer, z will deviate slightly from the direction of gravity, but we will herein label it as the vertical direction regardless and use the terms interchangeably. A sketch of the system

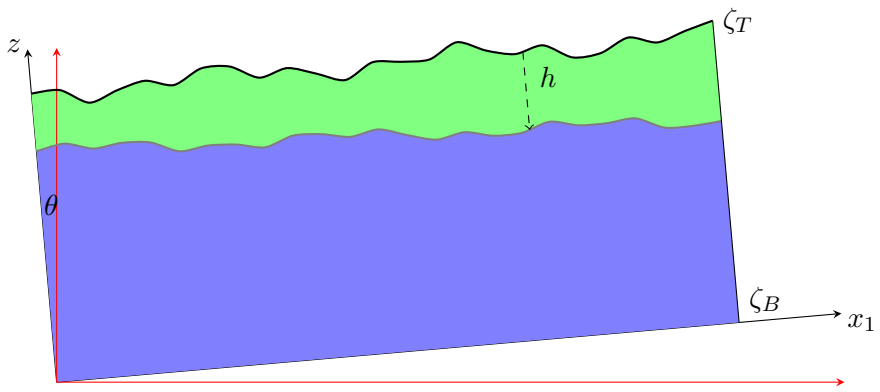

Figure 5.1: Sketch of a reservoir system used to derive the basic sharp-interface, vertical equilibrium model. Red axes represent the main coordinate frame, while black axes represent the coordinate used in the derivation, rotated by an angle θ. The system is only shown in 2D as we assume the cross-section is uniformly extruded in the x_2-direction.

under consideration is shown in Figure 5.1. The vertically integrated version of (2.12) is

$$\int_{\zeta_T}^{\zeta_B} \frac{\partial}{\partial t}(\rho_\alpha s_\alpha \phi)\, dz - \int_{\zeta_T}^{\zeta_B} \nabla \cdot \left[\rho_\alpha \lambda_\alpha \boldsymbol{K}(\nabla p_\alpha - \rho_\alpha \boldsymbol{g})\right] dz = \int_{\zeta_T}^{\zeta_B} q_\alpha\, dz. \tag{5.8}$$

Under the assumption of no flow across the top and bottom boundaries, Leibniz' rule can be used to move the divergence outside the integral. Assuming that rock properties in the principal directions are independent of each other, the permeability tensor can be decomposed into a lateral component (denoted $\|$) and a transversal component (denoted $\perp$),

$$\boldsymbol{K} = \begin{bmatrix} \boldsymbol{K}_\| & \boldsymbol{0} \\ \boldsymbol{0} & \boldsymbol{K}_\perp \end{bmatrix}. \tag{5.9}$$

Consequently, the Darcy flux can be decomposed, and lateral flow will only depend on the lateral pressure gradient [4], in which case the vertical component vanishes and we get

$$\frac{\partial}{\partial t}\int_{\zeta_T}^{\zeta_B} (\rho_\alpha s_\alpha \phi)\, dz - \nabla_\| \cdot \int_{\zeta_T}^{\zeta_B} \rho_\alpha \lambda_\alpha \boldsymbol{K}_\| (\nabla_\| p_\alpha - \rho_\alpha \boldsymbol{g}_\|)\, dz = \int_{\zeta_T}^{\zeta_B} q_\alpha\, dz. \tag{5.10}$$

At the time of writing, the VE framework developed in the `mrst-co2lab` module of MRST [60] only supports horizontal compressibility, not vertical compressibility, as the latter is usually much smaller and can often be neglected in practice [5]. Hence, phase densities ρ_α do not vary with depth and are thus independent of p_α given by (5.7). The phase density of a VE column is then simply the fine-scale density sampled at a reference depth z_0, i.e., $\rho_\alpha = \rho_\alpha(x, y, z = z_0)$, which can be chosen arbitrarily. The resulting VE model is simplified. We refer to [5] for an exposition of a VE model with vertical compressibility included.

The lateral components of the pressure gradient and gravity term are independent

of the transversal direction, yielding

$$\frac{\partial}{\partial t}\Big[\rho_\alpha \int_{\zeta_T}^{\zeta_B} (s_\alpha \phi)\, dz\Big] - \nabla_\parallel \cdot \Big[\rho_\alpha \int_{\zeta_T}^{\zeta_B} \boldsymbol{K}_\parallel \lambda_\alpha\, dz\Big](\nabla_\parallel P_\alpha - \rho_\alpha(\boldsymbol{g}_\perp + \boldsymbol{g}_\parallel)) = \int_{\zeta_T}^{\zeta_B} q_\alpha\, dz, \tag{5.11}$$

where $\boldsymbol{g}_\perp = g_\perp \cos(\theta)\nabla_\parallel \zeta_T$. Next, we define the following upscaled variables

$$\Phi = \int_{\zeta_T}^{\zeta_B} \phi\, dz, \tag{5.12}$$

$$S_\alpha = \Phi^{-1} \int_{\zeta_T}^{\zeta_B} \phi s_\alpha\, dz, \tag{5.13}$$

$$\boldsymbol{\kappa} = \int_{\zeta_T}^{\zeta_B} \boldsymbol{K}_\parallel\, dz, \tag{5.14}$$

$$\boldsymbol{\Lambda}_\alpha = \boldsymbol{\kappa}^{-1} \int_{\zeta_T}^{\zeta_B} \frac{k_{r\alpha}}{\mu_\alpha} \boldsymbol{K}_\parallel\, dz, \tag{5.15}$$

$$Q_\alpha = \int_{\zeta_T}^{\zeta_B} q_\alpha\, dz, \tag{5.16}$$

which produces the following upscaled form of the governing multiphase flow equations

$$\begin{aligned} \frac{\partial}{\partial t}(\rho_\alpha \Phi S_\alpha) + \nabla_\parallel \cdot (\rho_\alpha \boldsymbol{F}_\alpha) &= Q_\alpha, \\ \boldsymbol{F}_\alpha = -\boldsymbol{\Lambda}_\alpha \boldsymbol{\kappa}(\nabla_\parallel P_\alpha - \rho_\alpha(g_\perp \cos(\theta)\nabla_\parallel \zeta_T + \boldsymbol{g}_\parallel)). & \end{aligned} \tag{5.17}$$

All upscaled quantities, except phase pressures, are based on vertical integration over the aquifer thickness to remove the dependence on the vertical coordinate. Upscaled phase pressures are instead sampled at a reference depth since vertical averaging of pressure is not a physically meaningful variable [6].

The final set of equations has been reduced from three dimensions to two and has the same structural form as the original fine-scale equations. Hence, the numerical solution strategy is principally the same as for the fine-scale complement, in which case the implementation is straightforward as long as the upscaled variables have analytical expressions [52]. For more general problems, e.g., with a capillary fringe, numerical integration or tabulation is needed to compute the upscaled quantities before the equation system can be solved [47].

A question to ask is when the VE assumption is valid. A qualitative measure from [28] suggests the following segregation time for a two-phase CO2–brine system,

$$t_s \sim \frac{H\phi\mu_w}{k^*_{rw} k_z \Delta\rho g}. \tag{5.18}$$

Here, H is the aquifer thickness, μ_w the viscosity of brine, k^*_{rw} the characteristic relative permeability of brine, and k_z the vertical permeability of the layer. If $t_s \ll T$, where T is simulation time, and additionally, t_s is shorter than discretization step Δt_i at time

step i, vertical equilibrium is a reasonable approximation. If the latter is not satisfied, a VE model will overestimate the vertical migration speed of the non-wetting phase [9].

Despite a dimensionality reduction, the VE model is not destined to yield worse results than the full-dimensional model in regions where the VE assumption applies. In fact, a VE model may give better solutions in these regions than a full-dimensional model because the former is not limited by a finite grid resolution in the vertical direction, but rather relies on analytical reconstruction of mobile plume depth [49, 40]. This gives a "correct" value of the relative permeability at the sharp interface [62]. Since the coarse equations have the same form as the full-dimensional equations, the principle of mass conservation still applies for the coarse variables. Thus, (5.17) satisfies a conservation law in the new solution space and therefore has a physical interpretation [52]. The fact that the coarse-scale variables depend on the fine-scale variables induces some challenges, as the fine-scale variables somehow need to be estimated. A common approach is to approximate the fine-scale variables through reconstruction operations [52]. This will usually involve evaluating nonlinear equations [47], but for particularly simple problems, the reconstructions have simple analytical expressions. This will be outlined in the upcoming section.

In the fully-implicit VE scheme, one of the phase pressures is chosen as a primary variable, typically the water pressure P_w. As for the second primary variable, there are two alternatives, either CO2 saturation, S_n, yielding the S-formulation, or the height of mobile CO2, h, yielding the h-formulation. The difference between the two formulations is minor from a computational point of view [47], but the S-formulation retains most of the elements from conventional reservoir simulators and is therefore a more practical choice for implementing new physical effects [48].

5.1 Reconstruction

Conceptually, the solution of the upscaled system of equations portrays a two-dimensional surface from which the upscaled variables have been derived. Fine-scale variables representing the three-dimensional solution can be reconstructed from the coarse-scale variables using the VE assumption [52]. More importantly than visualizing the solution, reconstructed variables are needed to evaluate upscaled quantities. If not stated otherwise, reconstructed variables are denoted by a hat.

Let P_α denote the upscaled pressure for phase α at a reference depth, herein chosen to be the top surface of an aquifer plane with dip angle θ. The fine-scale pressure is reconstructed by knowledge of the depth of the top surface ζ_T and the weight of the underlying column of fluid,

$$\hat{p}_\alpha(x, y, z) = P_\alpha(x, y) + \rho_\alpha(x, y)g(z - \zeta_T(x, y)) \cos(\theta). \tag{5.19}$$

As for the other variables, the reconstruction complexity depends on the assumptions behind the model [52]. In general, one cannot expect upscaled quantities like relative

permeability and capillary pressure to have analytical expressions [47].

There are specific cases for which the procedure is simplified; for a sharp-interface model, we expect the phase saturations to be piecewise constant with a discontinuity at the sharp interface. Hence, fine-scale saturation can be directly retrieved from a closed-form expression solely dependent on the height of the column of mobile CO2, h. For a fine cell with top depth T and bottom depth B, as part of a VE column with top depth ζ_T, the reconstruction becomes

$$\hat{s}_n(h) = \begin{cases} 1, & \zeta_T + h > B, \\ \frac{\zeta_T + h - T}{B - T}, & T \leq \zeta_T + h \leq B, \\ 0, & \zeta_T + h < T. \end{cases} \tag{5.20}$$

The simplified sharp-interface model will be presented shortly.

5.2 Capillary pressure

The VE equivalent of capillary pressure is referred to as pseudo-capillary pressure and equals the difference between the upscaled non-wetting and wetting phase pressures. The pseudo-capillary pressure ensues directly from the hydrostatic pressure (5.7) and can be uniquely described by a reference depth z_0 and the depth h of the interface between CO2 and brine. Let p_e be the capillary entry pressure at h. The pseudo-capillary pressure is defined in terms of the reference height, here chosen to be the level of the caprock $z_0 = \zeta_T$, and is given by

$$P_c(h) = p_e + (\rho_w - \rho_n) g h(x, y) \cos(\theta), \tag{5.21}$$

where the dependence of the upscaled variable on the lateral coordinates is emphasized through $h = h(x, y)$.

The pseudo-capillary pressure serves a similar role in the VE framework as the fine-scale capillary pressure does in the full-dimensional framework in the sense that it provides a coarse-scale closure equation,

$$P_c(S_w) = P_n - P_w. \tag{5.22}$$

In the S-formulation, the coarse saturation S_w is then obtained by inverse computation of the coarse capillary–saturation function, $P_c^{-1}(S_w)$.

5.3 Residual saturation

To implement residual saturation in the VE framework, an additional state variable is required, namely the maximum non-wetting saturation obtained in each cell through simulation history, labeled $S_{n,\max}$. This variable is used to keep track of the maximum

height $h_{\max}$ of the layer of residual CO2 residing below the mobile plume and above the region of pure brine. In a sharp-interface model, $S_{n,\max}$ is related to $h_{\max}$ through

$$S_{n,\max} = \frac{h_{\max}}{H}(1 - s_{wr}), \tag{5.23}$$

where s_{wr} is the residual water saturation.

In the general case of a capillary fringe, S_n is not given by an analytical expression of h and $h_{\max}$ but requires a nonlinear computation of $P_c^{-1}(S_w)$. Analytical relations are only provided in the simplified case of negligible capillary forces, as will be outlined in the next section.

5.4 Development of the sharp-interface model

Following the work of [51], a measure of the importance of capillary forces, or more precisely the extent of any developed capillary fringe is given by

$$L_{cf} = \frac{1}{\Delta\rho g \cos(\theta)} \frac{\partial p_c}{\partial s_n}. \tag{5.24}$$

Hence, a large density difference between supercritical CO2 and brine at subsurface conditions yields a thin capillary fringe. Moreover, if the capillary pressure is weak, it varies weakly with saturation, i.e., $\frac{\partial p_c}{\partial s_n}$ is small. In this case, (5.24) is usually several orders of magnitude smaller than the thickness of the formation, and often comparable to the vertical extent of fine cells, in which case it can be neglected in practice. This is often satisfied for large-scale CO2 storage, in which case the transition between CO2 and brine can adequately be represented as an infinitely thin zone [18]. This limiting case is referred to as a sharp-interface assumption, characterized by a discontinuous transition from maximum brine saturation below the interface to maximum CO2 saturation above the interface.

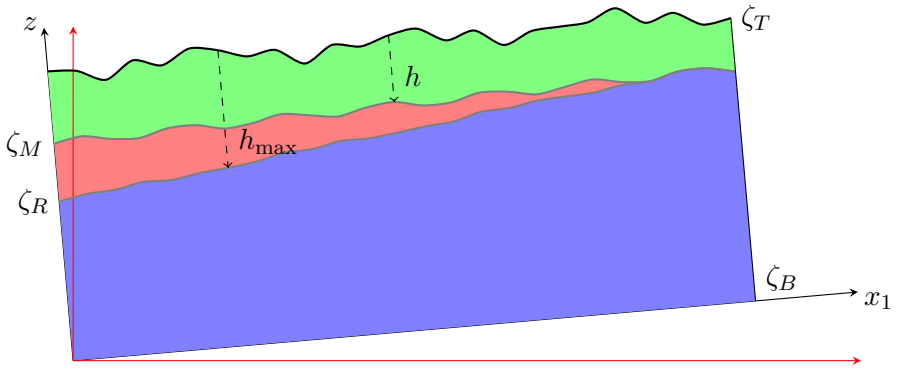

Figure 5.2: Sketch of a reservoir system used to derive the sharp-interface VE model with residual saturation. The definitions are similar to that of Figure 5.1 but now including a red zone of residual CO2 saturation of height $h_{\max}$. The sharp interface is located at ζ_M and the historically maximum depth reached is ζ_R.

Only a few years after injection stops, residual trapping contributes significantly

to the trapping inventory [41]. Residually saturated CO2 is vertically segregated from the historically deepest point reached (denoted ζ_R) and up to the bottom of the mobile plume (denoted ζ_M). Inclusion of residual saturation only amounts to modest changes in the sharp-interface model, and the modified reservoir system is illustrated in Figure 5.2. A VE column is separated into three regions delineated by sharp interfaces; a bottom region extending from ζ_B to ζ_R consisting of pure brine, a middle region extending from ζ_R to ζ_M of brine and residual CO2, and a top region from ζ_M to ζ_T of mobile CO2 and residual brine. The capillary pressure at the sharp interface located at ζ_M equals the entry pressure required for CO2 to invade the brine-filled region.

In a sharp-interface model, the fine-scale saturation is reconstructed by a simple analytical expression in terms of plume height $h = \zeta_M - \zeta_T$ and the historically maximum plume height $h_{\max} = \zeta_R - \zeta_T$ without having to rely on the inverse computation of a fine-scale capillary-saturation function, $s^{cap}(p_c)$. Let $H = \zeta_B - \zeta_T$ be the height of the VE column. The variables h and $h_{\max}$ are related to the primary variables S_n and $S_{n,\max}$ through

$$\begin{aligned} h &= H\frac{S_n(1-s_{wr}) - S_{n,\max}s_{nr}}{(1-s_{wr})(1-s_{nr}-s_{wr})}, \\ h_{\max} &= H\frac{S_{n,\max}}{1-s_{wr}}. \end{aligned} \tag{5.25}$$

This enables us to swap primary variables from S_n and P_w to h and P_w, in which case we adopt the h-formulation of the governing equations.

The closed-form expression of reconstructed CO2 saturation in terms of h and $h_{\max}$ is given by

$$\hat{s}_n = \begin{cases} 1-s_{wr}, & z < \zeta_T + h, \\ s_{nr}, & \zeta_T + h \le z \le \zeta_T + h_{\max}, \\ 0, & z > \zeta_T + h_{\max}. \end{cases} \tag{5.26}$$

For discrete cells, z spans a finite interval, and the reconstructed saturation for a cell is given by an average of $1-s_{wr}$ and s_{nr}, depending on how much of the cell covers mobile and residual CO2. Let $\zeta_M = \zeta_T + h$ and $\zeta_R = \zeta_T + h_{\max}$ be the depth and maximum depth of the top plume, respectively. Using the following scaling factors for the mobile and residual parts,

$$a_M = \max\Big[\min\big(\frac{\zeta_M - T}{H},\, 1\big),\, 0\Big], \tag{5.27}$$

$$a_R = 1 - a_M - \max\Big[\min\big(\frac{B - \zeta_R}{H},\, 1\big),\, 0\Big], \tag{5.28}$$

where T and B denote the top and bottom of the fine cell, the reconstructed fine-scale saturation becomes

$$\hat{s}_n = a_M(1-s_{wr}) + a_R s_{nr}. \tag{5.29}$$

The model equations become slightly more involved with the inclusion of residual saturation but are still subject to simplification when considering a sharp interface and

assuming that the rock parameters have negligible variation in the vertical direction [4]. From the fact that only one fluid phase is mobile on each side of the interface, the fine-scale mobilities in the respective regions are simplified,

$$\lambda_w(s_w) = \begin{cases} 0, & s_w = s_{wr} \\ \lambda_{w,e}, & s_w = 1 - s_{nr} \\ 1, & s_w = 1, \end{cases} \tag{5.30}$$

and similarly for CO2-mobility

$$\lambda_n(s_n) = \begin{cases} 0, & s_n < 1 - s_{wr} \\ \lambda_{n,e}, & s_n = 1 - s_{wr}. \end{cases} \tag{5.31}$$

Here, $\lambda_{\alpha,e} = k_{r\alpha,e}/\mu_\alpha$ are mobilities calculated from the endpoint relative permeabilities $k_{r\alpha,e}$, which are evaluated at their respective maximum saturations in a system where both phases are present. In general $k_{r\alpha,e} < 1$ due to additional resistance from the residual content of the other phase [24].

To facilitate a simple and appealing derivation of the sharp-interface model, we assume that porosity is independent of depth and the permeability tensor has independent lateral and transversal components. The upscaled rock quantities then reduce to

$$\Phi = H\phi, \qquad \boldsymbol{\kappa} = H\boldsymbol{K}_\parallel. \tag{5.32}$$

Upscaled phase saturations become

$$S_\alpha = (H\phi)^{-1}\phi \int_{\zeta_T}^{\zeta_B} s_\alpha \, dz = \frac{1}{H}\int_{\zeta_T}^{\zeta_B} s_\alpha \, dz, \tag{5.33}$$

with wetting and non-wetting components

$$\begin{aligned} S_w &= \frac{1}{H}\Big[\int_{\zeta_T}^{\zeta_T+h} s_{wr}\, dz + \int_{\zeta_T+h}^{\zeta_T+h_{\max}} (1-s_{nr})\, dz + \int_{\zeta_T+h_{\max}}^{\zeta_B} 1\, dz\Big] \\ &= \frac{1}{H}\big[s_{wr}h + (1-s_{nr})(h_{\max}-h) + (H-h_{\max})\big], \\ S_n &= \frac{1}{H}\Big[\int_{\zeta_T}^{\zeta_T+h} (1-s_{wr})\, dz + \int_{\zeta_T+h}^{\zeta_T+h_{\max}} s_{nr}\, dz\Big] \\ &= \frac{1}{H}\big[(1-s_{wr})h + s_{nr}(h_{\max}-h)\big]. \end{aligned} \tag{5.34}$$

Coarse-scale mobilities are simplified to

$$\boldsymbol{\Lambda}_\alpha = (H\boldsymbol{K}_\parallel)^{-1}\boldsymbol{K}_\parallel \int_{\zeta_T}^{\zeta_B} \frac{k_{r\alpha}(\hat{s}_\alpha)}{\mu_\alpha}\, dz = \frac{1}{H}\int_{\zeta_T}^{\zeta_B} \frac{k_{r\alpha}}{\mu_\alpha}\, dz. \tag{5.35}$$

Because the reconstructed fine-scale saturations $\hat{s}_\alpha$ are piecewise constant in each VE

column, the integrals reduce to closed-form expressions,

$$\begin{aligned}\boldsymbol{\Lambda}_w &= \frac{1}{H\mu_w}\Big[\int_{\zeta_T}^{\zeta_T+h} k_{rw}(s_{wr})\,dz + \int_{\zeta_T+h}^{\zeta_T+h_{\max}} k_{rw}(1-s_{nr})\,dz + \int_{\zeta_T+h_{\max}}^{\zeta_B} k_{rw}(1)\,dz\Big]\\ &= \frac{1}{H\mu_w}\big[k_{rw,e}(h_{\max}-h) + (H-h)\big],\\ \boldsymbol{\Lambda}_n &= \frac{1}{H\mu_n}\Big[\int_{\zeta_T}^{\zeta_T+h} k_{rn}(1-s_{wr})\,dz + \int_{\zeta_T+h}^{\zeta_T+h_{\max}} k_{rn}(s_{nr})\,dz + \int_{\zeta_T+h_{\max}}^{\zeta_B} k_{rn}(0)\,dz\Big]\\ &= \frac{h}{H\mu_n}k_{rn,e}.\end{aligned} \tag{5.36}$$

In general, if the source term depends on the depth, i.e., $q_\alpha = q_\alpha(z)$, the upscaled quantity Q_α must be computed numerically.

Pseudo-capillary pressure is derived using the fact that the fine-scale capillary pressure should equal the entry pressure at the depth $\zeta_T + h$ of the CO2–brine interface. If upscaled CO2 pressure is defined at the top of a VE column and upscaled water pressure at the bottom (the standard used in this thesis), pseudo-capillary pressure in the sharp-interface model is given by

$$P_c = P_n - P_w = p_e - \rho_n g h - \rho_w g(H-h). \tag{5.37}$$

Because the height h of mobile CO2 in the sharp-interface assumption is directly related to the coarse saturations S_n and $S_{n,\max}$, the closure equations are trivial to compute. We now have a VE model of four equations satisfying the principle of mass conservation at the coarse scale that can be solved numerically.

In an S-formulated sharp-interface VE model, the upscaled variables have analytical expressions and the system of algebraic equations closely resembles conventional full-dimensional models [47], simplifying the implementation. Thus, the S-formulation is the preferred solution approach for the problems considered in this thesis.

The solution-reconstruction procedure in a fully-implicit VE model is summarized by the pseudo-code presented in Algorithm 1.

