



Norwegian University of
Science and Technology

Simulations of Underground Storage of Natural Gas

Ingeborg Gåseby Gjerde

Master of Science in Physics and Mathematics

Submission date: July 2016

Supervisor: Knut Andreas Lie, MATH

Co-supervisor: Halvor Møll Nilsen, Sintef

Norwegian University of Science and Technology
Department of Mathematical Sciences

Simulations of Underground Storage of Natural Gas

Ingeborg Gåseby Gjerde

July 22, 2016

Abstract

The purpose of this thesis is to simulate fluid flows in aquifers used for natural gas storage. An introduction is given to aquifer flow modelling, and two models are developed that describe two-phase flow in porous media. The first model is constructed using the standard equations for Darcy flow and conservation of mass. The second model is then developed by integrating these equations in the vertical direction, and using the assumption of a sharp-interface and vertical equilibrium to evaluate the integrals.

These two models were tested by running simulations on the aquifer model of Stenlille, a major gas storage facility in Denmark. The specifics of the aquifer model, such as the grid, fluid properties and so on, were based upon an existing ECLIPSE model. The goal was to improve the capability the simulations had in predicting pressures and water production rates for the wells in the aquifer. To this end, the ECLIPSE model of the aquifer was transferred into the MRST framework, where it was easier to experiment with the models themselves.

The simulations run in MRST, using both the black-oil model and the VE model, were found to give successful predictions of the well pressures. As the VE model is two-dimensional, run-times were significantly lowered when this model was utilized. Neither the simulations run in MRST nor ECLIPSE were found to give successful predictions of water production rates.

Sammendrag

Målet med denne oppgaven er å simulere fluidflyt i et akvifer som brukes til å lagre metan. En introduksjon vil bli gitt angående modellering av akviferer, og to modeller vil bli utviklet for to-fase flyt i porøse medier. Den første modellen er basert på en konserveringslov og et standard uttrykk for Darcy flyt. Den neste modellen videreutvikles basert på denne, ved å integrere ligningene vertikalt. Antagelser om vertikalt ekvilibrium og en skarp overgang mellom de to fluidfasene brukes til å evaluere integralene.

De to modellene ble testet ved å bruke dem til å simulere fluidflyt i Stenlille akviferet, som er et stort metanlager i Danmark. Den spesifikke modellen for dette akviferet ble basert på en tidligere utviklet ECLIPSE modell. Målet er å forbedre evnen simuleringene basert på denne modellen har til å forutsi trykkforandringer og vannproduksjon i brønnene.

For å kunne eksperimentere friere med selve akvifer modellen, ble den reimplementert i MRST rammeverket. Simuleringer ble så kjørt ved å bruke denne modellen. Det ble konkludert med at både standard-modellen og VE modellen, slik de ble implementert i MRST, kunne predikere trykkforandringer med rimelig nøyaktighet. Siden VE modellen i praksis er to-dimensjonal, hadde VE simuleringene betydelig lavere kjøretid. Hverken simuleringer i ECLIPSE, eller i MRST, klarte å predikere vannproduksjon i brønnene.

Preface

This work fulfils the Master's thesis requirement for the Master's Programme in Applied Physics and Mathematics at NTNU. The work has been performed at SINTEF ICT's Department of Applied Mathematics in Oslo.

I would like to thank my advisors, Knut-Andreas Lie and Halvor Møll Nilsen at SINTEF ICT, for their guidance and valuable feedback. I would also like to thank Rebecca Allen and Odd Andersen for answering my numerous questions about vertical equilibrium models, and Bård Skaffestad for helping me unravel the mysteries of ECLIPSE.

Lastly, a great thanks to my family and friends, as I could not have completed my degree without you.

Contents

1	Introduction	1
2	Aquifer Flow Modeling	4
2.1	An Introduction to Aquifers	4
2.2	Porous Media Flow	5
2.2.1	Law of Mass Conservation	5
2.2.2	Darcy’s Law	7
2.2.3	Two-Phase Flow in Porous Media	7
2.2.4	The Black Oil Model	10
2.2.5	Model Reformulations	11
2.3	A Vertical Equilibrium Model for Aquifer Flow	13
2.3.1	Deriving the Vertically Integrated Model	14
2.3.2	Rewriting Into a Fractional Flow Formulation	17
3	Numerical Methods and Implementation	20
3.1	Discretization	20
3.2	Newtons Method and Automatic Differentiation	24
3.3	Well Equations	25
3.4	The Solver	26
3.4.1	The GMRES Method	26
3.4.2	The CPR Preconditioner	27
3.5	The MATLAB Reservoir Simulation Toolbox	29
3.5.1	Integration of Eclipse Models	31
3.5.2	The co2lab module	31
4	Description of the Stenlille Model	33
4.1	The Three-Dimensional Model	34
4.1.1	Grid and Boundary Conditions	34
4.1.2	Rock Properties	40
4.1.3	Fluid Properties	42

4.1.4	Initial Conditions