Algorithm 1: Fully-implicit sharp-interface model

Result: Solution vector $\{\boldsymbol{P_w}, \boldsymbol{S_n}\}$ of multiphase flow equations

initialize fine state: $\{p_w^0,\ s_n^0\}$;

vertically integrate to get coarse state: $\{P_w^0,\ S_n^0\}$;

define upscaled rock properties;

$\Phi = \int_{\zeta_T}^{\zeta_B} \phi\, dz$;

$\boldsymbol{\kappa} = \int_{\zeta_T}^{\zeta_B} \boldsymbol{K}_{\parallel}\, dz$;

for $t = 0$ *to* T **do**

compute phase densities;

$\rho_\alpha = \rho_\alpha^{t+1}(x, y)$;

reconstruct fine-scale saturations;

$h = h(S_n^{t+1}, S_{n,\max}^{t+1}), \quad h_{\max} = h_{\max}(S_{n,\max}^{t+1})$;

$\hat{s}_n = a_M(h)(1 - s_{wr}) + a_R(h, h_{\max}) s_{nr}$;

$\hat{s}_w = 1 - \hat{s}_n$;

compute the following upscaled variables;

$P_c^{t+1}(S_w^{t+1}) = p_e - \rho_n g h - \rho_w g(H - h)$;

$P_n^{t+1} = P_w^{t+1} + P_c^{t+1}$;

$S_w^{t+1} = 1 - S_n^{t+1}$;

$\boldsymbol{\Lambda}_\alpha^{t+1} = \boldsymbol{\kappa}^{-1} \int_{\zeta_T}^{\zeta_B} \boldsymbol{K}_{\parallel} \lambda_\alpha(\hat{s}_\alpha)\, dz$;

$Q_\alpha^{t+1} = \int_{\zeta_T}^{\zeta_B} q_\alpha\, dz$;

compute coarse Darcy flux;

$U_\alpha^{t+1} = -\boldsymbol{\Lambda}_\alpha^{t+1} \boldsymbol{\kappa} (\nabla_{\parallel} P_\alpha^{t+1} - \rho_\alpha^{t+1} (g_\perp \cos(\theta) \nabla_{\parallel} \zeta_T + \boldsymbol{g}_{\parallel}))$;

solve Equation 5.17 for new state $\{P_w^{t+1},\ S_n^{t+1}\}$;

end

reconstruct fine-scale solution $\{\hat{\boldsymbol{p}}_{\boldsymbol{w}},\ \hat{\boldsymbol{s}}_{\boldsymbol{n}}\}$ algebraically from Equation 5.19 and Equation 5.29;

Chapter 6

Hybrid models

The scientific literature includes numerous studies of vertical equilibrium models that yield excellent results compared to conventional 3D reservoir simulators for large-scale migration in homogeneous formations [52]. However, for multilayered systems with semi-permeable barriers, significant vertical flow may be present even during late stages of migration and the VE assumption is potentially violated. Hybrid models are developed to provide a middle ground by combining the computational efficiency of VE models with the accuracy of fine-scale models. The idea is to utilize VE models in regions where the time of segregation (5.18) is sufficiently short, and 3D models in remaining regions [26]. The resulting reservoir grid becomes mixed-dimensional with regions modeled by different dimensionality. The governing coarse and fine equations of a hybrid model are similar to those of a pure VE and full-dimensional model, respectively, but evaluation of properties requires special treatment at the transition between regions of different discretizations to ensure a consistent mapping between fine and coarse variables [44].

The literature contains several studies of hybrid models applied to different reservoir systems and with different physical effects included, e.g., [10, 26, 43, 11] to mention some. To our knowledge, there are no acknowledged studies of fully-implicit sharp-interface hybrid models that include residual saturation and capillary entry pressure. The latter has a great impact on diffuse leakage through low-permeable interbeddings, for instance through thin shales at the Utsira formation [16]. Accumulation of CO2 below low-permeable barriers – and possible penetration – is not solely dependent on entry pressure, but also greatly influenced by residual trapping that favors less penetration by mitigating pressure build-up [23]. Residual saturation and capillary pressure have been included in several hybrid models from the literature, e.g., [10, 28, 27], reporting great agreement with full-dimensional models. They are based on the IMPES method which is generally less robust than a fully-implicit method for long-term CO2 migration with complex flow [52]. Hence, it is valuable to implement similar effects in a fully-implicit hybrid model.

A challenge with multilayered systems is homogeneous regions with semi-low permeability in which the vertical flow of CO2 is too slow to assume instantaneous gravity segregation of CO2 and brine. Still, vertical dynamics are predominant and it

is reasonable to simplify the modeling by assuming strictly vertical migration until the phases have segregated. Guo et al. [27] developed a hybrid model that relaxes the VE assumption by solving the vertical dynamics as one-dimensional problems rather than assuming instantaneous segregation, coupling coarse-scale pressure equations with fine-scale transport equations for vertical flow. Herein, we retain the fully-implicit system and apply a different relaxation approach for columns where the VE assumption is potentially violated. The implementation of a relaxed vertical equilibrium (RVE) assumption will be presented shortly.

6.1 Transition zones

A challenge of hybrid models is how to handle fluxes between mixed-dimensional **discretization regions** (an assembly of adjacent cells with matching interfaces of non-zero transmissibility and similar discretization types) in a consistent way. For interfaces between two fine cells or two VE cells of similar type, the flux is evaluated as described in the governing equations for the full-dimensional and VE models, respectively. What is not as trivial is how the flux is evaluated at the transition between *different* discretization regions. Discretization regions for a full-dimensional, a VE and a hybrid model are shown in Figure 6.1.

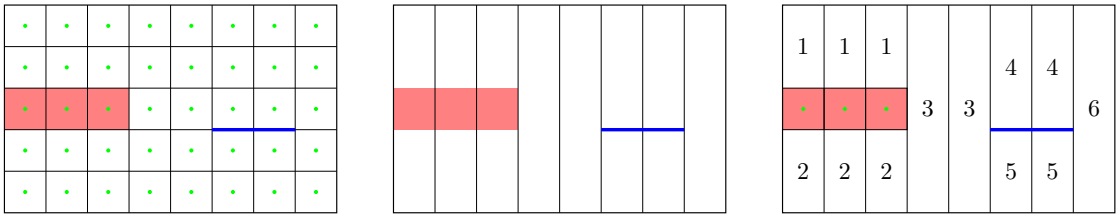

(a) Full-dimensional **(b)** VE **(c)** Hybrid

Figure 6.1: Discretization regions for a) a full-dimensional, b) a VE, and c) a hybrid model. Semi-permeable cells and faces are marked in red and blue, respectively. Black lines delineate cells, where fine cells are marked with green dots. Numbering represents VE cells as part of same discretization region.

The horizontal transition between a fine cell and a coarse (VE) cell is characterized by a discontinuous top and/or bottom surface, and so is the interface of VE cells of different heights. To obtain a consistent flux at the interface, we introduce a set of virtual cells in the larger neighboring cell that match the vertical extent of the smaller neighboring cell, exemplified in Figure 6.2 for the transition between two VE columns of different heights.

Proper modeling of diffuse leakage through sealing barriers is essential for reliable simulations of multilayered systems in the middle-to-long term. Møyner and Nilsen [44] performed a study of the hybrid model on the Utsira formation in which diffuse leakage through three semi-permeable barriers was represented by cell faces, i.e., as vertical VE–to–VE connections, with a fixed transmissibility multiplier of 10^{-4}. In other words, the impact of entry pressure and pressure gradients internally in the semi-permeable

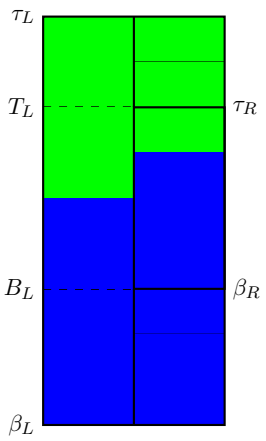

Figure 6.2: Illustration of a transition zone between two VE columns of different discretization. Green color represents CO2 and blue color represents water. Dashed lines are introduced virtual cells, τ and β are top and bottom of parent cell, while T and B are top and bottom of virtual cells. For VE columns with no virtual cells added, $\tau = T$ and $\beta = B$. Subscript L refers to the left column, and R to the right column. The right column is the smallest one and has no virtual cells added, i.e., $T_R = \tau_R$ and $B_R = \beta_R$.

barriers was represented indirectly by the transmissibility, which has been shown to be an inaccurate approximation [1]. A better solution would be to have the flexibility to represent the semi-permeable barriers as full-dimensional cells. This is substantiated by the fact that the thickness of intraformational semi-permeable barriers may span several meters [16, 66]. For the thickest layers, internal flow could turn out to be significant, and a cell representation is most appropriate. In either case, approximating the entry pressure by a transmissibility multiplier will incur errors.

With a full-dimensional rather than a lower-dimensional face representation of semi-permeable barriers, horizontal fluxes are calculated either as VE–to–fine or VE–to–VE horizontal transition regions. For vertical transitions, VE–to–VE is replaced by VE–to–fine, but since the interface is only across a single face we only need to introduce a single virtual fine cell for VE columns in the transition zone, as illustrated in Figure 6.3. Additional computations arise from using full-dimensional cells to represent sealing barriers. The benefit is that important physical effects such as capillary pressure, viscous forces, and heterogeneous permeability are accounted for explicitly, and transmissibility is correctly related to the permeability of the layer through its original formulation (3.5), not approximately through a multiplier.

Diffuse leakage can be modeled either as VE–to–VE transitions using lower-dimensional faces (Figure 6.3a) or as VE–to–fine transitions using full-dimensional cells (Figure 6.3b). In either case, virtual fine cells are introduced above and below the sealing barrier (white cells with numbers 2 and 4) where reconstructed variables are calculated at the center of the virtual cells. We refer to [44] for a thorough description of different discretization regions.

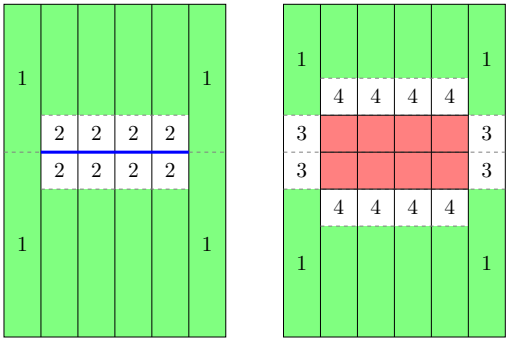

(a) Faces **(b)** Cells

Figure 6.3: Transition regions for a semi-permeable barrier modeled as a) faces and b) full-dimensional cells. Virtual cells are marked by dashed lines. VE cells are shaded in green. Semi-permeable barriers represented as faces are marked in blue, and as cells marked in red. Numbers represent the following transition zones: 1) VE–to–VE horizontal, 2) VE–to–VE vertical, 3) VE–to–fine horizontal, and 4) VE–to–fine vertical.

6.2 Capillary pressure

The pseudo-capillary pressure (5.37) is unchanged by including residual saturation, because the fine-scale capillary pressure still equals the entry pressure at the CO2–brine interface, and the hydrostatic pressure distribution of a phase is unaffected by the residual content of another immobile phase. Reconstruction of capillary pressure at the center of virtual fine cells simply uses the capillary function on the reconstructed saturation in the virtual cells, $\hat{p}_c = p_c(\hat{s}_w)$. For a sharp-interface model, there is no need for numerical integration since the expression is analytically given by $\hat{p}_c = p_e$. The pseudo-capillary pressure in a virtual VE cell, $\hat{P}_c$, is calculated from (5.37) with h and H corresponding to the local height of the mobile plume and height of the virtual cell, respectively.

6.3 Residual saturation

Implementation of residual saturation in a hybrid model is essentially the same as for conventional VE models, but with particular care at transitions between regions of different discretizations. The main factor that invalidates the conventional implementation of residual saturation is the long time scale for diffuse leakage, which is likely to yield advancing plume tips due to slow gravity drainage even during late stages of the simulation [26]. Because the tail of the plume moves faster than the tip [33], the migrating plume may be immobilized before it reaches the top of the column. The vertical distribution of CO2 given in Figure 5.2 is no longer globally applicable, as there may be additional residual parts that have penetrated through semi-permeable barriers, isolated from the top plume.

A full-dimensional model is more appropriate in the vicinity of semi-permeable barriers [26], but with numerous semi-permeable barriers, the number of full-dimensional cells quickly grows too large for a hybrid model to yield any significant advantages.

Alternatively, we retain a reduced dimensionality but relax the VE assumption. This is inspired by the work of [26] and [9], which resolved slow vertical drainage as independent one-dimensional problems as part of a pseudo-VE model. Herein, the relaxation is based on a naive treatment of coarse fluxes from the relevant transition zones that enables us to discern the top segregated plume from the isolated plumes originating from semi-permeable barriers.

6.3.1 Immobilization of plume

In the VE assumption, residual CO2 is distributed from the bottom of the mobile part of the plume with height h and down to the maximum depth reached, with height $h_{\max}$. However, for multilayered systems we expect CO2 to spread into multiple separate plumes among the semi-permeable barriers. In addition to the buoyantly segregated plume below the caprock, residual CO2 must also be traced back to individual semi-permeable barriers, where diffuse leakage and retardation of the leading edge of the plume may immobilize the plume before reaching the caprock [33].

The expressions for h and $h_{\max}$ in (5.25) does not account for additional late-stage migrating plumes. This may lead to an overestimate of the mobile plume at the top of a VE column. Fortunately, it is possible to account for these additional plumes, potentially immobilized, in the VE framework by a naive consideration of fluxes. The idea is to locate VE columns transitioning to a lower layer (that be VE–to–fine or VE–to–VE transitions), and keep track of the accumulated upward fluxes across the interfaces. If CO2 penetrating a semi-permeable barrier enters a high-permeable region where VE is justified, strict vertical migration is imposed [28]. Hence, we can explicitly compute how much CO2 in the associated columns originates from the bottom of the semi-permeable barriers.

The edges of semi-permeable barriers are also regions where one may expect late-stage vertical flow. To this end, we need to keep track of fluxes across *horizontal* VE–to–VE transition zones, as these are responsible for late buoyant migration beyond the edges. Combined with the coarse saturations S_n and $S_{n,\max}$ we have enough information to reproduce the vertical saturation distribution in RVE columns.

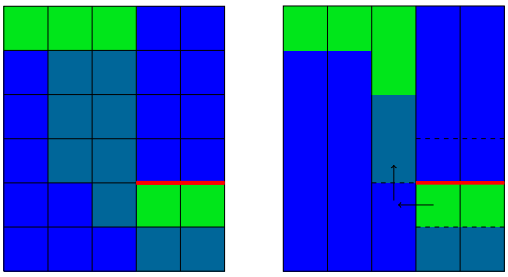

(a) Full-dimensional **(b)** Hybrid

Figure 6.4: Distribution of CO2 resulting from a horizontal flux at the endpoint of a semi-permeable barrier (marked in red) for a) a full-dimensional and b) a hybrid model. Weak green color represents residual CO2 while strong green represents mobile CO2. Dashed lines in the hybrid grid delineate virtual cells. Black arrows denote fluxes from semi-permeable layers.

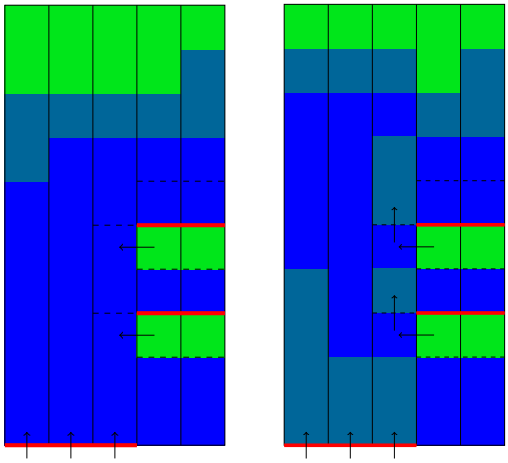

(a) Old **(b)** New

Figure 6.5: Distribution of CO2 resulting from bottom and horizontal fluxes from semi-permeable barriers (marked in red) for a) the old implementation and b) the new implementation of residual saturation for a multilayered system.

The VE assumption will by definition assume exclusive buoyant migration in RVE columns and is therefore prone to underestimate the width of the plume beyond the edge of sealing barriers, as illustrated in Figure 6.4. This is an inevitable artifact of the VE assumption. The residual part is reconstructed assuming the net horizontal flux from a VE column is converted to an equivalently large vertical flux upward from the top of the neighboring virtual VE column. Reconstructed saturation in a multilayered system is shown in Figure 6.5, exemplifying a case with multiple immobilized plumes in the same (middle) column, originating either from the bottom or the edges of semi-permeable barriers.

6.3.2 Implementation

To incorporate residual saturation in the hybrid model, but without the extra computational burden of representing VE cells as fine cells, we introduce two new state variables,

$$u^b_\Sigma = \sum_{i=1}^{n} u^b_i \Delta t_i, \qquad u^b_{\Sigma,\max} = \sum_{i=1}^{k} u^b_i \Delta t_i. \tag{6.1}$$

The former, u^b_Σ, is the cumulative sum of CO2-phase fluxes up from the bottom of a VE column transitioning to a low-permeable barrier below, summed from the start of the simulation and up to the current time step n. The second variable, $u^b_{\Sigma,\max}$, describes the same but accumulation is only up to time step $k \leq n$, at which the current $S_{n,\max}$ was reached. Similarly, we store the horizontal fluxes across VE columns of different discretizations,

$$u^h_\Sigma = \sum_{i=1}^{t} u^h_i \Delta t_i, \qquad u^h_{\Sigma,\max} = \sum_{i=1}^{k} u^h_i \Delta t_i. \tag{6.2}$$

The variables $u^*_{\Sigma,\max}$ $(* = b, h)$ are necessary to ensure a consistent maximum saturation $S^T_{n,\max}$ of the CO2 plume residing below the caprock, distinguished from late-migrating

plumes emanating from semi-permeable barriers.

The original expressions (5.25) for h and $h_{\max}$ are retained for all cells except columns with an RVE assumption. There are three types of VE cells that obtain a modified expression and are relabeled "RVE" cells.

- VE cells transitioning to a layer below in a VE–to–VE or VE–to–fine vertical transition region, queried by $B = \beta$ and $T \neq \tau$, where B and T are the bottom and top depths of the virtual fine cells as part of the VE cell above the interface, using the convention in Figure 6.2.

- VE cells transitioning to neighboring VE cells of different discretization, represented as a VE–to–VE horizontal transition zone. We only select the virtual VE cells *not* sharing top surface with the parent VE cell, queried $T \neq \tau$. For the two virtual VE cells shown in Figure 6.4b (middle column), only the lower one satisfies this criterion.

- VE cells that satisfy the two above requirements, associative of rising plumes both as diffuse leakage and from edges of semi-permeable barriers, as illustrated in Figure 6.5.

There are two regions to distinguish: mobile CO2, $S_{n,\text{mob}}$, buoyantly segregated at the top with maximum saturation reached $S^T_{n,\max}$, and net late-stage CO2 migrating from semi-permeable barriers, S_u. The latter is calculated from the accumulated fluxes from semi-permeable barriers connected to a VE cell. For a VE cell j, the calculation is

$$S_u(j) = \frac{u^b_\Sigma + \sum_{i\in\mathcal{D}_h(j)} u^h_{\Sigma,i}}{\Phi V} = S_{u,b} + \sum_{i\in\mathcal{D}_h(j)} S_{u,hi}, \tag{6.3}$$

where $u^h_{\Sigma,i}$ is the accumulated flux from the i'th barrier out of $\mathcal{D}_h$ horizontally neighboring semi-permeable barriers, with $S_{u,hi}$ the corresponding saturation, and $S_{u,b}$ the saturation resulting from diffuse leakage from a semi-permeable barrier at the bottom of the VE cell.

Similarly, we have that the accumulated flux from semi-permeable barriers up to the time step where the current $S_{n,\max}$ value was first reached yields a saturation

$$S_{u,\max}(j) = \frac{u^b_{\Sigma,\max} + \sum_{i\in\mathcal{D}_h(j)} u^h_{\Sigma,\max,i}}{\Phi V}. \tag{6.4}$$

The maximum saturation from the top plume is given by

$$S^T_{n,\max} = S_{n,\max} - S_{u,\max}, \tag{6.5}$$

from which the maximum reached height of the top plume evaluates to

$$h_T = \max\left(H\frac{S^T_{n,\max}}{1 - s_{wr}},\, 0\right). \tag{6.6}$$

The height of the plume rising from a semi-permeable barrier i is given by the expression

$$d_{ki} = \max\Big(\bar{H}_i(1 - \bar{S}_{F,ki}),\ h_T\Big), \tag{6.7}$$

where for a semi-permeable barrier at the bottom of the column we have $ki = b$ and

$$\bar{H}_i = H, \qquad \bar{S}_{F,b} = S_{u,b}/s_{nr},$$

where H is the height of the parent VE cell. For semi-permeable barriers at VE–to–VE horizontal transition regions we have $k = h$ and

$$\bar{H}_i = H_i, \qquad \bar{S}_{F,hi} = S_{u,hi}/(\frac{H_i}{H}s_{nr}), \tag{6.8}$$

where $H_i = T_i - \tau$ is the height from the top of the parent VE cell, τ, to the top of the virtual VE cell i, T_i. The expression for normalized plume saturation $\bar{S}_{F,hi}$ follows from the fact that the maximum amount of immobilized CO2 in the parent cell originating from semi-permeable barrier i is $\frac{H_i}{H}s_{nr}$, so that $\bar{S}_{F,hi} \in [0,\ 1]$.

The height of all individual plumes stacked on top of each other from the bottom, labeled h_B, is given by

$$h_B = \max(H - \Delta h_B,\ h_T), \qquad \Delta h_B = (H - d_b) + \sum_{i \in \mathcal{D}_h(j)} (H_i - d_{hi}), \tag{6.9}$$

where $h_B = h_T$ implies that the VE column is entirely filled with residual saturation. If this is the case, the phases are effectively buoyantly segregated, VE is satisfied, and CO2 from additional fluxes will immediately accumulate below the top.

With these variables, we can define the net height reached by all plumes in an RVE cell to be

$$h_{\max} = h_T + (H - h_B). \tag{6.10}$$

Finally, total residual content in the column is

$$S_{nr,tot} = s_{nr}\frac{h_{\max}}{H}, \tag{6.11}$$

which yields the mobile saturation

$$S_{n,\mathrm{mob}} = S_n - S_{nr,tot}, \tag{6.12}$$

and the height of the mobile plume on the top

$$h = H\Big(\frac{S_{n,\mathrm{mob}}}{1 - s_{wr} - s_{nr}}\Big). \tag{6.13}$$

With modified expressions for h and $h_{\max}$, the coarse mobilities and associated phase fluxes are also modified for RVE columns.

6.4 Dissolution

Dissolution of CO2 has been shown to be significant trapping mechanism for a wide range of time scales and for conditions typically found on the Norwegian Continental Shelf [31]. Recent experiments reveal that this is also true for the FluidFlower rig, particularly in the form of convective mixing [56, 68]. Hence, to provide realistic simulations, our models should beneficially implement dissolution. This is substantiated by the fact that convective mixing is a common phenomenon for formations satisfying VE [57, 4]. Even though convective dissolution is a process acting on the pore scale, it is possible to implement it in a black-oil formulation at the Darcy flow regime by representing it as an upscaled dissolution rate [57]. Extending the fully-implicit hybrid model with dissolution has not been a priority of this thesis, and is instead postponed to future work. We will return to the issue of implementing dissolution in a hybrid setting in Chapter 8.

6.5 Reconstruction

For hybrid models, reconstruction of fine-scale variables follows the same principle as for pure VE models, but the expressions are modified for VE columns by accounting for the hysteretic variable $S_{n,\max}$. Moreover, for relaxed VE columns, we also need to account for accumulated fluxes from semi-permeable barriers to locate isolated late-migrating plumes.

6.5.1 Fine-scale saturation

The reconstructed CO2 saturation still has an analytical expression in the sharp-interface setting but now includes a dependency on the residual parts left in the wake of migrating plumes, which for RVE columns are described in terms of the depths d_b and d_{hi} given by (6.7).

Reconstruction of fine-scale CO2 saturation given by (5.29) requires modifications for RVE columns to account for immobilized plumes. The residual content is now given by

$$a_R = \min(a_{RT} + a_{RS},\ 1 - a_M), \tag{6.14}$$

where a_{RT} represents the residual part for the top plume (i.e., a_R in (5.29)) and a_{RS} the residual part for plumes originating from sealing barriers. Let T and B be the top and bottom depths of fine cell c of a VE column with height $H = B - T$, let $H_{tk} = \tau + H_k$ denote the top depth of the virtual cell bounded by sealing barrier k and let $h_{tk} = \tau + h_k$ be the top depth of the plume rising from sealing barrier k. There are four scenarios in which this plume may contribute to the residual content of fine cell c,

- Cell c is located partly below the bottom of the plume and not above the top of the plume, given by the indicator $I_b^k = (h_{tk} \leq T < H_{tk}) \ \wedge \ (B > H_{tk})$.

- Cell c is located partly above the top of the plume and not below the bottom of the plume. Indicator: $I_t^k = (h_{tk} < B \leq H_{tk}) \ \wedge \ (T < h_{tk})$.
- Cell c covers the entire plume. Indicator: $I_o^k = (T < h_{tk}) \ \wedge \ (B > H_{tk})$.
- The entire cell c is inside the plume. Indicator: $I_i^k = (h_{tk} \leq T < H_{tk}) \ \wedge \ (h_{tk} < B \leq H_{tk})$.

With the help of these indicators, we have

$$a_{RS} = \frac{1}{H} \sum_{k \in \mathcal{D}_h} [(H_{tk} - T)I_b^k + (B - h_{tk})I_t^k + (H_{tk} - h_{tk})I_o^k + HI_i^k]. \tag{6.15}$$

The expression for the mobile ratio a_M is unchanged for RVE columns.

6.5.2 Fine-scale phase pressure

Reconstruction of fine-scale pressure in hybrid models needs to account for a nonuniform distribution of CO2 and brine, but still reproduce a hydrostatic profile in the limits when VE is applicable. In this regard, [28] proposes a more general saturation-weighted reconstruction,

$$\hat{p}_w = P_w - \int_{z_r}^{\beta} ((\bar{s}_n \rho_n + \bar{s}_w \rho_w)g + \bar{s}_n \frac{\partial p_c}{\partial z}) \, dz, \tag{6.16}$$

where z_r is the point of interest where pressure should be reconstructed and $\bar{s}_\alpha$ are phase saturations normalized by the residual content.

The reconstruction (6.16) is specific for models where *two* mobile phases may be present inside a cell, e.g., in the pseudo-VE model from [28]. In the sharp interface, relaxed VE framework developed in this thesis, weighting by saturation becomes redundant because there will only be a single mobile phase present in each of the three distinct parts of a VE column shown in Figure 5.2. The same is true for RVE columns (see Figure 6.5). With these simplifications, (6.16) becomes

$$\begin{aligned} \hat{p}_w &= P_w - \int_{z_r}^{d} \rho_n g \, dz - \int_{d}^{\beta} \rho_w g \, dz \\ &= P_w - \rho_n g(d - z_r) - \rho_w g(\beta - d), \end{aligned} \tag{6.17}$$

where $d = \max(\tau + h, z_r)$ is the maximum of the depth of the mobile plume $\tau + h$ and the point of interest z_r, and β is the bottom of the column. We emphasize that z_r is chosen to be the center of virtual fine cells and the bottom of virtual VE cells. In virtual fine cells, CO2 pressure is evaluated as $\hat{p}_n = \hat{p}_w + p_c(\hat{s}_w)$, where $\hat{s}_w$ is the reconstructed brine saturation. The coarse CO2 pressure in a virtual VE cell is given by $\hat{P}_n = \hat{P}_w + \hat{P}_c(\hat{h})$, where $\hat{P}_w$ is the coarse water pressure reconstructed at the bottom of the virtual VE cell and $\hat{P}_c$ is the pseudo-capillary pressure in the virtual VE cell that depends on the local height $\hat{h} = \tau + h - T$ of the plume in this cell.

6.6 Coupled equations

The governing equations in a hybrid framework couple the fine-scale and coarse-scale formulations. From the fact that the upscaled equations in a VE model have principally the same form as the fine-scale equations, we are in a position to formulate a general set of algebraic equations for a hybrid model. For this, we need two indicator functions. The first one, $C(i)$, returns the column-index of cell i in the hybrid grid, where all cells vertically stacked on top of each other share the same column-index. Secondly, the indicator $V(i)$ returns the discretization region of cell i. As an example, two cells j and k that are vertically stacked but delimited by a sealing face share the same column-index, i.e., $C(j) = C(k)$, but are from different VE regions, i.e., $V(j) \neq V(k)$. All fine cells c_f are assigned the same discretization, by default $V(c_f) = 1$. Moreover, let $N(i)$ denote the set of neighboring cells connected to cell i. With this in hand, the discrete hybrid equations are compactly written as

$$\frac{\Gamma_{i,\alpha}^{n+1} - \Gamma_{i,\alpha}^{n}}{\Delta t_n} - \sum_{j \in N(i)} f_{ij}(\Pi_{\alpha}^{n+1}) = \Psi_i^{n+1}, \tag{6.18}$$

The definitions of the accumulation and source terms now depend on the particular cell i by

$$\begin{aligned} \Gamma_{i,\alpha} &= \begin{cases} \rho_\alpha \phi_i s_{i,\alpha}, & V(i) = 1, \\ \rho_\alpha \Phi_i S_{i,\alpha}, & V(i) \neq 1, \end{cases} \\ \Psi_i &= \begin{cases} q_i, & V(i) = 1, \\ Q_i, & V(i) \neq 1. \end{cases} \end{aligned} \tag{6.19}$$

The variable set Π_α used to compute the flux terms f_{ij} depends on the column-indices and discretization regions of cell i and j, and are either given as fine-scale variables, coarse-scale variables or a combination thereof,

$$\Pi_{ij,\alpha} = \begin{cases} (p_\alpha, \rho_\alpha, \lambda_\alpha(s_\alpha), z), & \text{internal fine–to–fine}, \\ (P_\alpha, \rho_\alpha, \Lambda_\alpha(\hat{s}_\alpha), Z), & \text{internal VE–to–VE}, \\ (\hat{p}_\alpha^{ij}, \rho_\alpha^{ij}, \lambda_\alpha(\hat{s}_\alpha^{ij}), z), & \text{VE–to–fine} \vee \text{vertical VE–to–VE}, \\ (\hat{P}_\alpha^{ij}, \rho_\alpha^{ij}, \Lambda_\alpha(\hat{s}_\alpha^{ij}), Z), & \text{horizontal VE–to–VE}. \end{cases} \tag{6.20}$$

Superscript ij refers to variables computed in a transition involving virtual cells, either cell i or j. Note that the coarse mobility Λ_α depends on the reconstructed saturation $\hat{s}_\alpha$ through integration of the fine-scale relative permeability $k_{r\alpha}$, which has an analytical expression in a sharp-interface setting. The coarse depth Z represents the top and bottom of the formation for gas phase and liquid phase, respectively. Reconstructed saturation in the virtual cell is based on an average value over the cell (see Section 6.5.1),

while reconstructed pressure is evaluated at the center of virtual fine cells and at the bottom of virtual VE cells, as illustrated in Figure 6.6.

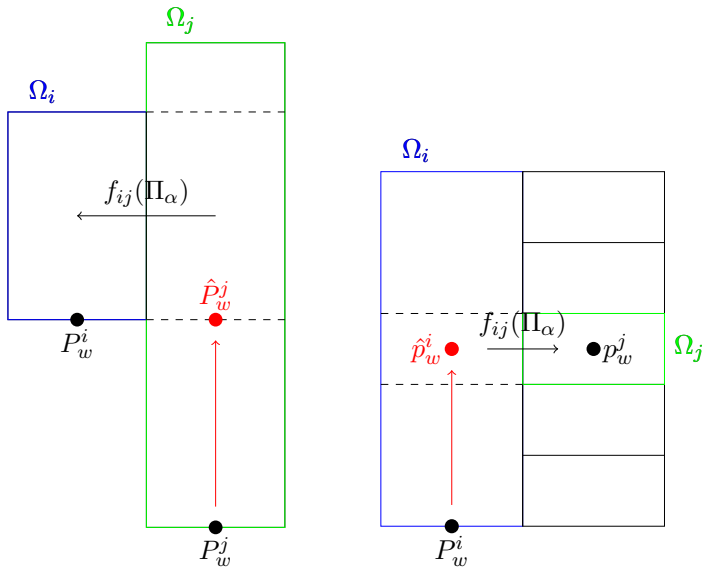

(a) VE-to-VE **(b)** VE-to-fine

Figure 6.6: Reconstruction of water pressure (denoted with hat) in a) a VE–to–VE horizontal transition zone and b) a VE–to–fine horizontal transition zone. Here, $f_{ij}(\Pi_\alpha)$ represents the flux function between cells i and j. Note how we distinguish between coarse variables P_w and fine-scale variables p_w.

6.7 Selection of discretization regions

Manually dividing the computational grid into discretization regions is a cumbersome process, especially for real reservoirs of complex geometry. Therefore, we propose an algorithm for automatic selection of discretization regions according to geometry and petrophysical properties.

First, we decompose the grid into separate columns according to the connectivity c_{ij} between cells. This is provided by the following sparse column-connectivity matrix that represents a graph partitioning,

$$(A_c)_{ij} = \begin{cases} 1, & c_{ij} \ \wedge \ C(i) = C(j), \\ 0, & \text{else}, \end{cases} \tag{6.21}$$

where c_{ij} is a boolean representing the connectivity between cells i and j, defined as

$$c_{ij} = \begin{cases} 1, & T_{ij} > T_\epsilon \ \wedge \ \Gamma_{ij} \notin \Gamma_{\mathcal{D}}, \\ 0, & \text{else}, \end{cases} \tag{6.22}$$

i.e., the cells are "connected" if the transmissibility T_{ij} at their common interface Γ_{ij} exceeds a prescribed value T_ϵ and if that interface is not part of a sealing barrier, $\Gamma_{\mathcal{D}}$. In

other words, all "non-connected" cells $c_{ij} = 0$ and horizontal connections $C(i) \neq C(j)$ are removed from the graph. We then select the disconnected components of the graph using a depth-first search from the Boost Graph Library [58], which is then used to assign cells into distinct column sections.

Next, we need to merge column sections that belong to the same discretization regions. The partitioning of different discretization regions is based on a cell connectivity matrix representing the coarse connections between cells,

$$(A_d)_{ij} = \begin{cases} 1, & M(i, D(j)) > 1, \\ 0, & \text{else}, \end{cases} \tag{6.23}$$

where $D(j)$ is the discretization region of cell j and $M(i, J)$ is the number of cells with discretization region J that cell i is connected to with a non-zero transmissibility. By merging all cells associated with the connected components of the graph, we obtain the discretization regions V of the grid.

As an example, if VE cell i is an internal VE column, its neighboring VE cells have the same discretization region J, implying that $M(i, J) > 1$ and $(A_d)_{ij} = 1$. Cell i then belongs to the same discretization regions as its neighbors so the columns are merged. The selection algorithm is described in more details in [44].

6.8 Solution strategy

Once the primary variables are defined for each cell (that being VE, fine, or virtual), and the appropriate expressions for the interface fluxes have been established, the system equations are ready to be solved. By using the S-formulation, the full problem can be solved monolithically as a set of nonlinear algebraic equations. A pseudo-code for the hybrid solution scheme is provided in Algorithm 2.

Algorithm 2: Hybrid sharp-interface model

Result: Solution vectors $\{\boldsymbol{p_w}, \boldsymbol{s_n}\}$ and $\{\boldsymbol{P_w}, \boldsymbol{S_n}\}$ of multiphase flow equations
initialize states for fine and coarse parts of domain: $\{p_w^0,\ s_n^0\}$;
vertically integrate in VE regions to get coarse state: $\{P_w^0,\ S_n^0\}$;
define state-independent upscaled properties;
$\Phi = \int_{\zeta_T}^{\zeta_B} \phi\, dz$;
$\boldsymbol{\kappa} = \int_{\zeta_T}^{\zeta_B} \boldsymbol{K}_{\parallel}\, dz$;
for $t = 0$ *to* T **do**
compute phase densities $\rho_\alpha = \rho_\alpha^{t+1}(x, y)$;
VE cells: calculate $h = h(S_n^{t+1}, S_{n,\max}^{t+1})$ and $h_{\max} = h_{\max}(S_{n,\max}^{t+1})$;
RVE cells: calculate $h = h(S_{n,\mathrm{mob}}^{t+1})$ and $h_{\max} = h_{\max}(h_B^{t+1}, h_T^{t+1})$;
virtual cells: compute $a_M = a_M(h)$ from (5.27), $a_R = a_R(h, h_{\max})$ from (6.14);
if *Horizontal VE–to–VE* **then**
reconstruct coarse phase saturations in virtual VE cells;
$\hat{S}_n = a_M(1 - s_{wr}) + a_R s_{nr}$;
$\hat{S}_w = 1 - \hat{S}_n$;
calculate water height Δz and CO2 plume height Δh below *bottom* B of virtual VE cell;
$\Delta h = \max(\tau + h - B,\ 0)$;
$\Delta z = \beta - B - \Delta h$;
reconstruct water pressure at bottom of virtual VE cells;
$\hat{P}_w = P_w^{t+1} - \rho_w g \Delta z - \rho_n g \Delta h$;
end
if *VE–to–Fine* $\vee$ *Vertical VE–to–VE* **then**
reconstruct phase saturations in virtual fine cells;
$\hat{s}_n = a_M(1 - s_{wr}) + a_R s_{nr}$;
$\hat{s}_w = 1 - \hat{s}_n$;
calculate water height Δz and CO2 plume height Δh below *center* C of virtual fine cell;
$\Delta h = \max(\tau + h - C,\ 0)$;
$\Delta z = \beta - C - \Delta h$;
reconstruct water pressure at center of virtual fine cells;
$\hat{p}_w = P_w^{t+1} - \rho_w g \Delta z - \rho_n g \Delta h$;
end
compute upscaled variables: P_c^{t+1}, $\boldsymbol{\Lambda}_\alpha^{t+1}$, Q_α^{t+1};
compute fine and coarse Darcy fluxes: $\boldsymbol{u}_\alpha^{t+1}$ and $\boldsymbol{U}_\alpha^{t+1}$;
solve Equation 2.11 and Equation 5.17 for states $\{p_w^{t+1},\ s_n^{t+1}\}$ and $\{P_w^{t+1},\ S_n^{t+1}\}$;
end
VE cells: reconstruct $\{\hat{\boldsymbol{p}}_w,\ \hat{\boldsymbol{s}}_n\}$ from Equation 6.17 and Equation 5.29;

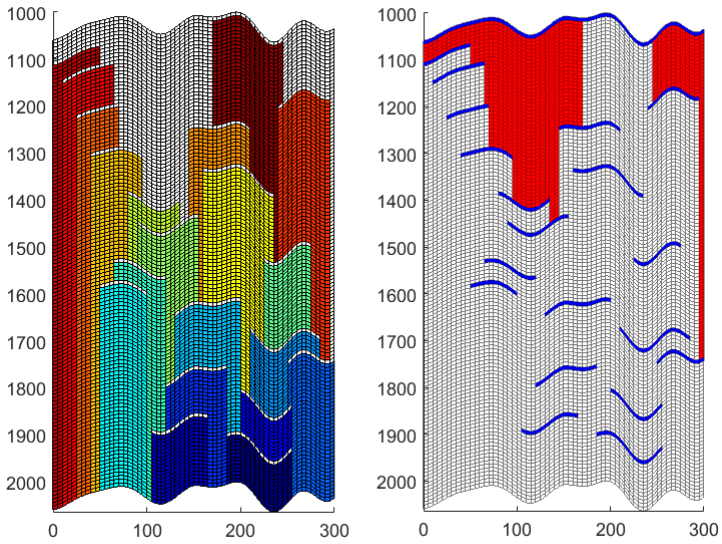

Figure 6.7: Left: Top surface subgrids (colored) under sealing barriers of a heterogeneous formation. The top boundary is open and is thus not assigned any top surface grid. Right: Remaining regions not part of a top surface subgrid, either discretized as full-dimensional cells (blue) or VE cells (red).

6.9 Trap analysis

Structural traps in a formation are located as described in Section 2.1.7. For an exposition of how a trap analysis is conducted for a conventional VE model in MRST, we refer the reader to [4] and [50]. There are more computations in a hybrid framework due to the presence of multiple semi-permeable barriers, but the calculations follow the same principle. That is, we extract the top surface of subgrids bounded above by sealing barriers, perform separate trap analysis for each of them, and then assemble the trapped masses to give a global trapping inventory. The top surfaces are mutually exclusive, facilitating a parallelized trap analysis. Extraction of top surface subgrids for a multilayered system is shown in Figure 6.7. In this thesis, *semi-permeable barriers* are treated to have sufficiently low permeability and high entry pressure to be considered as structural traps in the trapping inventory. In the following, we briefly go through how the six mass categories are calculated for a hybrid model, particularly the modifications for RVE columns.

Categorization of trapped CO2 in a VE column is found by separating the column into two parts, a part that is above the spill-point and considered structurally trapped and a part below the spill-point constituting the free plume. Let h and $h_{\max}$ be the height of the mobile and residual part of the plume, respectively, V_ϕ the pore volume and $z_t = \min\big[\max(z_{sp} - \tau,\, 0),\, H\big]$, where z_{sp} is the spill-point depth of the associated trap and H the height of the VE column. The residual structurally trapped height of the top plume is given by

$$h_{srt} = \max\big(\min(z_t, h_{\max}) - \min(z_t, h),\, 0\big), \tag{6.24}$$

and the residual non-structurally trapped height is

$$h_{rt} = \max\left(h_{\max} - \max(h, z_t),\ 0\right). \tag{6.25}$$

The mobile mass and residual mass of structurally trapped CO2 are given by, respectively

$$M_{sm} = V_\phi \rho_n \min(z_t, h)(1 - s_{wr} - s_{nr}), \tag{6.26}$$

$$M_{sr} = V_\phi \rho_n \big(\min(z_t, h) + h_{srt}\big) s_{nr}. \tag{6.27}$$

The mobile mass and residual mass of the free plume are given by, respectively,

$$M_{fm} = V_\phi \rho_n \min(h - z_t, 0)(1 - s_{wr} - s_{nr}), \tag{6.28}$$

$$M_{fr} = V_\phi \rho_n \min(h - z_t, 0) s_{nr}. \tag{6.29}$$

Lastly, residual mass left in the wake of the plume, non-structurally trapped, reads

$$M_r = V_\phi \rho_n h_{rt} s_{nr}. \tag{6.30}$$

For additional details of trapping for conventional VE columns, we refer to the description in [46].

Calculating volumes of CO2 residually trapped in the wake of the plume is less trivial for the fine regions of a hybrid model since the saturation can take any value between the endpoints $[s_{nr},\ 1 - s_{wr}]$. As the saturation decreases towards the residual value, the relative permeability drops to zero, implying that it will take an infinite amount of time to reach the residual value. Hence, the indicator $s_n \leq s_{nr}$ for residually trapped CO2 will never be satisfied, even though it is *effectively* residually trapped. To account for this, we apply a buffer ϵ_{res} to the residual value upon query, i.e., $s_n \leq \epsilon_{res} s_{nr}$. After experimenting, we found $\epsilon_{res} = 1.5$ to be a good estimate. That said, the residual buffer only determines the relative distribution of residual CO2 inside and outside the migrating plume. Since residual CO2 is trapped with high security, regardless of subcategory, the relative distribution is not of utter importance. We retain the distinction for completeness, though.

Generation of a trapping inventory for RVE columns (superscript rve) becomes more intricate, since we need to account for additional plumes emanating from sealing barriers. The additional plumes are assumed to be immobilized unless they have reached the buoyantly segregated plume under the sealing barrier above. Thus, the structurally trapped mobile masses and free masses are unchanged, i.e., $M_{sm}^{rve} = M_{sm}$, $M_{fm}^{rve} = M_{fm}$ and $M_{fr}^{rve} = M_{fr}$.

We refer to Figure 6.8 for visual support of the following derivation. To calculate structurally trapped volumes for RVE columns we need to assemble parts of each plume that are located above the spill-point depth z_{sp} of the column. Let H_i and h_i be the heights from τ to sealing barrier i and the top of the plume originating from sealing barrier i, respectively. The height of the structurally trapped part of the plume

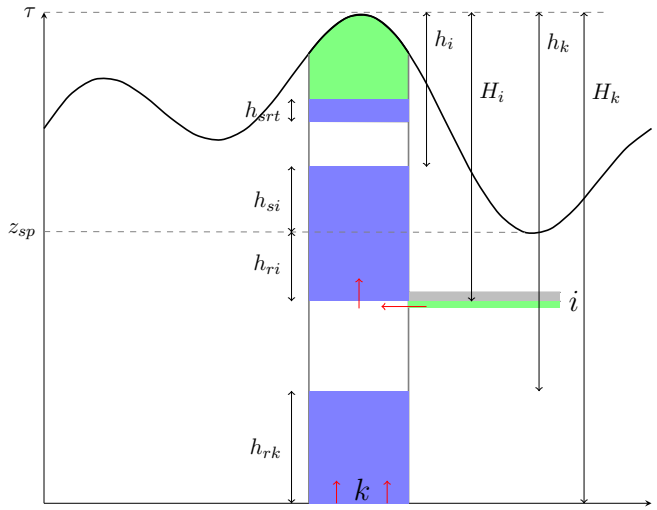

Figure 6.8: Illustrative example of an RVE column with height $H = \beta - \tau$ along with definitions of plume heights used to derive a hybrid trapping inventory. The left boundary is closed, explaining the depth of the spill-point z_{sp}. Mobile plume is shown in green, residual CO2 in blue. Sealing barrier i is responsible for a separate plume of height $h = h_{si} + h_{ri}$, and sealing barrier k causes diffuse leakage and a separate plume of height $h = h_{rk}$. Since the latter is at the bottom of the RVE column, we have $H_k = H$. Red arrows denote fluxes of CO2 from sealing barriers.

originating from sealing barrier i is then given by

$$h_{si} = z_t - \min(z_t,\, h_i) - \max(z_t - H_i,\, 0). \tag{6.31}$$

The net height of the residual parts of an RVE column that is structurally trapped is then given by a sum of the associated parts from the topmost plume and the remaining plumes,

$$h_{sr} = h_{srt} + \sum_{i \in \mathcal{D}_s} h_{si}, \tag{6.32}$$

where $\mathcal{D}_s$ is the set of sealing barriers connected to the RVE column. The net residual mass of structurally trapped CO2 in RVE columns then becomes

$$M_{sr}^{rve} = V_\phi \rho_n \big(h_{sr} + \min(z_t,\, h)\big) s_{nr}. \tag{6.33}$$

The height of the free part of the plume originating from sealing barrier i is given by

$$h_{ri} = \max(H_i,\, z_t) - \max(h_i,\, z_t), \tag{6.34}$$

and the net height of residual *non-structurally* trapped parts of plumes becomes,

$$h_r = h_{rt} + \sum_{i \in \mathcal{D}_s} h_{ri}. \tag{6.35}$$

This gives the net residual mass of non-structurally trapped CO2 left in the wake of migrating plumes,

$$M_r^{rve} = V_\phi \rho_n h_r s_{nr}. \tag{6.36}$$

Chapter 7

Results

In this chapter, we outline the discrete setup for a set of CO2 injection scenarios and present the results of the numerical simulations. We first consider a simple synthetic, rectilinear 2D formation, for which we demonstrate how to set up and simulate a hybrid model in MRST. The applicability of hybrid models is accentuated by simulating on a sloped 2D domain with a stochastic distribution of semi-permeable barriers. We conclude with an application to real data from the 11th SPE benchmark.

We remark that this chapter assumes the reader is familiar with how black-oil models are implemented in MRST's AD–OO framework. For an exposition, we refer the reader to tutorials in the MRST textbooks [38] and [39].

7.1 Two-dimensional rectilinear domain

As our first test case, we consider a highly simplified, rectilinear formation consisting of a single semi-permeable barrier interbedded in a high-permeable rock with permeability 100 mD, for which VE modeling is reasonable [7]. Hence, the problem is well suited for a hybrid model. The simulation grid is illustrated in Figure 7.1. There is a cut-off in the low-permeable barrier to mimic a complete fracture through which CO2 leaks immediately.

In the following, we provide a concise description of setting up a hybrid model and

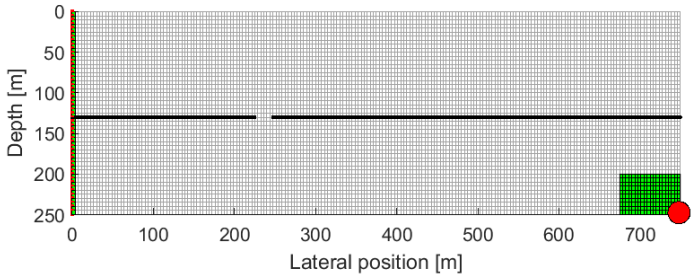

Figure 7.1: Synthetic two-dimensional horizontal grid used for simulations in Section 7.1. Imposed full-dimensional cells are marked in green. The horizontal sealing barrier is marked by thick, black faces. The left boundary, marked in red, is open. The other boundaries are closed. The red dot marks the location of the injection well.

running a simulation in MRST. Herein, we restrict ourselves to the specifics related to the hybrid model. For a thorough demonstration, we refer to [38].

The rectilinear grid is generated by calling `cartGrid(size, dim)` with desired dimensions and sizes. In a hybrid framework, fine cells part of a VE column must be agglomerated. This is performed internally in the top-level function `convertToHybridModel` that takes as input the fine-scale model and the set of faces and cells in the original full-dimensional grid that represent sealing barriers, and returns a hybrid representation of it. First, we construct a boolean vector `face_category` with zeros for all faces representing sealing barriers or other imposed fine regions. Together with the fine-scale grid `G`, we compute coarse discretization regions for each fine cell by

```
cell_category = findCategoriesMultiVE(G, face_category);
```

Cells are connected by a graph through `face_category`, and assigned a distinct category by running a depth-first search as detailed in Section 6.7. Based on the `cell_category` vector, we can construct a hybrid grid through

```
CG = generateCoarseGridMultiVE(G, cell_category);
```

The hybrid grid includes a partition vector that maps from fine to coarse cells, `CG.partition`, which is used to construct a coarse model based on the TPFA scheme.

```
model_coarse = upscaleModelTPFA(model, CG.partition);
```

The coarse model defines upscaled petrophysical properties that we extract by `rock_c = model_coarse.rock`. We now have a coarse grid and a coarse rock object that together with the fluid object form a complete hybrid model containing a description of the governing equations permissible in a mixed-dimensional setting.

```
model_hybrid = WaterGasHybridModel(CG, rock_c, fluid);
```

The new model, `WaterGasHybridModel`, is a modified version of MRST's original `TwoPhaseWaterGasModel` for a two-phase CO2–brine system that adds operators for calculating heights of virtual cells, functionality for handling sealing barriers and region support for capillary pressure. The expression for phase fluxes depends on the discretization of adjacent cells. Therefore, we change the setup of the governing equations, `equationsWaterGas`, to a modified version `equationsWaterGasHybrid` accommodating a hybrid discretization. Fluxes over interfaces of any transition zone are compactly computed by the helper function `computeHybridFluxesRVE`, which also computes heights $\boldsymbol{h_T}$ and $\boldsymbol{h_B}$ for relaxed VE columns.

Next, we create hybrid representatives of a schedule and an initial state. Boundary conditions in the fine model are mapped to the coarse model using `CG.partition`, and

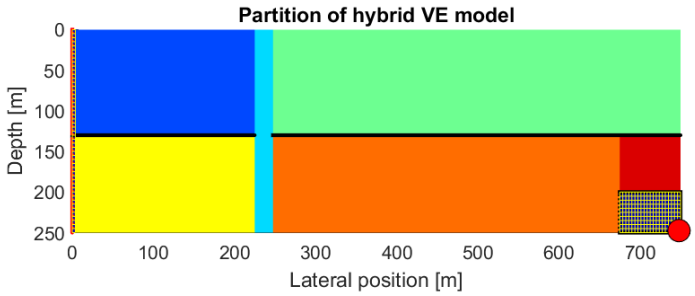

Figure 7.2: Partitioning of the synthetic two-dimensional horizontal domain into different discretization regions. VE cells are marked by pure colors. Full-dimensional cells are black and bounded by yellow edges.

the upscaled schedule is computed by the function `upscaleSchedule`. The initial state is converted to a hybrid permissible counterpart by upscaling the state variables for VE cells. The primary variables, coarse saturation and coarse pressure, are computed as pore-weighted averages of associated fine cells. The hybrid initial state is computed by the function `upscaleState`.

Finally, the hybrid physical model, upscaled state and upscaled schedule can be packed into a simulation object `simulateScheduleAD` that runs a complete simulation. It supports optional arguments for outputting information about time steps, convergence rates, reservoir states, etc. [38]. After simulation, fine-scale variables for each time step are reconstructed by calling the function `convertHybridStates`.

The MATLAB codes for the hybrid framework are provided in the folders `models` and `discretization` of the GitHub repository https://github.com/vetlenev/Master-Thesis, and are projected to synthesize with the `hybrid-ve` module of MRST.

The hybrid partitioning is shown in Figure 7.2. Notice that full-dimensional cells are imposed near the well and at the open left boundary. The latter is imposed to ensure the hydrostatic pressure distribution is fulfilled on the fine scale.

Two sets of simulations are performed. In the first one, we set a transmissibility multiplier of 10^{-8} for the semi-permeable barrier to approximate it as fully sealing. CO2 is injected over 40 years with subsequent 360 years of migration. The solutions after injection and at end of simulation are shown in Figure 7.3.

The thickness and shape of the plume under the sealing barrier correspond well between the full-dimensional and the hybrid models. We note that residual saturation left behind in the vertical passage through the chimney, modeled as a VE–to–VE horizontal transition region, is captured by the hybrid model. However, the horizontal extent of the residual part is wider in the full-dimensional model than in the hybrid model. This is a direct consequence of the RVE assumption for these cells. CO2 exiting through the open chimney is assumed to migrate strictly vertically and segregate with the resident brine. In other words, the horizontal component of the flux through the chimney is missed in the hybrid model so that the residual part is only one column wide.

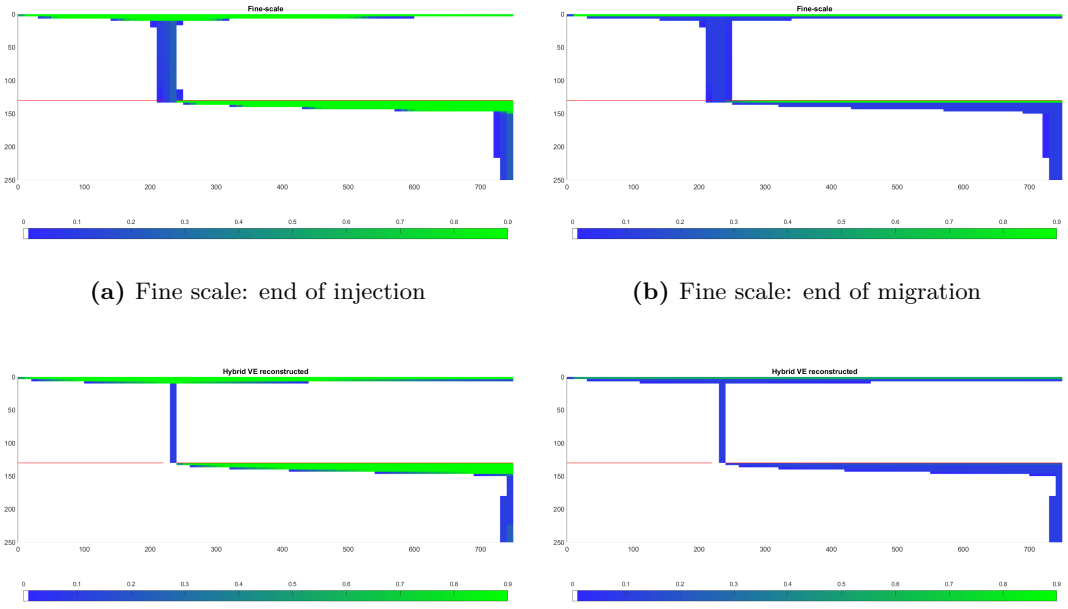

(a) Fine scale: end of injection

(b) Fine scale: end of migration

(c) Hybrid: end of injection

(d) Hybrid: end of migration

Figure 7.3: Comparison of CO2 saturation for a full-dimensional and hybrid model for a 2D domain with a single horizontal semi-permeable barrier (marked in red) embedded in a high-permeable region. White indicates pure brine, blue residual CO2 (with residual saturation $s_{nr} = 0.1$) and green max possible CO2 saturation ($1 - s_{wr} = 0.9$).

In the late-migration stage, there are notable differences in CO2 saturation immediately below the sealing barriers. In the fine-scale model, the saturation is approximately at the maximum value $s_n = 1 - s_{wr}$, while it is significantly lower in the hybrid model. The layer of CO2 becomes thin enough to only be represented by a single layer of cells in the vertical. The fact that a full-dimensional model is limited by the vertical resolution, the calculation of the depth of the sharp interface is vulnerable to errors, especially for horizontal sealing barriers [40]. For VE columns in the hybrid model though, phase fluxes at the sharp interface are accurately resolved, as discussed in Chapter 5.

To demonstrate the effect of vertical resolution, we conduct three simulations of increasing resolution. The relative permeability function is linear to accentuate the differences caused by grid resolution. Saturation distributions at the end of simulation are shown in Figure 7.4, and exited volumes of CO2 as a function of time is shown in Figure 7.5. During injection, there is hardly any difference between the models. After injection stops, particularly beyond time step 250, the models diverge and eventually flatten. The discrepancy is due to the limited vertical resolution failing to capture the true depth of the sharp interface. For the vertical resolutions shown in Figure 7.4, the sharp interface is less than one cell thick and its true height is not accurately resolved. By increasing the vertical resolution, the accuracy increases and the full-dimensional solution converges to that predicted by the analytical expression used for the hybrid model.

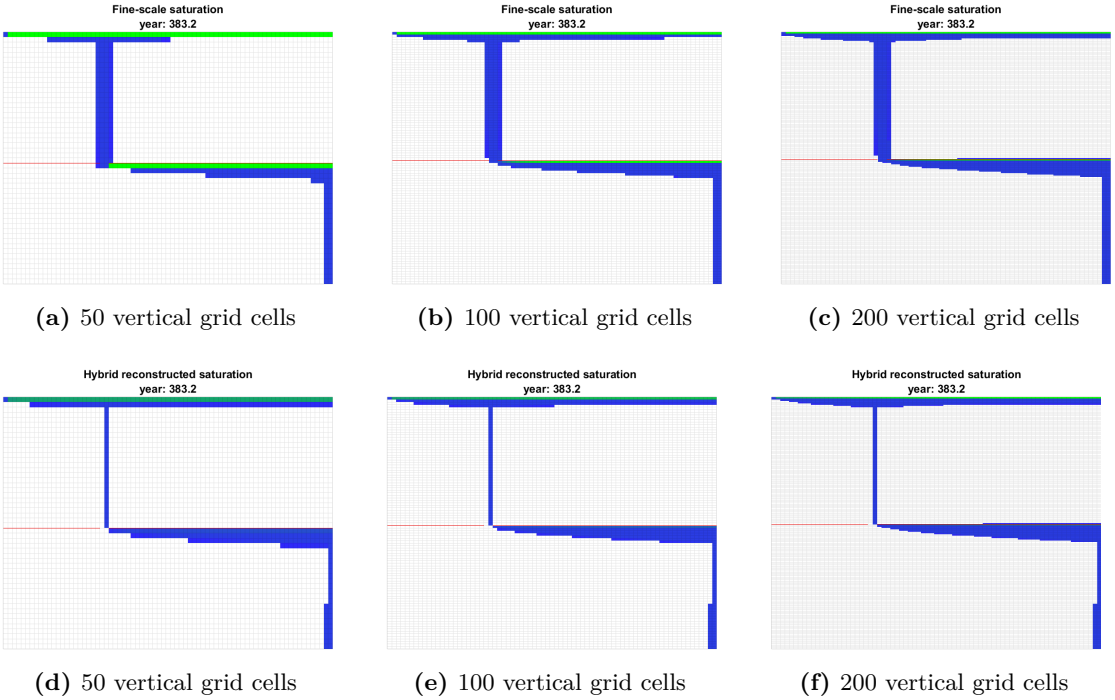

(a) 50 vertical grid cells **(b)** 100 vertical grid cells **(c)** 200 vertical grid cells

(d) 50 vertical grid cells **(e)** 100 vertical grid cells **(f)** 200 vertical grid cells

Figure 7.4: CO2 saturation after migration for three different vertical resolutions, illustrating how the vertical saturation in VE columns of a full-dimensional model converges to the solution for the hybrid model as the vertical resolution is refined. Top row: full-dimensional model. Bottom row: hybrid model.

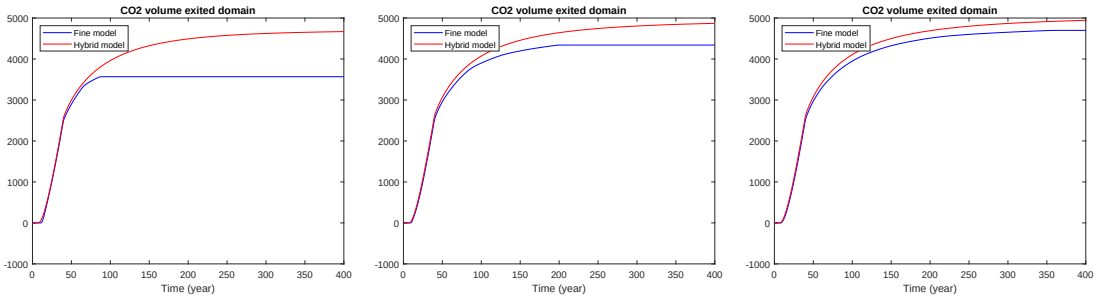

(a) 50 vertical grid cells **(b)** 100 vertical grid cells **(c)** 200 vertical grid cells

Figure 7.5: Comparison of exited volume (in m^3) for three different vertical resolutions for the 2D horizontal test case with a single horizontal sealing barrier.

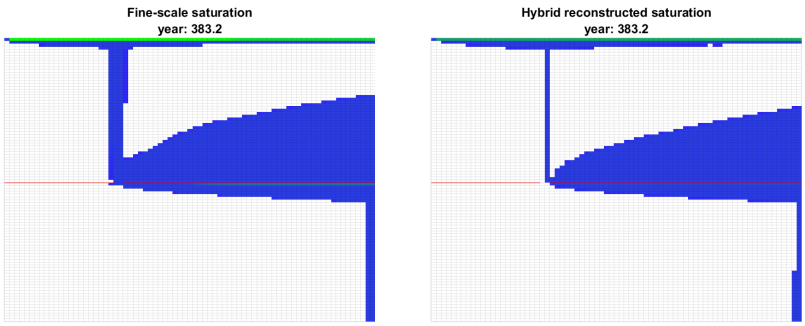

(a) Full-dimensional **(b)** Hybrid

Figure 7.6: Comparison of full-dimensional and hybrid model for a 2D rectilinear domain with a single horizontal semi-permeable barrier represented as a face with transmissibility multiplier 10^{-5}. Hundred cells are used in the vertical.

The second setting applies a transmissibility multiplier of 10^{-5} to facilitate diffuse leakage through the semi-permeable barrier. The saturation distribution at the end of simulation is shown in Figure 7.6. Clearly, the RVE assumption of the hybrid model is able to capture the slow vertical drainage through the semi-permeable barrier, but not perfectly; the immobilized plume is more widespread in the full-dimensional model.

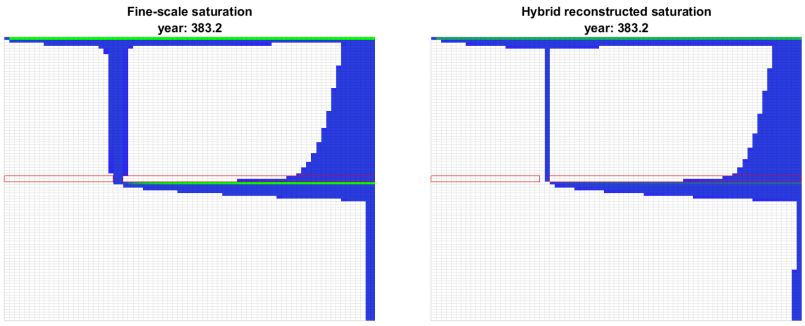

(a) Full-dimensional **(b)** Hybrid

Figure 7.7: Comparison of full-dimensional and hybrid model for a 2D rectilinear domain with a single horizontal semi-permeable barrier represented as full-dimensional cells. Hundred cells are used in the vertical.

A similar set of simulations are performed where the semi-permeable barrier is represented as full-dimensional cells instead of faces, as detailed in Section 6.1, and is shown in Figure 7.7. Diffuse leakage is now attributed to a capillary entry pressure, set to $5 \cdot 10^4$ Pa, rather than a transmissibility multiplier. Overall, both the mobile plume and the residual part left in the wake coincide well. In contrast to a transmissibility multiplier, using full-dimensional cells with capillary exclusion, CO2 only penetrates through the right part of the semi-permeable barrier. This is as expected since CO2 is not able to penetrate the sealing barrier before the entry pressure is exceeded. By applying a transmissibility multiplier instead, CO2 will in principle penetrate immediately upon contact with the barrier. This emphasizes the inaccuracies of modeling sealing barriers with tranmissibility multipliers.

7.2 Two-dimensional curvilinear domain

Next, we generate a 2D sloped reservoir to add complexity. To this end, we set up a heterogeneous formation with multiple low-permeable barriers of different thickness. We demonstrate the appropriateness of modeling such regions as relaxed VE columns.

The hybrid partitioning of the domain is provided in Figure 7.8. We use an adaptive grid generation that represents low-permeable barriers/lenses either as full-dimensional cells or lower-dimensional faces, depending on their thickness. The left boundary is open, and remaining boundaries are sealed. CO2 is injected over 40 years with subsequent 360 years of migration.

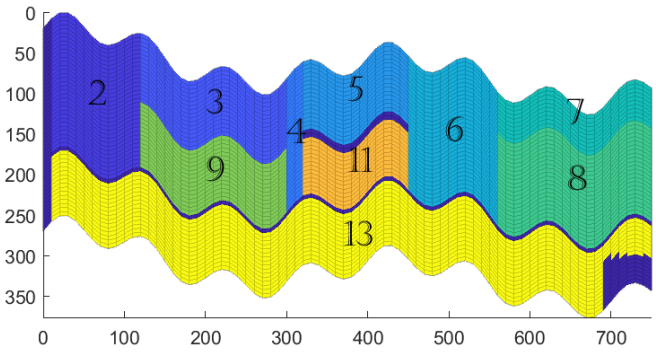

Figure 7.8: Partitioning of the synthetic two-dimensional sloped domain into different discretization regions marked by numbers. The dark blue regions are imposed fine cells. Spatial dimensions are given in units of meter.

Figure 7.9 reports the saturation distribution after injection and at the end of migration. The solution reveals a network of plumes that penetrate through semi-permeable barriers and also divert around them. A conventional VE model is unable to predict accumulation below semi-permeable interbeddings, as opposed to the hybrid model which is able to capture the most prominent plume paths. The deviation to the fine solution is primarily caused by two factors. First, the naive treatment of fluxes in transition regions to semi-permeable barriers incurs errors when tracing the history of migrating plumes. Second, a quadratic relative permeability combined with a limited vertical resolution is responsible for slower plume migration. This is evident in the plots as we observe the mobile plumes have migrated further upslope the semi-permeable barriers in the hybrid model than in the full-dimensional model.

The trapping inventory for the fine and hybrid models are shown in Figure 7.10, including the normalized absolute difference. The maximum difference, occurring 60 years after injection stops, is $6.68 \cdot 10^{-4}$ MT (megatonn), which is about 16% of total CO2 mass. Most of the mass categories are in great agreement with the fine model, where the free plume and exited volumes are exceptions. Deviations in exited volumes are primarily due to a nonlinear relative permeability retarding the speed at the sharp interface, as well as the RVE assumption slightly overestimating the speed of vertical plume migration.

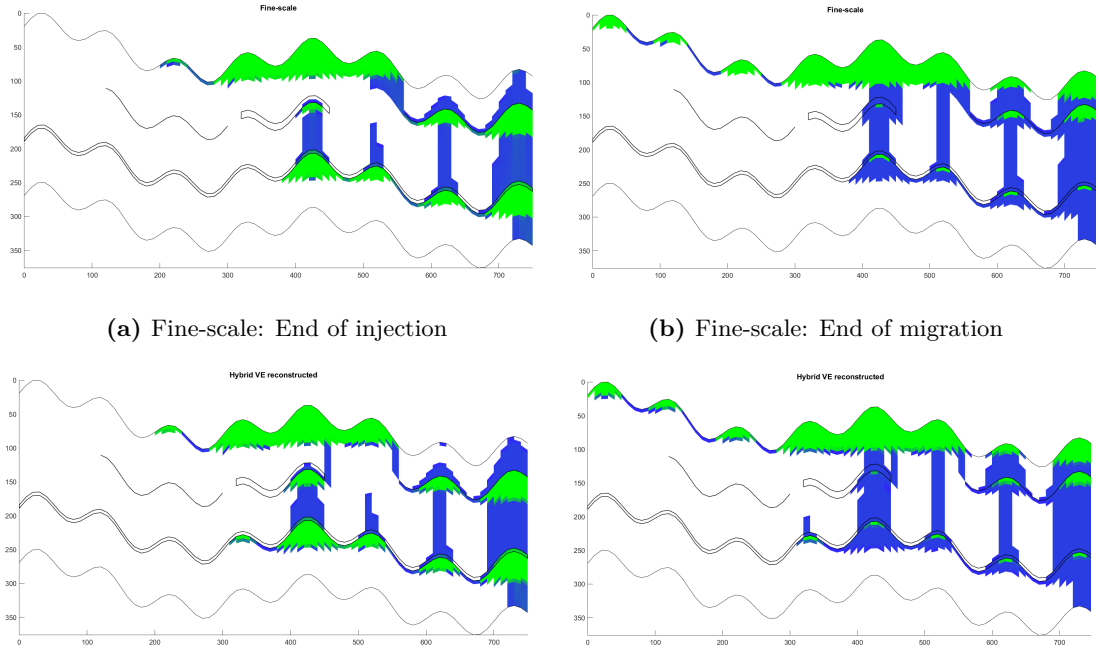

(a) Fine-scale: End of injection

(b) Fine-scale: End of migration

(c) Hybrid: End of injection

(d) Hybrid: End of migration

Figure 7.9: Comparison of CO2 saturation at end of injection and end of migration for a synthetic sloped, 2D domain with multiple sealing barriers. Quadratic relative permeability and 50 cells in the vertical direction are used. Sealing barriers are indicated by black faces.

To accentuate how the nonlinearity of the relative permeability affects upslope migration speed, we run a similar set of simulations but with linear fine-scale relative permeability that equals the upscaled relative permeability. With these settings, the maximum absolute difference in CO2 trapping is reduced from $6.68{\cdot}10^{-4}$ MT to $6.14{\cdot}10^{-4}$ MT. The top row of Figure 7.11 shows the corresponding CO2 saturation at the end of the simulation. The upslope migration now corresponds better between the models, though vertical resolution and the RVE assumption still incur some deviations in CO2 saturation.

Rerunning the simulations with linear relative permeability and twice as fine vertical resolution, the lateral migration upslope sealing barriers for the hybrid model is almost indistinguishable from the fine model, as shown in the bottom row of Figure 7.11. The cell in the full-dimensional model at which the sharp interface cuts through effectively represents a capillary fringe, whose spatial extent is dictated by the vertical resolution. Thus, a coarser resolution enhances the retardation effect, and we observe a greater deviation between the models. By increasing the vertical resolution, the lateral fine fluxes approximate the "correct" fluxes computed from the VE assumption [62]. The maximum absolute difference in CO2 trapping is further reduced to $4.43{\cdot}10^{-4}$ MT, which is about two-thirds of the maximum difference using quadratic relative permeability and half the vertical resolution, and 10% of the total mass of CO2. Deviations in exited volume are now mainly caused by errors in the RVE assumption.

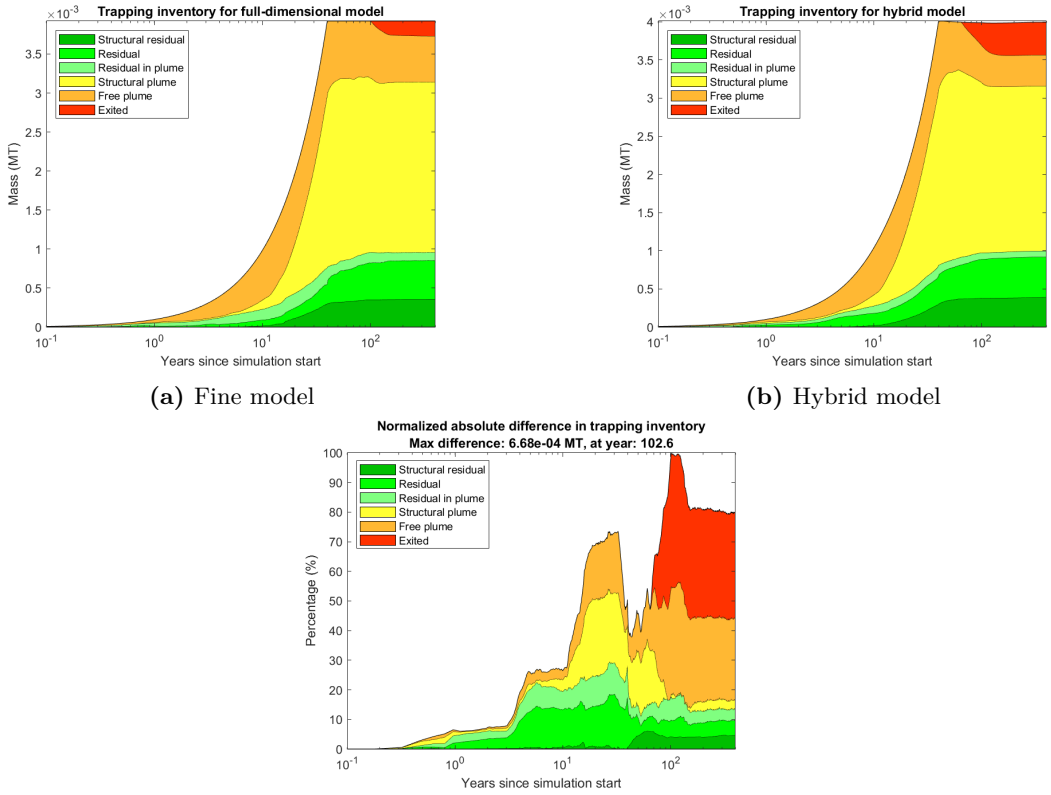

(a) Fine model

(b) Hybrid model

(c) Trapping difference

Figure 7.10: Trapping inventories for a synthetic sloped, multilayered 2D domain with quadratic relative permeability and 50 grid cells in the vertical direction.

7.3 Two-dimensional thick, multilayered system

Here, we present a vertically thick, heterogeneous reservoir with a pseudo-random distribution of sealing barriers that do not extend the full width of the reservoir. The geometry is motivated by vertical stacks of sealing barriers typical for formations on the Norwegian Continental Shelf, e.g., the Utsira formation [2]. Figure 7.12 shows the partitioning of discretization regions for a realization of 20 sealing barriers. The reservoir is located at a depth between 1000 and 2000 meters to merely simulate CO2 in a supercritical state. The curvatures of the formations are exaggerated to promote a trapping inventory where all mass categories are present. Entry pressure for sealing barriers and the high-permeable background are set to $7 \cdot 10^4$ Pa and $5 \cdot 10^3$ Pa, respectively. Brine and CO2 have linear relative permeabilities. No-flux boundary conditions are imposed on all sides except the top which is open with a brine pressure of $10^7 + \rho_w g z$ Pa.

Figure 7.13 shows CO2 saturation as injection stops (left column) and at the end of the simulation (right column). Semi-permeable barriers divert the main plume towards the left of the domain where it rises and eventually exits out the top boundary. This migration path, with its history traced by the residual content, is resolved by the hybrid model. The major difference is the *lateral* extent of the plume, which is wider in the full-dimensional model. Due to the RVE assumption in the hybrid model, more CO2 exits

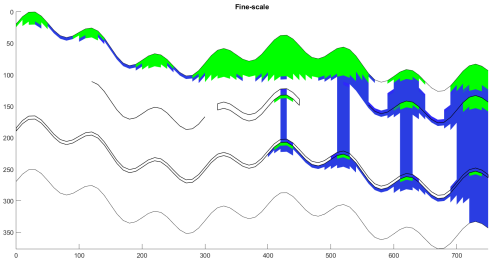

(a) Full-dimensional. 50 cells in vertical.

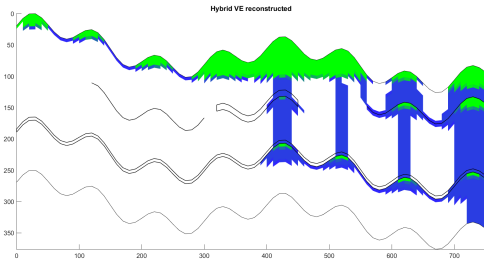

(b) Hybrid. 50 cells in vertical.

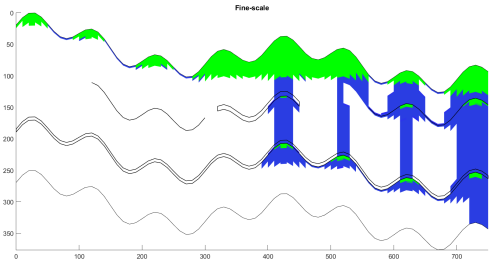

(c) Full-dimensional. 100 cells in vertical.

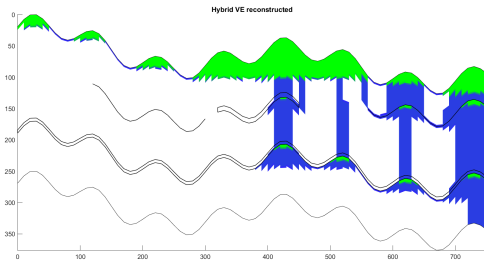

(d) Hybrid. 100 cells in vertical.

Figure 7.11: CO2 saturation at end of simulation for a synthetic sloped, multilayered 2D domain with linear relative permeability. Top row: coarser vertical resolution (50 cells). Bottom row: finer vertical resolution (100 cells).

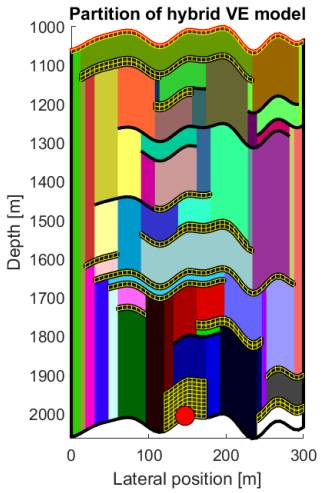

Figure 7.12: Discretization regions for a 2D multilayered formation with a stochastic distribution of sealing barriers with permeability 0.1 mD. The background region has permeability 100 mD. Fine cells are bounded by yellow edges.

the domain, as evident in the trapping inventories in Figure 7.14. The larger residual content in the full-dimensional model is attributed to our observations of increased lateral spread of CO2. Twenty years after injection stops, the maximum difference in trapping is reached, with exited volumes and the structural plume showing the largest deviations. Thus, despite the long-term structural trapping being in excellent agreement between the models, there are significant errors during migration *before* equilibrium is approached.

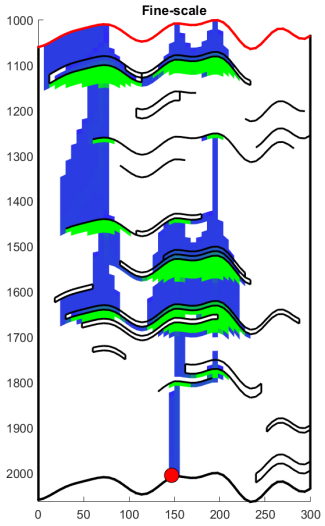

(a) Fine: End of injection

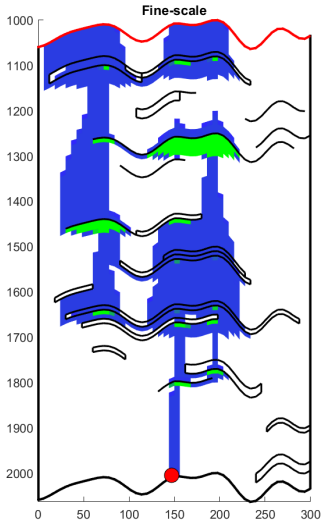

(b) Fine: End of migration

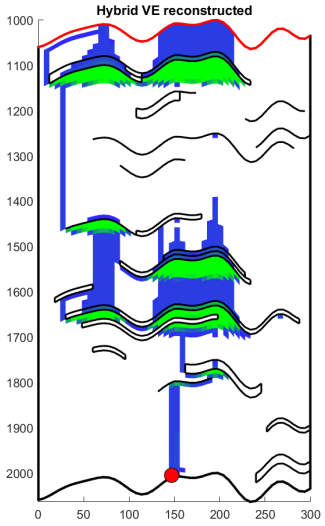

(c) Hybrid: End of injection

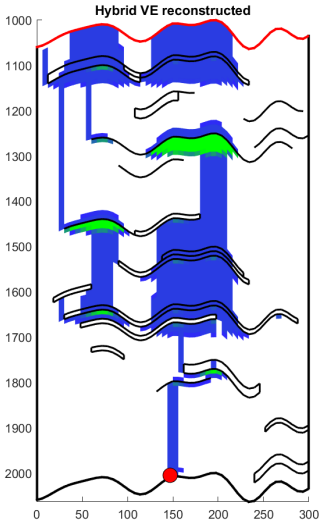

(d) Hybrid: End of migration

Figure 7.13: CO2 saturation at end of injection and end of migration for a tall, multilayered 2D domain with linear relative permeability. Semi-permeable barriers are highlighted in black.

The depths of the tip and average mobile plume are shown in Figure 7.15, where the average depth of the plume is calculated as saturation-weighted average for all cells where $s_n > s_{nr}$. The plume is only traced up to the top of the formation (1000 meters), because the fate of CO2 that leaves the domain is unknown. Just after a couple of months, both the tip and average mobile plumes have migrated further vertically for the hybrid model than for the fine model. This is as expected by the VE assumption, where

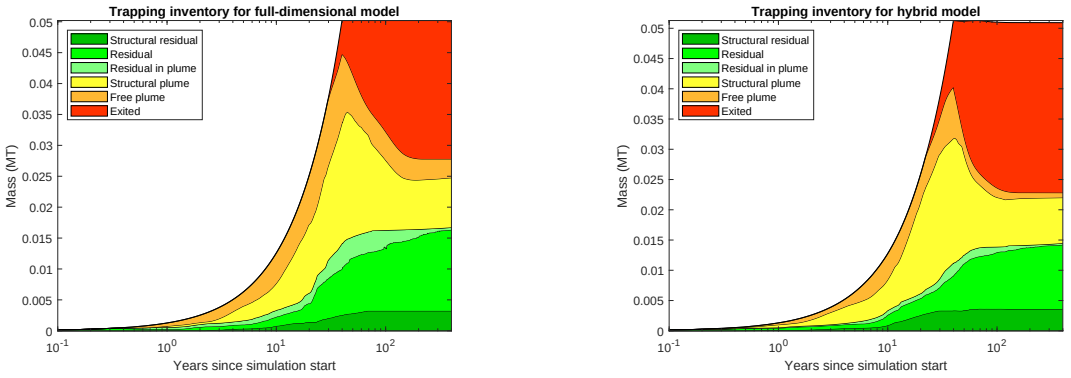

(a) Fine model

(b) Hybrid model

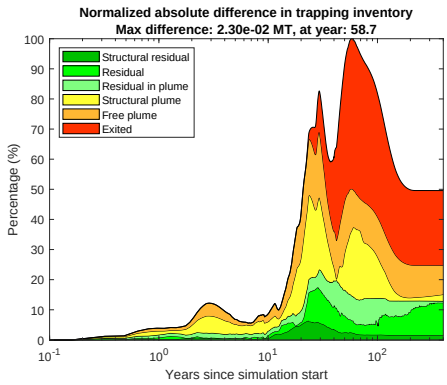

(c) Absolute difference

Figure 7.14: Trapping inventories for a synthetic thick, multilayered 2D domain with linear relative permeability.

CO2 and brine segregate instantaneously.

The depth of the tip of the plume is characterized by a stair-stepped pattern where each step represents immediate accumulation below an upcoming sealing barrier with subsequent leakage. This explains why the tip of the mobile plume exits the domain at an earlier time step for the hybrid model, i.e., it overestimates the speed of the tip. The functional form of the tip generally complies well, apart from a delay in the full-dimensional model, indicating that the migration paths of the plumes coincide.

The average depth of the plume is generally in good agreement between the models throughout the simulation, although the hybrid average is slightly and consistently ahead of the fine-scale average. This sustains until the tip of the plumes have reached the top of the formation. The calculations only account for plumes *inside* the domain, because the fate of CO2 that exits the domain is unknown. We have opted to omit this part due to the uncertainty it imposes. This explains why the average depth is not monotonically decreasing.

The same setup is simulated again but now with highly nonlinear relative permeability, i.e., using a Brooks-Corey exponent $n = 2.5$. The saturation distribution after injection ceases is shown in Figure 7.16. Compared with Figure 7.13, we observe a slight delay of the plume in the full-dimensional model when using a nonlinear relative

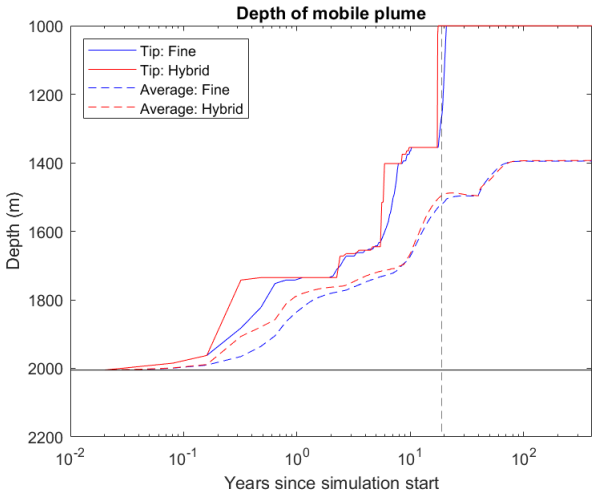

Figure 7.15: Depth of tip and average mobile plume over time for a 2D synthetic, thick formation with linear relative permeability. The horizontal black line represents the depth of the injection well. The dashed vertical line shows the point in time when the tip of the plume has reached the top of the formation. The average depth is only the global average up until this point.

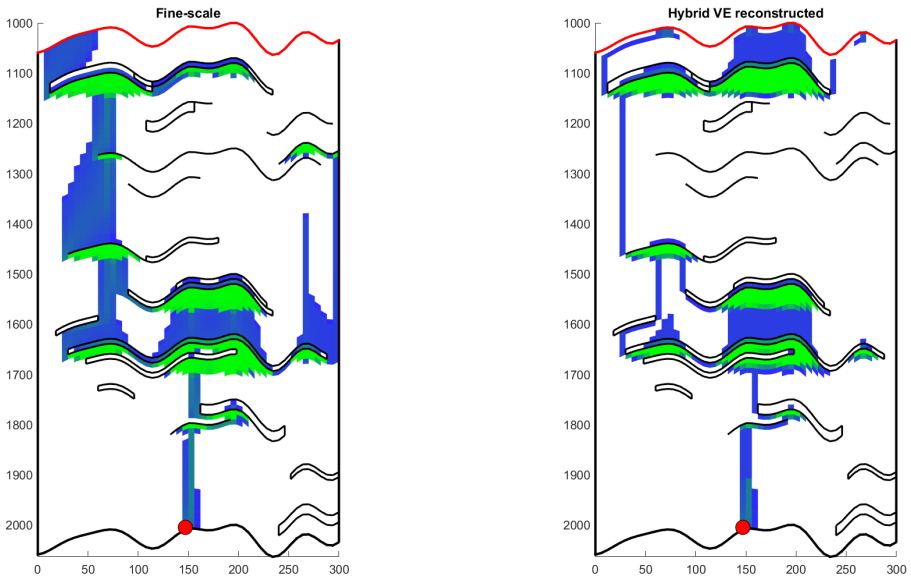

(a) Fine model **(b)** Hybrid model

Figure 7.16: CO2 saturation at end of injection for a thick, multilayered 2D domain with nonlinear relative permeability (Brooks-Corey exponent $n = 2.5$).

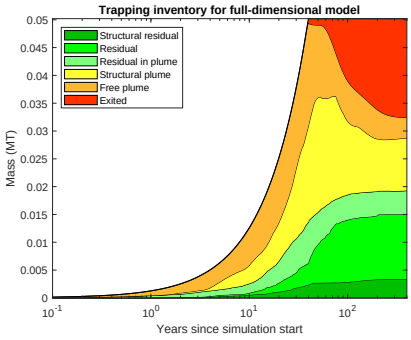

(a) Fine model

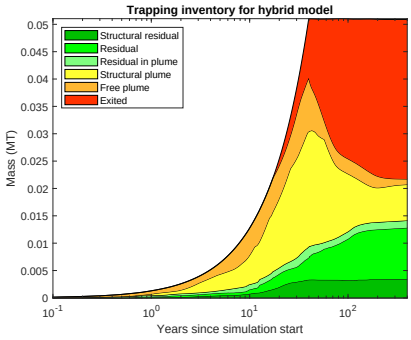

(b) Hybrid model

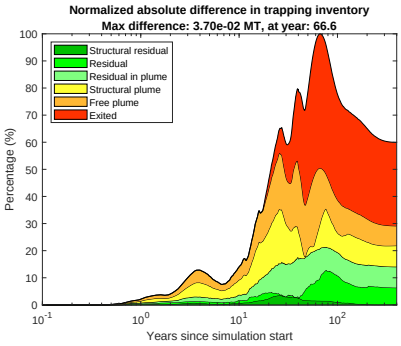

(c) Absolute difference

Figure 7.17: Trapping inventories for a synthetic thick, multilayered 2D domain with nonlinear relative permeability.

permeability, attributable to the reduced mobility. The nonlinear effect delays the spilling at the mid-left part of the domain sufficiently to cause increased accumulation and subsequent spilling at the mid-right structural trap, which is not resolved in the hybrid model. Overall, this results in an even larger deviation between CO2 volumes that have left the domain, as evident in the trapping inventory shown in Figure 7.17. The additional trap filled in the right part of the domain also gives a larger deviation in structurally trapped CO2. The maximum absolute difference in trapping for simulations with quadratic relative permeability has increased with 60% compared to simulations with linear relative permeability, indicating that the accuracy of our sharp-interface hybrid model is sensitive to the fine-scale relative permeability. There is a larger gap in time for when the tip of the plume leaves the model domain, as shown in Figure 7.18, yet again due to the reduced mobility.

Computational runtimes for a stochastic configuration of 20 layers are provided in Table 7.1. The hybrid model is not orders of magnitude faster than the full-dimensional model, but the median time of time steps is notably shorter for the hybrid model. This indicates that the hybrid model suffers from convergence issues for a few time steps. Figure 7.19a and Figure 7.19b show the cumulative runtime and the cumulative number of iterations, respectively. They reveal that there are a few time steps where the hybrid model spends a lot of iterations and runs slower. Simulation time for the full-dimensional model appears to be more consistent, but is longer on average. Another remarkable

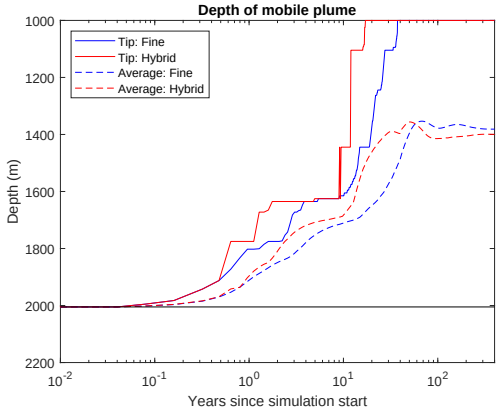

Figure 7.18: Depth of tip and average mobile plume over time for a 2D synthetic, thick formation with nonlinear relative permeability.

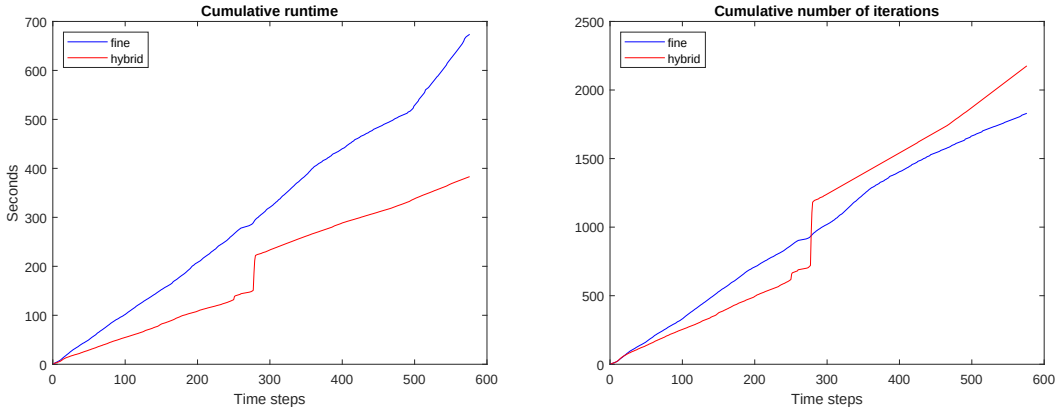

(a) Runtime **(b)** Iterations

Figure 7.19: Cumulative runtime and cumulative number of iterations for a pseudo-random distribution of twenty sealing barriers, and Brooks-Corey exponent $n = 2.5$ for relative permeability.

Model	n	Total time (s)	Median time (s)
Fine	1	1056	1.65
	2.5	674	1.06
Hybrid	1	520	0.44
	2.5	383	0.52

Table 7.1: Simulation times for fine-scale and hybrid models of a synthetic, tall domain using relative permeabilities with different Brooks-Corey exponent n, applying for both phases.

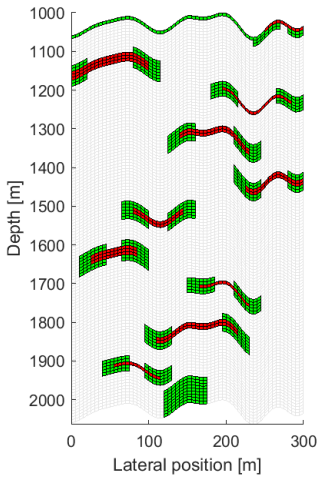

Figure 7.20: Discretization regions for a formation with ten sealing barriers. Red cells represent the sealing barriers, and green cells represent additional imposed fine cells.

observation is that the simulations with nonlinear relative permeablity run faster than those with linear relative permeability. In general, one would expect the opposite since the nonlinearity is computationally more demanding to handle.

A compromise that mitigates the errors from the RVE assumption at the expense of more computations is to impose additional fine cells in the domain to resolve more of the lateral spreading. More specifically, we include additional fine cells around the edges of sealing barriers. The discretization of a domain with ten sealing barriers is shown in Figure 7.20. Figure 7.21 compares the solution of a hybrid model, with and without added fine cells, with a full-dimensional solution. The hybrid solution without added cells clearly deviates from the full-dimensional solution, where the errors are primarily incurred by the RVE assumption. In contrast, the hybrid solution with added fine cells manages to trace a considerable amount of lateral spreading beyond sealing barriers, which qualitatively gives a more accurate solution.

7.4 Stochastic configurations of heterogeneous formation

To further investigate the sensitivity of our hybrid model on formation heterogeneity, we run stochastic configurations of the same domain as in Section 7.3 with sealing barriers of different permeability. The barriers have approximately uniform lengths and are well dispersed to facilitate statistically significant results. The great dispersion of sealing barriers will challenge our implementation of residual saturation and capillary exclusion.

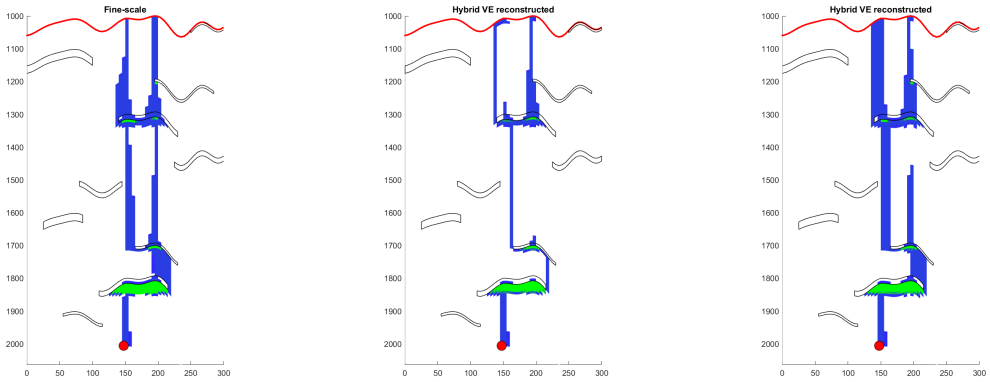

(a) Full-dimensional **(b)** Hybrid, without added cells **(c)** Hybrid, with added cells

Figure 7.21: CO2 saturation at end of simulation for a formation with ten sealing barriers. Two hybrid models are compared with a full-dimensional model. One of them uses the original hybrid framework (b), while the other includes additional fine cells around the tip of sealing barriers (c).

We will perform a qualitatively and quantitatively cross-simulation comparison including runtime, trap categorization, and depth of mobile plume. Since resolving the trailing edge of the plume in the full-dimensional model is not trivial, as explained in Section 6.9, we merge the categories *Residual* and *Residual in plume* to make the analysis more concise. In the following, we perform twenty stochastic simulations each for 20, 40, and 60 low-permeable barriers. The permeability field for a stochastic distribution of 60 barriers is shown in Figure 7.22a and the discretization regions in Figure 7.22b. There are two types of sealing barriers with permeability 0.01 and 0.1 mD and entry pressure $1.5 \cdot 10^5$ and $6 \cdot 10^4$ Pa, respectively. Injection stops after 20 years and simulation ends after 200 years.

Figure 7.23 shows the tip depth of the CO2 plume averaged across stochastic simulations with 20, 40, and 60 semi-permeable barriers. The tip of the plume in the hybrid model is consistently ahead of the tip in the fine-scale model, as anticipated from the VE assumption. We also notice how the tip of the plume delays as the number of sealing barriers increases due to more migration occurring laterally. Considering the time when the tip of the plume reaches the top of the formation (1000 meters depth), there is a larger deviation between simulations with 20 and 40 semi-permeable barriers than with 40 and 60 barriers. The former has approximately a four times larger difference in time than the latter, indicating a nonlinear dependency with the number of sealing barriers. The average absolute difference in plume depth accumulated up to current time step k between the fine-scale and hybrid model (hereby abbreviated $AADA$) is given by

$$AADA = \sum_{i=1}^{k} |(d_i)_h - (d_i)_f| \oslash \sum_{i=1}^{k} i, \tag{7.1}$$

where $\oslash$ represents elementwise division and $(d_i)_h$ and $(d_i)_f$ are the depths of the plume tip at time step i for the hybrid and full-dimensional model, respectively. The $AADA$ is shown by green curves in Figure 7.23, and quantifies in a visually pleasing way how

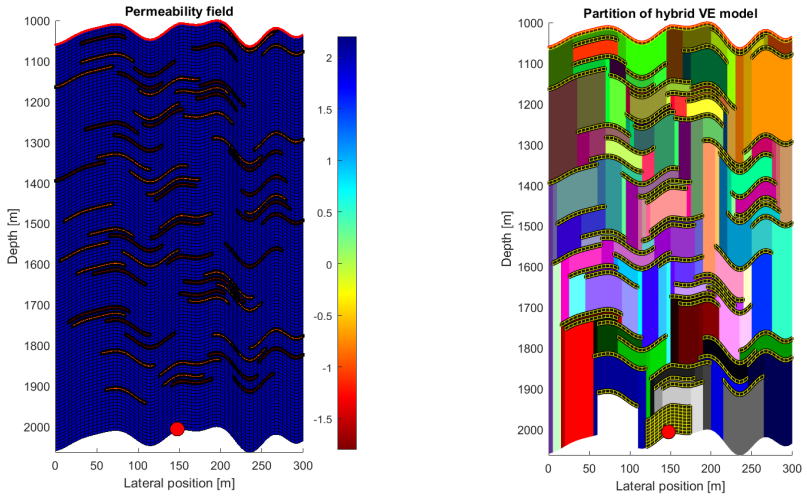

(a) Permeability field. **(b)** Discretization regions.

Figure 7.22: Stochastic distribution of 60 low-permeable barriers in a heterogeneous formation. a) Permeability field. Colorcodes represent the common logarithm of permeability, which is given in units of milli Darcy (mD). b) Partitioning of domain into discretization regions.

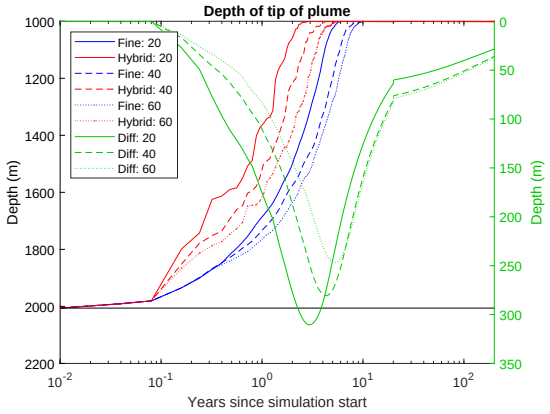

Figure 7.23: Depth of tip of plume for full-dimensional and hybrid model from simulations with a stochastic distribution of different number of semi-permeable barriers. Graphs with label *Diff* show average of absolute difference accumulated up to current time step ($AADA$).

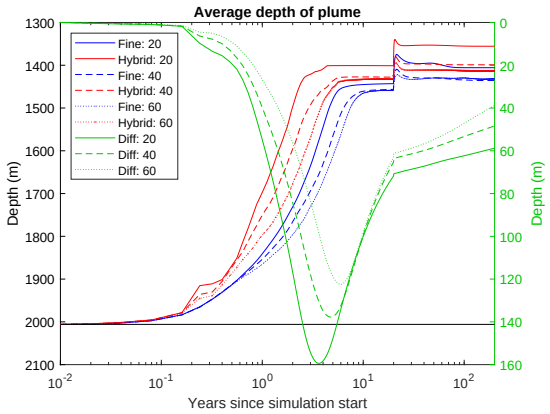

Figure 7.24: Average depth of plume for full-dimensional and hybrid model for stochastic distribution of different number of semi-permeable barriers.

Layers	Residual		Mobile		Exited		Net	
	Mean	St. dev.	Mean	St. dev.	Mean	St. dev.	Mean	St. dev.
20	48.1	10.6	75.3	51.4	17.6	4.5	27.7	5.5
40	56.8	6.9	44.0	24.0	27.9	6.5	38.8	5.4
60	50.1	8.4	52.9	26.6	30.6	7.3	40.5	6.4

Table 7.2: Scaled $L1$ error in percentage for three mass categories, including net masses, for the hybrid model for different number of semi-permeable barriers. Mean and standard deviation are reported.

the plume depths differ over time. The difference decreases beyond the point when the plume tip has reached the top for both models, because the subsequent absolute errors are zero. The maximum difference is largest for 20 semi-permeable barriers, and shows a decreasing trend for increased number of barriers.

The average depth of the plume for the same set of stochastic simulations is shown in Figure 7.24. Similarly to the depth of the plume tip, the average plume at a given time is deeper for more semi-permeable barriers. The apparent discontinuities around 20 years are due to immediate changes in force balance as injection stops. After the discontinuity, more CO2 exits the domain than migrates upwards. Since exited CO2 is not accounted for in the calculations, we observe a slight increase in average plume depth before it flattens as the spatial distribution of CO2 equilibrates. The $AADA$ appears to decrease at a constant rate for increased number of barriers. The difference in average plume depth between the two models at the end of the simulation quantifies how well the hybrid model reproduces the long-term distribution of mobile CO2. We see that the $AADA$ at the final time step is lower for more sealing barriers, where it is about one-third lower for 60 barriers than for 20 barriers.

Relative errors in CO2 mass categories summed over all time steps for the hybrid model are provided in Table 7.2, where the errors are relative to the full-dimensional model. Accounting for all categories, the net relative error is smallest for 20 semi-permeable barriers by a considerable margin, while being comparable in size for 40 and 60 semi-permeable barriers, indicating a positive correlation to the number of

semi-permeable barriers. The relative error for the mobile plume indicates a negative correlation, while errors in residual CO2 are comparable and do not reveal any correlation to the number of sealing barriers.

7.5 FluidFlower benchmark study

In this section we present models from MRST that are run on the new 11th SPE benchmark, using our semi-structured grid developed specifically for this case in addition to a fully-structured Cartesian grid. We first run a full-dimensional compositional simulator and then consider our extended fully-implicit hybrid model.

7.5.1 Compositional model

The compositional model supports dissolution of CO2 into the liquid phase. We use MRST to set up the simulations as described in Section 7.1. The model uses the FluidFlower grid generated in this thesis, as discussed in Section 4.1.

Concentration of CO2 in the liquid phase after 5 hours of injection is shown in Figure 7.25. The sealing effect of facie 1 is clear since hardly any CO2 enters the facie. Apart from the faults and facies 1 and 2, migration is predominantly vertical, indicating that a large part of the reservoir satisfies the VE assumption. Fluxes of CO2 are significant enough to penetrate the high-permeable pockets of the heterogeneous fault to fill the top-left part of the domain with dissolved CO2. In the semi-structured grid, CO2 penetrates the heterogeneous fault only through its top, while it also spreads laterally out the west of the fault in the pure Cartesian grid.

Concentration of CO2 in the liquid phase after two days for the semi-structured grid is shown in Figure 7.26a. At this stage, dissolution is prominent and widespread all over the domain. Convective mixing is apparent in the form of multiple gravity fingers that sink downwards due to the larger density of CO2-filled water. The fingers to the right of the high-permeable fault follow the orientation of the grid rather than the vertical, which is a consequence of the grid-orientation effect incurred by the TPFA scheme [38]. The dissolution rate is significant enough to transport dissolved CO2 away from the sharp interfaces and fill a large part of the domain.

Figure 7.26b shows the same simulation on a fully-structured Cartesian grid. This is discernible from the stair-stepping pattern at slopes and angles. The distribution of dissolved CO2 is qualitatively the same as in Figure 7.26a, apart from shallower gravity fingers descending strictly vertically, not subject to any grid-orientation effects.

Empirical observations of CO2 concentration in the physical rig in [56] after 5 hours and 48 hours are shown in Figure 7.25c and Figure 7.26c, respectively. Compared to our compositional simulations, the most prominent migration paths and fingers are resolved on both grids. The extent of the fingers is more accurately captured by the simulation on the semi-structured grid. While this simulation overestimates the convective mixing in the top-left of the domain, the simulation on the fully-structured grid appears to

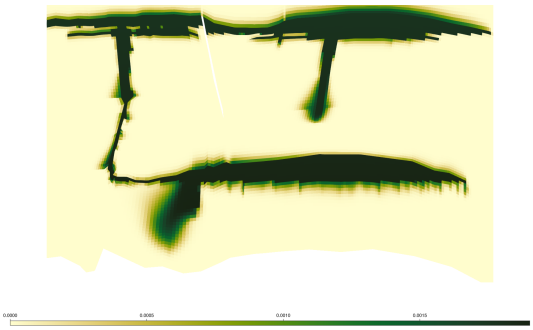

(a) Semi-structured grid

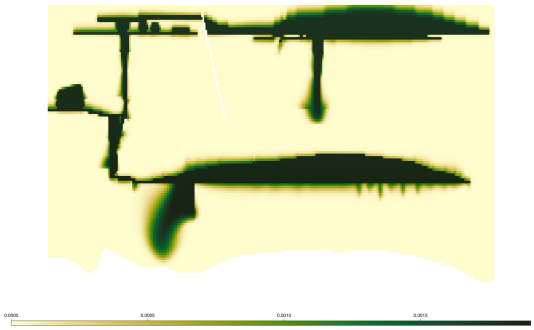

(b) Pure Cartesian grid

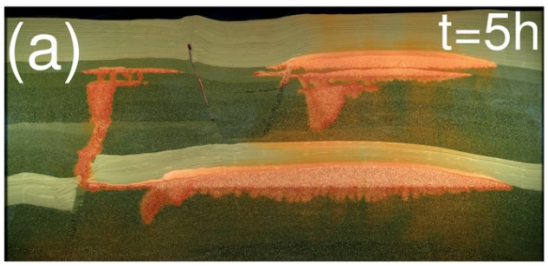

(c) Empirical observation

Figure 7.25: Concentration of CO2 in the liquid phase when injection stops (after 5 hours) for a compositional model. Upper-left figure is from a semi-structured grid adapted to FluidFlower. Upper-right figure is from a pure Cartesian grid. Bottom figure is empirical ground truth after 5 hours, from [56].

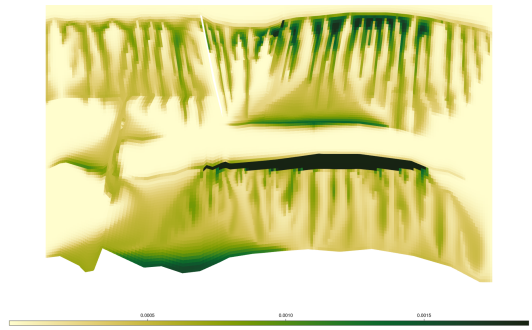

(a) Semi-structured grid

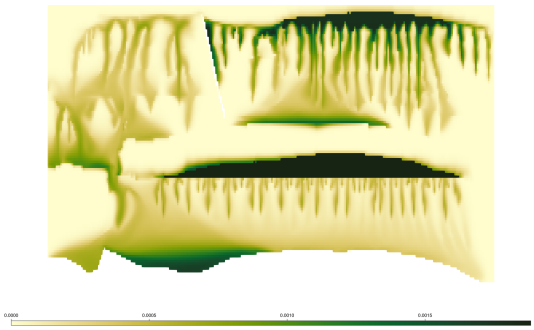

(b) Pure Cartesian grid

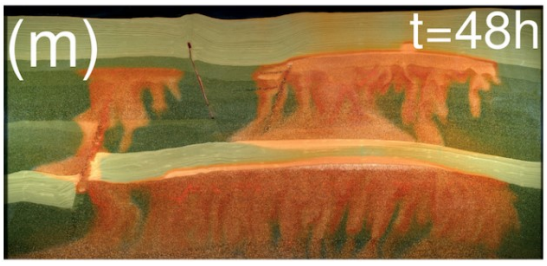

(c) Empirical observation

Figure 7.26: Concentration of CO2 in the liquid phase after 48 hours for a compositional model. Bottom figure taken from [56].

overestimate dissolved CO2 immediately to the left of the heterogeneous fault. Moreover, the thickness of mobile CO2 below the *lower* facie 1 is overestimated in simulations on the fully-structured grid, while the thickness below the *upper* facie 1 appears to be underestimated in simulations on the semi-structured grid. Overall, concentrations of CO2 from the two simulations are in good agreement with empirical observations.

7.5.2 Hybrid model

Next, we run our new hybrid model on the semi-structured grid for the FluidFlower and compare the results with a full-dimensional simulation. Important physical effects like dissolution and convective mixing are not modeled. Thus, comparing with empirical ground truth becomes irrelevant. We draw inspiration from petrophysical and fluid properties provided in Table 1 of [56] as well as simulations performed in [22], which reports mass distribution of CO2 after injection and migration. This shows that the primary seal, apart from the impermeable barriers, is facie 1. Among the remaining sand layers, it appears like facie 2 also inhibits vertical migration to some degree. We thus represent facie 2 as a semi-permeable barrier as well. That is, all facies in the structured parts of the grid are assumed to satisfy the VE assumption except facies 1 and 2. Faults and pinch-outs are modeled as full-dimensional cells regardless of their composition.

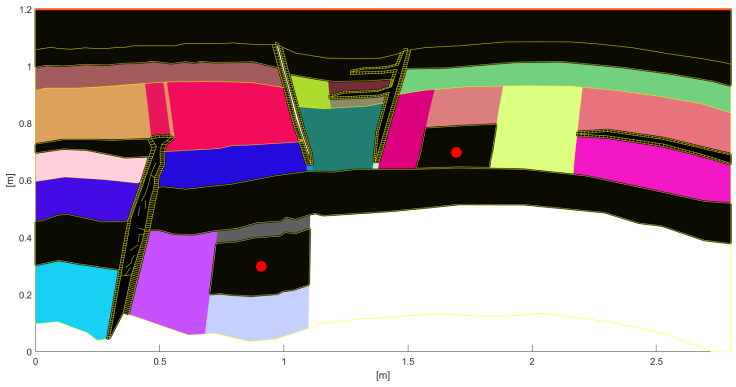

Figure 7.27: Partitioning of the semi-structured FluidFlower grid into different discretization regions. VE cells are marked by pure colors. Full-dimensional cells are dark and bounded by yellow edges. The two injection wells are marked by red dots.

Figure 7.27 shows a partition of the semi-structured grid into discretization regions. Notice that the cells at the boundary of a fault are modeled as full-dimensional "buffer" cells to give consistent fluxes at the transition between structured and unstructured regions. Regions in the vicinity of wells are also modeled with full-dimensional cells. Transitions between different facies in the structured parts of the grid, apart from facies 1 and 2, are discretized as vertical VE–to–VE. If the entry pressure of the upper facie is lower than the lower facie, CO2 will immediately rise to the top of the upper facie by the VE assumption, not accumulate below it. In the hybrid framework, this is modeled by removing the VE–to–VE transition and merging the associated VE regions, e.g., as

has been done in the light-yellow region to the right of the upper well.

Snapshots of fine-scale saturation at the end of the simulation for full-dimensional and hybrid models are shown in Figures 7.28a and 7.28b. We refer to Figure 4.1 for numbering of the facies. For the hybrid model, the capillary pressure does not exceed the entry pressure of facie 1 and CO2 accumulates below it, unable to penetrate even after injection ceases. For facie 2 on the other hand, only a single cell layer of CO2 is necessary to exceed the entry pressure and enter the layer. Accumulated CO2 reaching the top-right fault rapidly fills the top part due to its order of magnitude higher permeability.

By comparing with the full-dimensional solution, there is a clear mismatch in the final distribution of CO2. Whereas the full-dimensional model predicts solely accumulation below facie 2, significant amounts of CO2 penetrate facie 2 and spreads underneath facie 1 in the hybrid model. Through our implementation of residual saturation for RVE columns, it is possible to trace the plume back to facie 3, where the full-dimensional model reveals a significant lateral spreading of CO2.

To investigate the validity of the VE assumption for facie 3, we run a simulation of our hybrid model where facie 3 is discretized as full-dimensional cells instead of VE. The final saturation distribution is shown in Figure 7.28c. There is still a deviation in the final CO2 distribution, particularly below the upper facie 1 where CO2 has started accumulating in the hybrid model but not in the full-dimensional model. Still, the models are clearly in better agreement after invalidating the VE assumption for facie 3.

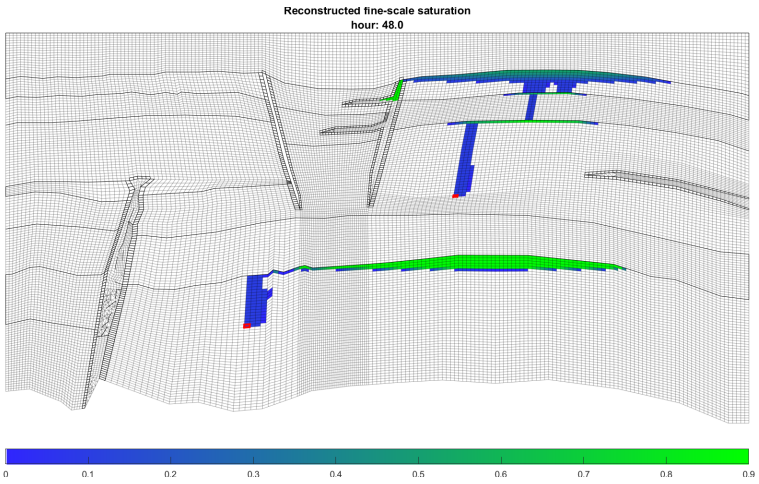

(a) Hybrid model.
Non-VE: facies 1 and 2.

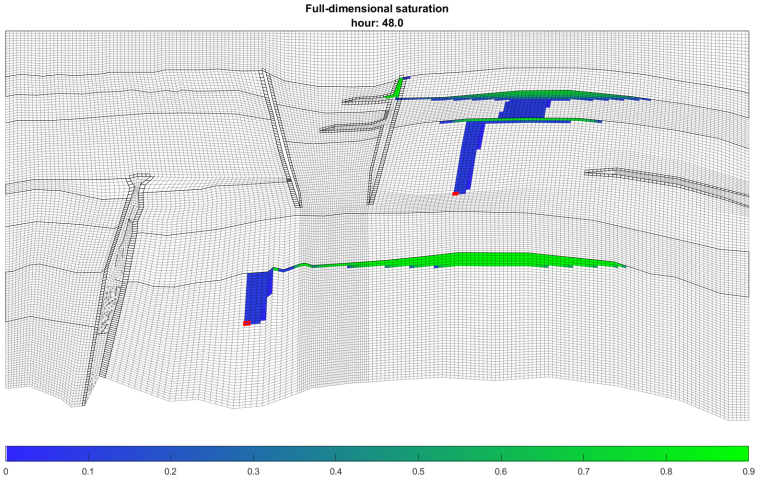

(b) Full-dimensional model.
Non-VE: facies 1 and 2.

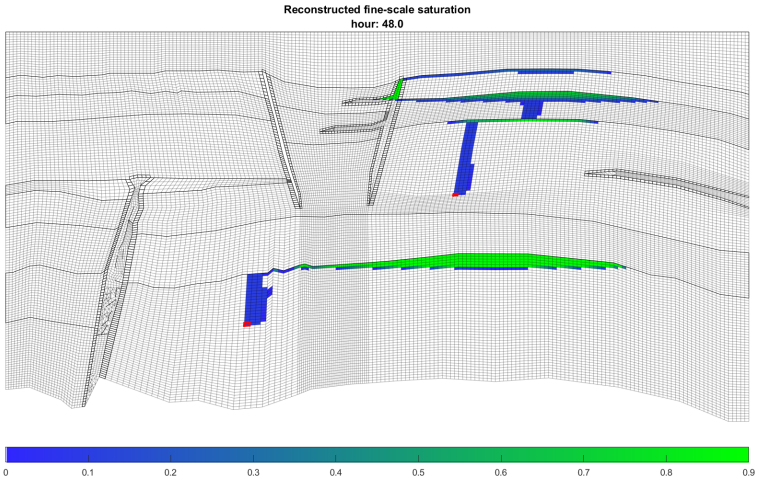

(c) Hybrid model.
Non-VE: facies 1, 2 and 3.

Figure 7.28: CO2 saturation in the FluidFlower at end of simulation (after 48 hours). Sealing barriers and buffer cells marked by thick black lines. Grid lines are shown to highlight the semi-structured cell topology. In a) and b), facies 1 and 2 are discretized with full-dimensional cells. In c), facie 3 is also discretized as full-dimensional.

Chapter 8

Discussion

The introduction of hybrid models for reservoir simulation provides new opportunities and insight into long-term CO2 storage simulation. In the following, we elaborate on the applicability of our hybrid model on heterogeneous formations in light of accuracy, computational efficiency, and degree of heterogeneity. We also discuss our contributions to the 11th SPE CSP, thereby the properties and potential of a semi-structured grid adapted for the FluidFlower, and results from a compositional and hybrid model.

In Section 7.1 we touched upon the effect of limited vertical resolution on resolving the sharp interface. The results demonstrated that a VE model may resolve the position of the sharp interface more accurately than a full-dimensional model. From this assessment, one should be critical of using the full-dimensional model as a reference solution. By using 100 cells in the vertical, the errors at the sharp interface for the full-dimensional model were minor. Since the other simulated setups in this thesis use at least this vertical resolution, we believe our full-dimensional models reasonably represent reference solutions [40]. Still, it is important to keep this artifact in mind and be aware that the full-dimensional solution is susceptible to errors near the sharp interface.

8.1 Comparison of hybrid and full-dimensional model

Results from the synthetic test cases contextualize the flexibility of the fully-implicit method by adding physical effects. For example, the introduction of residual saturation and RVE columns does not require any modifications of the governing equations, only a modification of the heights h and $h_{\max}$ used to evaluate some of the variables in the equation. In Section 7.1 we compared a face-representation and a cell-representation of thin sealing barriers. Although the former is less computationally expensive for heterogeneous formations, there is an inherent vagueness in approximating capillary exclusion with a transmissibility multiplier. The blocking effect of entry pressure is lost, as observed by comparing Figure 7.6 and Figure 7.7. Also, the fact that the multiplier must be prescribed based on heuristics makes it hard to justify in practice.

The results show that the governing flow paths of the hybrid model match that of the full-dimensional model. The deviations are attributable to the inevitable artifacts of the (relaxed) VE assumption for reconstructing residual CO2, i.e., by assuming strictly

vertical migration. This raises the question of the validity of the VE assumption in the high-permeable background regions of our synthetic reservoirs. To investigate this we estimate the time of segregation using (5.18). With a permeability of 100 mD and an average height between sealing barriers of 100 m, we find

$$t_s \approx 1.9 \,\text{years}.$$

With a simulation time of $T = 400$ years, we have that $t_s \ll T$. On the other hand, time steps are notably shorter, which challenges the VE assumption and helps explain the unresolved lateral spreading in the hybrid model before vertical segregation.

In Section 7.3 we remarked that the errors from the RVE assumption could be reduced by imposing additional fine cells around the tip of sealing barriers. Qualitatively, the simulation showed a significant increase in accuracy. However, the simulation time was nearly doubled by adding extra fine cells, and simulations for other configurations (not provided here) indicate a large variance to this end. To uncover any benefits of this approach, a stochastic analysis needs to be conducted, which we have not proceeded with herein.

Errors incurred by the RVE assumption could critically make the difference between accumulation below and penetration through a semi-permeable barrier, as observed in the right part of the domain in Figure 7.16. The observed trend in Figure 7.23 of decreasing $AADA$ in plume tip depth for increased number of semi-permeable barriers is as expected from the fact that the height of high-permeable VE regions, which are separated by sealing barriers, decreases for increased number of sealing barriers. In turn, this enhances the validity of the VE assumption according to (5.18). This is substantiated in Figure 7.23 by considering the time when the tip of the plume reaches the top of the formation. There is a significantly larger gap in time between the hybrid and full-dimensional model for few sealing barriers compared to many.

If the average location of the plume for the two models approximates the same depth during late-stage migration, it indicates that the distribution of mobile CO2 *inside* the domain complies well between the models. From Figure 7.24 we noticed that the $AADA$ at late-stage migration is smaller for more semi-permeable barriers. This suggests that the final distribution of mobile CO2 is not as accurately reproduced for few semi-permeable barriers as for many barriers. Table 7.2 supports this statement, which shows a significantly larger error in mobile CO2 for 20 semi-permeable barriers compared to 40 and 60 barriers. We believe this is attributed to the VE assumption being less appropriate for fewer semi-permeable barriers due to larger formation heights of the VE regions. Because there are fewer barriers to inhibit vertical migration, the hybrid model is more prone to miss structural trapping below other sealing barriers caused by the RVE assumption inevitably missing lateral spreading. By including more sealing barriers, the segregation time and lateral spread are not as extensive since mobile CO2 will quickly accumulate below the next sealing barrier, reducing the error from the RVE assumption. This is exemplified in one of the stochastic simulations of 20 sealing

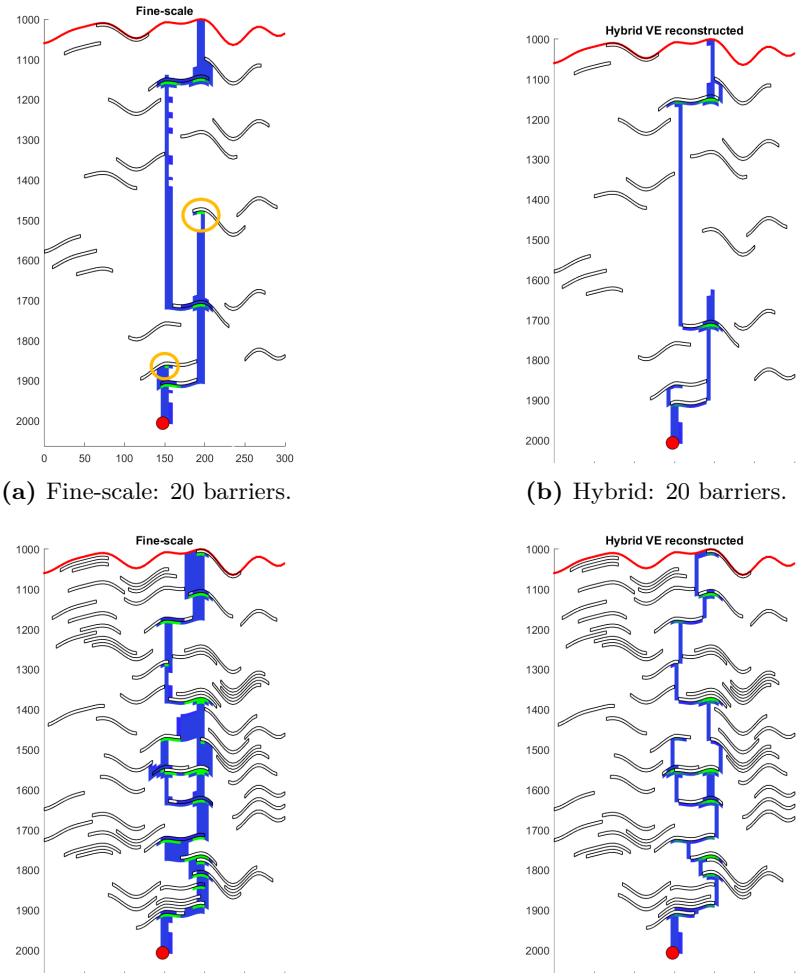

(a) Fine-scale: 20 barriers.

(b) Hybrid: 20 barriers.

(c) Fine-scale: 60 barriers.

(d) Hybrid: 60 barriers.

Figure 8.1: Saturation at final time step (400 years) from two selected stochastic simulations, for 20 and 60 semi-permeable barriers, respectively. Upper row: 20 barriers. Lower row: 60 barriers. Left column: full-dimensional model. Right column: hybrid model.

barriers, shown in the top row of Figure 8.1, where we observe that the lateral spread beyond the bottom barriers causes accumulation of mobile CO2 below two subsequent sealing barriers (marked by orange circles). This is not observed for the simulation of 60 sealing barriers (bottom row of Figure 8.1), partly explained by the short vertical distance between the barriers.

An additional factor contributing to the error in mobile CO2 is the small thickness of CO2 accumulating below sealing barriers, which for many barriers is represented by only one or a few vertical cells. Hence, the full-dimensional model is vulnerable to errors due to limited vertical resolution. This is illustrated in Figure 8.1 where we notice that under some of the barriers there are significantly more mobile CO2 in the full-dimensional model than the hybrid model, both for 20 and 60 sealing barriers. Additionally, limited spatial resolution may induce small errors in spill-point depths that collectively may give a significant error in mobile CO2 below sealing barriers [59, 50], an error that is therefore sensitive to the number and distribution of sealing barriers.

Layers	Model	Total time (s)	Median time (s)		Iterations	
		Mean	Mean	St. dev.	Sum	Mean
20	Fine	373	0.48	0.23	1066	1.9
	Hybrid	249	0.35	0.16	799	1.4
40	Fine	283	0.36	0.13	1213	2.1
	Hybrid	194	0.26	0.06	853	1.5
60	Fine	422	0.53	0.24	1364	2.4
	Hybrid	313	0.40	0.17	967	1.7

Table 8.1: Average of runtimes and number of nonlinear iterations for twenty stochastic simulations for each of the given number of semi-permeable barriers.

An interesting question is why the relative errors in residual CO2 (see Table 7.2) do not indicate any correlation to the number of sealing barriers, as opposed to mobile and exited CO2. By increasing the number of sealing barriers, a larger portion of the domain is subject to RVE modeling. Hence, we expect more accumulated errors from the RVE assumption, which give larger *absolute* error in residual trapping, but not necessarily larger *relative* error. If anything, it appears that more barriers give a slightly larger error, perhaps due to a more complex maze of barriers making it difficult to accurately capture individual plumes. This is supported by the saturation distribution in Figure 8.1, showing that the trace of plumes is not as accurately resolved for 60 barriers as for 20 barriers. In general, though, we need more data to substantiate this statement and draw any conclusions about how the error in residual CO2 correlates to the number of sealing barriers.

The seemingly large *magnitude* of the errors is also rooted in the RVE assumption. If viscous forces are non-negligible, lateral spreading of CO2 beyond sealing barriers becomes significant, and subsequent vertical migration may occupy more than one column. In the RVE assumption, CO2 migrating beyond the edge of a sealing barrier will only occupy one column (e.g., compare the plumes in Figure 8.1). Thus, in RVE columns it is common to obtain a relative error of at least 50%, a possible explanation for why the errors in residual trapped CO2 are close to this value. Despite inaccuracies, the RVE assumption manages to trace the history of individual plumes separated from the main plume. Exited volumes represent the largest mass category during the post-injection stage, which makes the net absolute error in trapping most sensitive to this mass category. This is the reason why the scaled $L1$ error for the sum of *all* mass categories is smallest for 20 semi-permeable barriers (see Table 7.2).

In Section 7.3, we observed that the performance of the hybrid model appears to be quite sensitive to the exact configuration of sealing barriers, with a few difficult time steps being responsible for mediocre runtimes. The result in Table 7.1 is not adequately representative of the performance difference between the models, because we obtained significantly different values by running a simulation with another configuration. To gain more insight into the performance of the hybrid model, we compare simulation times for the stochastic simulations in Section 7.4.

Average runtimes are provided in Table 8.1. The hybrid simulations run faster than

the full-dimensional simulations, both in terms of total time and median time, but not orders of magnitude faster. The reason is that there are very few cells in the synthetic cases considered in this thesis compared to what is customary for three-dimensional problems. It is likely that a major part of the simulation time is spent assembling the equations rather than solving the system of linearized equations. The real advantage of our hybrid model would materialize if it was applied to a realistic three-dimensional reservoir typically with hundredths of thousands to millions of cells [44].

Interestingly, Table 8.1 shows that simulations with 40 semi-permeable barriers run significantly faster than simulations with 20 and 60 barriers. The larger standard deviation suggests that the particular configurations of 20 and 60 semi-permeable barriers are more sensitive to convergence issues of the numerical solver. This substantiates our claim that the hybrid model is particularly sensitive to the exact configuration of sealing barriers. To investigate this, we look at the number of nonlinear iterations provided in Table 8.1. The number increases in conjunction with increased number of semi-permeable barriers due to more complex, non-equilibrium flow. Generally, the hybrid model uses fewer iterations since it does not explicitly model the fast vertical flow that usually requires time step cuts and more iterations to converge, in agreement with other studies of hybrid models [44, 43]. The iteration statistics do not show any correlation with the observed differences in runtime between 40 barriers and 20 and 60 barriers, but rather show an increasing trend for increased number of semi-permeable barriers. Hence, the significantly shorter runtime for 40 barriers is attributable to less expensive computations of the linearized equations during an iteration in the Newton solver. The difference is also observed in the full-dimensional solver, verifying that it is not an artifact of the hybrid model.

What exactly causes the severe convergence issues for the hybrid model is yet to be understood. We suspect two possible issues. The first is related to mass conservation. By careful inspection of the hybrid trapping inventories shown in Chapter 7, e.g., Figure 7.10b and Figure 7.14b, we notice that the net mass after injection ceases does not remain constant at the total injected mass, i.e., mass is not fully conserved. Since a finite volume scheme relies on the principle of mass conservation, this could cause convergence problems for the nonlinear solver. Another possible explanation is minor errors in computed phase fluxes for RVE columns. Our modified expressions for h and $h_{\max}$ in Section 6.3.2 involve evaluations of nonlinear min- and max-functions on standard doubles (non-AD variables) that may yield erroneous tracking of derivatives [39, pp. 200-254]. In turn, this may challenge the nonlinear solver. Finding the root cause requires more detailed analyses, and warrants further investigation.

Perhaps surprisingly, our results do not indicate that the hybrid simulations with few semi-permeable barriers have any greater computational advantage over the full-dimensional simulator than those with many semi-permeable barriers. However, the difference in the number of fine cells is not too significant for the number of sealing barriers considered herein, at least not compared to the resolution of the global grid. We

would need to include even more sealing barriers, or simulate for a domain with orders-of-magnitude more cells (e.g., a three-dimensional reservoir), to witness any significant computational advantage. Nonetheless, the net trapping error indicates a lower accuracy for increased number of semi-permeable barriers, attributable to accumulated errors from the RVE assumption. Combined with the fact that the computational advantage of a hybrid model generally decreases for increasing number of full-dimensional cells [10], it is reasonable to state that the benefit of hybrid models is quite sensitive to the number and distribution of sealing barriers. The fact that a CO2–brine system is highly sensitive to changes in reservoir parameters makes a stochastic modeling approach especially valuable to this end [52, 22].

8.2 FluidFlower

In the following we discuss our contributions to the 11th SPE CSP. We put the semi-structured grid adapted to FluidFlower into context by comparing simulations with conventional Cartesian grids and results from the literature. This leads to a consideration of numerical errors. We also reflect upon the potential of hybrid modeling for FluidFlower.

8.2.1 Semi-structured grid

Generating a semi-structured grid for the FluidFlower is unfortunately not an automated process, but once generated, we have a flexible, multi-model applicable grid for the FluidFlower. The full-dimensional version is detailed enough for a compositional simulator to yield accurate results, and the semi-structured discretization permits hybrid modeling. In contrast to conventional Cartesian grids applied in the literature [22], our semi-structured grid provides a discretization that conforms perfectly to the boundary of barriers and faults. Unfortunately, this comes at a cost. Parts of the grid become skewed, which introduces grid-orientation errors for TPFA schemes that will be discussed shortly. Moreover, the angled faults are responsible for nonuniform grid spacing in otherwise homogeneous, structured facies. For instance, in the bottom of the domain between 1.1 meters and 1.4 meters (see Figure 4.5), the cells are unnecessarily small, but this is inevitable in order to permit a structured discretization.

8.2.2 Hybrid modeling

The results in Section 7.5.2 showed a significant difference in CO2 saturation between the full-dimensional and hybrid models. We mentioned that the VE assumption might be violated in facie 2. Quantitatively, this is substantiated by calculating the time of segregation for facie 2. Plugging in the relevant petrophysical, geometrical, and fluid properties into (5.18), we get

$$t_s \approx 205\ s.$$

The characteristic relative permeability is unknown, so we set $k^*_{rw} = 0.1$ to give an estimate of the upper bound of segregation time. Time steps during late injection are $\Delta t_i = 150\,s$, and simulation time is $T = 48\,hr$. That is, $t_s \ll T$ is indeed satisfied, while the requirement $t_s < \Delta t_i$ is not. Since (5.18) only gives an estimate of segregation time, it is reasonable to believe that the VE assumption in facie 2 is not adequately justified.

We need to stress that the motivation for running our hybrid model on the FluidFlower is to demonstrate its versatility and potential, not to give as good match as possible with state-of-the-art simulators. First, simulation times are not a strong argument for using the hybrid model, since convergence issues prevent the fast simulation times we hope for. Convergence problems are presumably a combination of irregular cells that give transmissibilities very close to zero and a low density and high compressibility of gaseous CO2 that makes the system more challenging to solve numerically [21]. Second, full-dimensional cells cover too much of the domain for the computational advantages of a hybrid model to outweigh the reduced accuracy when applied to a heterogeneous formation like FluidFlower. And finally, dissolution and capillary fringe are not accounted for. Whereas the former is a dominant process, the latter is expected to be vanishingly small due to highly nonlinear capillary pressure functions and small simulation times [56, 22].

That said, our hybrid model is flexible in the sense that the grid is easily adjustable by merging VE regions or replacing them with full-dimensional cells. This was exemplified in Section 7.5.2 where we invalidated the VE assumption for facie 3 based on significant differences in lateral spreading between the models. Now the question is, what is responsible for the differences in CO2 saturation observed below facie 1? Since facie 3 is now resolved with the same fine-scale model, the root cause is traced back to the VE assumption in the region below facie 3 and above the upper right well, which may slightly overestimate the speed of the plume. Another sensible explanation is the nonlinearity of the fine-scale relative permeability that retards the advancing tip of the plume. This may inhibit plume migration sufficiently to prevent CO2 from accumulating below facie 1. The effect of a nonlinear relative permeability becomes negligible for buoyancy-dominated problems [44]. However, full-dimensional simulations reveal that viscous forces are significant during injection, which enhances the effect.

The computational demand of running compositional simulations on the FluidFlower is large [56]. Short time-stepping is needed to account for the rapid vertical migration occurring in the highest-permeable facies. The short time-stepping is further restricted by including compositional effects, due to additional nonlinearities. In turn, this challenges the VE assumption, and one should reassess the validity of VE regions. Nonetheless, we believe a compositional hybrid model would be an excellent candidate for mitigating the computational challenge.

An analysis of the real benefits of applying a hybrid model on the FluidFlower is not conducted herein. We merely try to convey the versatility and applicability of our extended hybrid model on more complex grids than conventional Cartesian. At the

time one manages to implement dissolution and resolve the convergence problems, the hybrid model will become more relevant for simulations on the FluidFlower. Since the computational gain of a hybrid model increases with the number of cells in the grid, we anticipate that the hybrid model will be particularly advantageous for simulations on the three-dimensional version of the FluidFlower [21]. Moreover, the hybrid model will make history-matching of plumes feasible [48] and permit greater refinement of lateral dimensions for increased accuracy.

8.2.3 Numerical errors

Despite the level of detail and flexibility of the semi-structured MRST grid generated specifically for the FluidFlower project, there are some disadvantages that should be mentioned. In Figure 4.5, we observe that the grid becomes skewed as we move from the west and east boundaries to the impermeable and high-permeable faults, respectively. Mathematically, the skewness is expressed as $n_{i,k} \nparallel c_{i,k}$ for cells i and k. Recall from Section 3.1 that this geometry introduces grid-orientation errors for the TPFA scheme that gives inconsistent numerical solutions. By using structured, Cartesian subgrids whose faces are aligned with the geometry of the layering, this numerical artifact is inevitably caused by the deviated faults not aligned in the axial directions. Alternatives for mitigating the grid-orientation effect that is not considered in this thesis include consistent discretization methods such as MPFA-O and nonlinear-TPFA [39].

Alternatively, a compromise that retains the structured grid representation and removes the grid-orientation errors is to discretize the faults by a stair-stepped pattern, with the disadvantage that the faults are only represented approximately [38]. This was investigated in Section 7.5.1 where we used a fully-structured Cartesian grid. By comparing CO2 concentration in the two grids, the grid-orientation error for the semi-structured grid is clear, particularly to the right of the high-permeable fault. Here, buoyancy-driven migration follows the local skewness of the grid, not the true vertical as expected, as depicted in Figure 7.25. This is further demonstrated in Figure 7.26, where we observe gravity fingers spreading along skewed cells. The grid-orientation effect is also responsible for some numerical diffusion [14], noticeable for the fingers in the lower part of the domain. On the other hand, slopes are smoother and better resolved in the semi-structured grid. This becomes especially apparent for buoyantly segregated CO2 tracing the topography of confining facies, which for several layers in the pure Cartesian grid are imprecisely approximated to be strictly horizontal. The same deficiency is apparent in other simulators in the literature [22, 63], whose severity depends on the grid resolution. To precisely resolve flow along boundaries of layering and faults will necessitate extremely fine resolution to remove the stair-stepping effect, which significantly increases the computational challenge. However, compositional flow in the FluidFlower is exceptionally demanding for nonlinear solvers, which restricts the level of grid refinement that is feasible for simulations. This grid refinement is avoided in our semi-structured grid; the layering is precisely discretized even with a rather coarse

resolution. This is an argument for developing alternatives to pure Cartesian grids.

In addition to grid-orientation effects, accurate modeling of convective dissolution is sensitive to cell geometry and resolution, both horizontally and vertically [22]. Since cells in the semi-structured grid have varying shape and size, this contributes to the considerable differences in viscous fingering between the two grids in Figure 7.26. Qualitatively, it seems like the overestimated thickness of mobile CO2 below facie 1 incurred on the pure Cartesian grid is larger than the underestimate incurred on the semi-structured grid. From this observation, combined with fingering shapes in the semi-structured grid being in better agreement with empirical results, it is reasonable to believe that the rate of dissolution is more accurate for simulations on the semi-structured grid.

As mentioned, one of the challenges of applying a hybrid model on the FluidFlower is that the high-permeable sand layers only satisfy VE vaguely, due to simulation time steps being very small. The VE assumption is further challenged by the fact that lateral pressure gradients for the meter-scale FluidFlower are relatively larger compared to subsurface reservoirs [36]. As such, FluidFlower is not necessarily a perfect candidate for hybrid modeling. Another challenge is the irregular shape of some of the VE columns that gives significant grid-orientation effects. This could be improved by testing our hybrid framework on a pure Cartesian grid and making a comparison, but this has not been considered in this thesis.

The compositional simulator run in this thesis, in addition to simulators from the literature [22, 56, 63], sheds light on the numerical challenges in simulating CO2 storage. The large variance in results obtained by state-of-the-art simulators is partly due to different functional relationships describing physical properties [22], for instance, employed dissolution model, but nonetheless emphasizes the need for comparative studies to improve the understanding of numerical errors. To this end, we believe the work of this thesis contributes to this cause by expanding the spectrum of numerical methods, algorithms, and grids for CO2 storage in general and FluidFlower in particular.

8.3 Improvements

Although our extended hybrid model shows promising results on multilayered systems, there are still unresolved issues. First, a more comprehensive stochastic analysis is needed to verify correlations between the accuracy of our hybrid model and the distribution of sealing barriers. First and foremost this involves more data. Another suggestion is to investigate relations to other parameters such as entry pressure, permeability, and thickness of sealing barriers.

Second, we have only considered homogeneous properties for each geological layer. To facilitate more realistic scenarios, subscale heterogeneities should be included. This challenges the RVE assumption for our hybrid model. In the absence of novel modeling ideas, one has to rely on fine-scale discretization here [69].

Third, cells in the discretized domain are static, with no feature for dynamically adapting to the state of the system, e.g., turning full-dimensional cells into VE cells if buoyant forces become dominant in the associated region. This feature is inspired by the promising results in [10, 11].

Experiments from the FluidFlower rig suggest that dissolution happens almost instantaneously in the early stages of injection [56], and can thus be modeled at equilibrium. Later, in regions where CO2 becomes buoyantly segregated on top of brine, VE applies and dissolution is more appropriately modeled as rate-dependent [4, 57]. Extending the hybrid framework with a compositional model incurs additional numerical challenges [45], though it is possible to apply simplifications to the model to mitigate the cost. For instance, AbdAllah et al. [68] justified the use of an IMPES scheme where dissolution was only modeled in the transport equation, not the pressure equation. The study showed promising results. An attempt was made to implement a combined rate-driven - instantaneous dissolution model [57], but without success due to non-convergent solutions. We believe that combining compositional hybrid models implemented in, e.g.,[43] and [10], with the extensions developed in this thesis is a promising approach to synthesize dissolution, residual saturation, and capillary pressure in a hybrid framework.

Chapter 9

Conclusion

Using MRST as a digital toolbox for reservoir simulation, we have extended a fully-implicit hybrid model with physical effects including residual saturation and capillary exclusion, where the former was implemented based on an alternative relaxed VE assumption inspired by [26]. For geological CO2 storage, the large temporal and spatial scales and nonlinear couplings of the governing equations pose challenges for numerical solvers. A fully-implicit formulation with Newton's method is more robust than a sequential splitting strategy by enabling long time steps. This flexibility has been leveraged in this thesis when simulating heterogeneous reservoirs of varying complexity, at the laboratory scale and subsurface scale, where significant time-step cutting was necessary to guarantee convergence. The solution of our extended hybrid model generally matches well with full-dimensional simulations for semi-complex heterogeneous reservoirs. The errors are primarily attributed to the RVE assumption.

Results from the synthetic test cases have demonstrated the appropriateness and benefits of using a hybrid model to simulate CO2 migration in semi-complex heterogeneous formations, compared to conventional VE and full-dimensional models. Simulation times for our hybrid model are shorter than simulation times for full-dimensional models while at the same time showing good agreement with the full-dimensional solution. A great increase in runtime is expected by applying our hybrid model on realistic three-dimensional formations with an orders-of-magnitude increase in number of cells compared to the synthetic two-dimensional grids considered herein. We have demonstrated the effect of vertical resolution and relative permeability function on the accuracy of the solution.

The location of the mobile plume is a vital factor dictating storage security. To this end, we have calculated the location of the tip and average mobile plume for both models. The tip of the plume in the hybrid model migrates faster than in the full-dimensional model, as anticipated from the VE assumption. The difference in depth of the plume tip decreases the more sealing barriers are included, and the $AADA$ is smallest for 60 barriers, both the tip depth and average depth.

Our stochastic simulations show that computation times are consistently shorter for the hybrid model, though the RVE assumption incurs additional errors in CO2 distribution, whose mass categories appear to be differently correlated to the number

of sealing barriers. Increasing the number decreases the error in mobile CO2, while it shows no particular correlation to errors in residual CO2. The former is supported by calculations of $AADA$, which is negatively correlated to the number of sealing barriers at the end of the simulation. Different categories show different correlations, but the net trapping error indicates greater accuracy for fewer sealing barriers. We recommend a more thorough stochastic analysis to justify the observed correlations.

As a second objective of the thesis, we have performed simulations on the FluidFlower, as part of the 11th SPE CSP, for a compositional model and a hybrid model. All programmatic routines, from data extraction to simulation, are performed solely in MRST. Specifically, in Section 4.1 we demonstrated how to generate a detailed semi-structured grid to represent polygonal data of the FluidFlower. We believe the composite grid is an important contribution to the ongoing research on the FluidFlower project. Although the grid is subject to grid-orientation errors, the high level of detail accurately reproduces the geology of the rig while also being compatible with hybrid modeling. To put simulations on the grid into context, we compared it with a fully-structured Cartesian grid not suffering from grid-orientation effects. The results for the compositional, full-dimensional model were in good agreement with empirical observations, for both grids.

The versatility of our extended hybrid model was accentuated by applying it to the semi-structured FluidFlower grid. By specifying facies that are expected to satisfy the VE assumption, we can adaptively modify the discretization of relevant regions, as exemplified in Figure 7.28. We compared our hybrid model with a corresponding full-dimensional model without dissolution, and the results were decent. The observed deviations were as expected. Conclusively, a complete extension of our hybrid model, synthesizing dissolution with residual saturation and capillary exclusion, is a promising candidate for the FluidFlower. This is particularly true for the three-dimensional version, which is anticipated to greatly reduce the computational demand of compositional simulations. We believe this will be a valuable contribution to the 11th SPE CSP and give new insight into the numerical challenges of CO2 storage simulation.

Bibliography

[1] Nurafiqah Abdullah et al. 'The study of the effect of fault transmissibility on the reservoir production using reservoir simulation—Cornea Field, Western Australia'. In: *Journal of Petroleum Exploration and Production Technology* 10 (2020), pp. 739–753.

[2] Takashi Akai et al. 'Numerical modelling of long-term CO2 storage mechanisms in saline aquifers using the Sleipner benchmark dataset'. In: *International Journal of Greenhouse Gas Control* 110 (2021), p. 103405.

[3] Ali Hasan Ali et al. 'A comparison of finite difference and finite volume methods with numerical simulations: Burgers equation model'. In: *Complexity* 2022 (2022).

[4] Odd Andersen. 'Simplified models for numerical simulation of geological CO2 storage'. PhD thesis. University of Bergen, 2017. URL: https://bora.uib.no/bora-xmlui/handle/1956/15477.

[5] Odd Andersen, Sarah Eileen Gasda and Halvor M Nilsen. 'Vertically averaged equations with variable density for CO 2 flow in porous media'. In: *Transport in Porous Media* 107.1 (2015), pp. 95–127.

[6] Karl W Bandilla, Michael A Celia and Evan Leister. 'Impact of model complexity on CO2 plume modeling at Sleipner'. In: *Energy Procedia* 63 (2014), pp. 3405–3415.

[7] Karl W Bandilla, Bo Guo and Michael A Celia. 'A guideline for appropriate application of vertically-integrated modeling approaches for geologic carbon storage modeling'. In: *International Journal of Greenhouse Gas Control* 91 (2019), p. 102808.

[8] Jacob Bear. *Dynamics of fluids in porous media*. Courier Corporation, 2013.

[9] Beatrix Becker et al. 'A pseudo-vertical equilibrium model for slow gravity drainage dynamics'. In: *Water Resources Research* 53.12 (2017), pp. 10491–10507.

[10] Beatrix Becker et al. 'An Adaptive Hybrid Vertical Equilibrium/Full-Dimensional Model for Compositional Multiphase Flow'. In: *Water Resources Research* 58.1 (2022), e2021WR030990.

[11] Beatrix Becker et al. 'An adaptive multiphysics model coupling vertical equilibrium and full multidimensions for multiphase flow in porous media'. In: *Water Resources Research* 54.7 (2018), pp. 4347–4360.

[12] Runar Lie Berge, Øystein Strengehagen Klemetsdal and Knut-Andreas Lie. 'Unstructured Voronoi grids conforming to lower dimensional objects'. In: *Computational Geosciences* 23 (2019), pp. 169–188.

[13] University of Bergen. *Flow in Porous Media Matters: The FluidFlower Concept*. URL: https://fluidflower.w.uib.no/. Accessed: 26.04.2023.

[14] CW Brand, JE Heinemann and K Aziz. 'The grid orientation effect in reservoir simulation'. In: *SPE Symposium on Reservoir Simulation*. OnePetro. 1991.

[15] Yann Brenier and Jérôme Jaffré. 'Upstream differencing for multiphase flow in reservoir simulation'. In: *SIAM journal on numerical analysis* 28.3 (1991), pp. 685–696.

[16] Andrew J Cavanagh and R Stuart Haszeldine. 'The Sleipner storage site: Capillary flow modeling of a layered CO2 plume requires fractured shale barriers within the Utsira Formation'. In: *International Journal of Greenhouse Gas Control* 21 (2014), pp. 101–112.

[17] Michael Anthony Celia et al. 'Status of CO2 storage in deep saline aquifers with emphasis on modeling approaches and practical simulations'. In: *Water Resources Research* 51.9 (2015), pp. 6846–6892.

[18] Benjamin Court et al. 'Applicability of vertical-equilibrium and sharp-interface assumptions in CO2 sequestration modeling'. In: *International Journal of Greenhouse Gas Control* 10 (2012), pp. 134–147.

[19] CO2 DataShare. *Datasets*. URL: https://co2datashare.org/. Accessed: 26.04.2023.

[20] Martin A Ferno et al. 'Room-scale CO2 injections in a physical reservoir model with faults'. In: *arXiv preprint arXiv:2301.06397* (2023).

[21] Bernd Flemisch et al. 'Draft Description: The 11th Society of Petroleum Engineers Comparative Solution Project'. 2023.

[22] Bernd Flemisch et al. 'The FluidFlower international benchmark study: Process, modeling results, and comparison to experimental data'. In: *arXiv preprint arXiv:2302.10986* (2023).

[23] Sarah E Gasda, Jan M Nordbotten and Michael A Celia. 'The impact of local-scale processes on large-scale CO2 migration and immobilization'. In: *Energy Procedia* 4 (2011), pp. 3896–3903.

[24] Sarah E Gasda, Jan M Nordbotten and Michael A Celia. 'Vertical equilibrium with sub-scale analytical methods for geological CO 2 sequestration'. In: *Computational Geosciences* 13 (2009), pp. 469–481.

[25] Christophe Geuzaine and Jean-Francois Remacle. *Gmsh 4.11.1*. 2022. URL: https://gmsh.info/doc/texinfo/gmsh.html.

[26] Bo Guo. 'Fluid injection and migration in the subsurface: Reduced-order models and multiscale modeling approaches'. PhD thesis. Princeton University, 2016.

[27] Bo Guo et al. 'A multiscale multilayer vertically integrated model with vertical dynamics for CO2 sequestration in layered geological formations'. In: *Water Resources Research* 52.8 (2016), pp. 6490–6505.

[28] Bo Guo et al. 'A vertically integrated model with vertical dynamics for CO2 storage'. In: *Water Resources Research* 50.8 (2014), pp. 6269–6284.

[29] Bo Guo et al. 'Flow regime analysis for geologic CO2 sequestration and other subsurface fluid injections'. In: *International Journal of Greenhouse Gas Control* 53 (2016), pp. 284–291.

[30] Malin Haugen et al. 'Physical variability in meter-scale laboratory CO_2 injections in faulted geometries'. In: *arXiv preprint arXiv:2301.07347* (2023).

[31] Herbert E Huppert and Jerome A Neufeld. 'The fluid mechanics of carbon dioxide sequestration'. In: *Annual review of fluid mechanics* 46 (2014), pp. 255–272.

[32] Fernando Sandro Velasco Hurtado, Clovis Raimundo Maliska and AFC Silva. 'On the factors influencing the grid orientation effect in reservoir simulation'. In: *19th International Congress of Mechanical Engineering, November*. 2007, pp. 5–7.

[33] Ruben Juanes, Christopher W MacMinn and Michael L Szulczewski. 'The footprint of the CO 2 plume during carbon dioxide storage in saline aquifers: storage efficiency for capillary trapping at the basin scale'. In: *Transport in porous media* 82 (2010), pp. 19–30.

[34] Ruben Juanes et al. 'Impact of relative permeability hysteresis on geological CO2 storage'. In: *Water resources research* 42.12 (2006).

[35] Eirik Keilegavlen et al. 'PoroTwin: A digital twin for a FluidFlower rig'. In: *arXiv preprint arXiv:2212.00865* (2022).

[36] Anthony R Kovscek, Jan Martin Nordbotten and Martin A Ferno. 'Scaling up FluidFlower results for carbon dioxide storage in geological media'. In: *arXiv preprint arXiv:2301.09853* (2023).

[37] Hans Petter Langtangen and Svein Linge. *Finite difference computing with PDEs: a modern software approach*. Springer Nature, 2017.

[38] Knut-Andreas Lie. *An introduction to reservoir simulation using MATLAB/GNU Octave: User guide for the MATLAB Reservoir Simulation Toolbox (MRST)*. Cambridge University Press, 2019.

[39] Knut-Andreas Lie and Olav Møyner. *Advanced modelling with the MATLAB reservoir simulation toolbox*. Cambridge University Press, 2021.

[40] I Ligaarden and Halvor Møll Nilsen. 'Numerical aspects of using vertical equilibrium models for simulating CO2 sequestration'. In: *ECMOR XII-12th European conference on the mathematics of oil recovery*. EAGE Publications BV. 2010, cp–163.

[41] Bert Metz et al. 'Report: Carbon Dioxide Capture and Storage'. In: (2005), p. 431.

[42] Olav Møyner. 'Next Generation Multiscale Methods for Reservoir Simulation'. PhD thesis. NTNU, 2016. URL: https://ntnuopen.ntnu.no/ntnu-xmlui/handle/11250/2431831.

[43] Olav Møyner, Odd Andersen and Halvor M Nilsen. 'Multi-model hybrid compositional simulator with application to segregated flow'. In: *Computational Geosciences* 24.2 (2020), pp. 775–787.

[44] Olav Møyner and Halvor Møll Nilsen. 'Multiresolution coupled vertical equilibrium model for fast flexible simulation of CO2 storage'. In: *Computational Geosciences* 23.1 (2019), pp. 1–20.

[45] Olav Møyner and Hamdi A Tchelepi. 'A mass-conservative sequential implicit multiscale method for isothermal equation-of-state compositional problems'. In: *SPE Journal* 23.06 (2018), pp. 2376–2393.

[46] Halvor Møll Nilsen, Knut-Andreas Lie and Odd Andersen. 'Analysis of CO2 trapping capacities and long-term migration for geological formations in the Norwegian North Sea using MRST-co2lab'. In: *Computers & Geosciences* 79 (2015), pp. 15–26.

[47] Halvor Møll Nilsen, Knut-Andreas Lie and Odd Andersen. 'Fully-implicit simulation of vertical-equilibrium models with hysteresis and capillary fringe'. In: *Computational Geosciences* 20 (2016), pp. 49–67.

[48] Halvor Møll Nilsen, Knut-Andreas Lie and Odd Andersen. 'Robust simulation of sharp-interface models for fast estimation of CO2 trapping capacity in large-scale aquifer systems'. In: *Computational Geosciences* 20.1 (2016), pp. 93–113.

[49] Halvor Møll Nilsen et al. 'Field-case simulation of CO2-plume migration using vertical-equilibrium models'. In: *Energy Procedia* 4 (2011), pp. 3801–3808.

[50] Halvor Møll Nilsen et al. 'Spill-point analysis and structural trapping capacity in saline aquifers using MRST-co2lab'. In: *Computers & Geosciences* 75 (2015), pp. 33–43.

[51] Jan M Nordbotten and Helge K Dahle. 'Impact of the capillary fringe in vertically integrated models for CO2 storage'. In: *Water Resources Research* 47.2 (2011).

[52] Jan M. Nordbotten and Michael A. Celia. *Geological Storage of CO2: Modeling Approaches for Large-Scale Simulation*. John Wiley & Sons, 2012. ISBN: 9780470889466.

[53] Stack Overflow. *Algorithm to generate equally distributed points inside a polygon*. Accessed: 2023-03-20. 2012. URL: https://stackoverflow.com/questions/11178414/algorithm-to-generate-equally-distributed-points-in-a-polygon.

[54] Society of Petroleum Engineers. *The 11th Society of Petroleum Engineers Comparative Solution Project*. URL: https://spe.org/en/csp/. Accessed: 26.04.2023.

[55] Hans Petter Langtangen. *Finite Difference Computing with Exponential Decay Models*. Springer Nature, 2016.

[56] Lluis Salo-Salgado et al. 'Direct comparison of numerical simulations and experiments of CO2 injection and migration in geologic media: Value of local data and predictability'. In: *arXiv preprint arXiv:2301.08875* (2023).

[57] Tor Harald Sandve et al. 'Convective Dissolution in Field Scale Co2 Storage Simulations Using the OPM Flow Simulator'. In: *TCCS–11. CO2 Capture, Transport and Storage. Trondheim 22nd–23rd June 2021 Short Papers from the 11th International Trondheim CCS Conference*. SINTEF Academic Press. 2021.

[58] Jeremy G Siek, Lie-Quan Lee and Andrew Lumsdaine. *The Boost Graph Library: User Guide and Reference Manual, The*. Pearson Education, 2001.

[59] SINTEF. *Effects of Data Resolution: Estimation of Trap Analysis Accuracy*. 2013. URL: https://www.sintef.no/projectweb/mrst/modules/co2lab/traps/estimation-of-trap-analysis-accuracy/. Accessed: 21.12.2022.

[60] SINTEF. *MRST-co2lab*. 2016. URL: https://www.sintef.no/projectweb/mrst/modules/co2lab/. Accessed: 13.05.2023.

[61] SINTEF. *MRST: Transforming research*. URL: https://www.sintef.no/projectweb/mrst/. Accessed: 21.12.2022.

[62] SINTEF. *Vertical-Equilibrium Models*. 2014. URL: https://www.sintef.no/projectweb/mrst/modules/co2lab/ve-models/. Accessed: 21.12.2022.

[63] Yuhang Wang et al. 'Simulation of CO2 Storage using a Parameterization Method for Essential Trapping Physics: FluidFlower Benchmark Study'. In: *arXiv preprint arXiv:2301.08004* (2023).

[64] Francesca Watson, Odd Andersen and Halvor Møll Nilsen. 'Rapid Optimisation of the New Sleipner Benchmark Model'. In: *TCCS–11. CO2 Capture, Transport and Storage. Trondheim 22nd–23rd June 2021. Short Papers from the 11th International Trondheim CCS Conference*. SINTEF Academic Press. 2021.

[65] Mark Wilkinson et al. 'CO2–mineral reaction in a natural analogue for CO2 storage—implications for modeling'. In: *Journal of Sedimentary Research* 79.7 (2009), pp. 486–494.

[66] GA Williams and RA Chadwick. 'Influence of reservoir-scale heterogeneities on the growth, evolution and migration of a CO2 plume at the Sleipner Field, Norwegian North Sea'. In: *International Journal of Greenhouse Gas Control* 106 (2021), p. 103260.

[67] Zhibing Yang et al. ‘Modeling immiscible two-phase flow in rough fractures from capillary to viscous fingering’. In: *Water Resources Research* 55.3 (2019), pp. 2033–2056.

[68] AbdAllah A Youssef, Qi Shao and SK Matthäi. ‘Simplified Numeric Simulation Approach for $CO_{2,g}$-Water Flow and Trapping at Near-Surface Conditions’. In: *arXiv preprint arXiv:2301.09049* (2023).

[69] Tianyuan Zheng, Bo Guo and Haibing Shao. ‘A hybrid multiscale framework coupling multilayer dynamic reconstruction and full-dimensional models for CO2 storage in deep saline aquifers’. In: *Journal of Hydrology* 600 (2021), p. 126649.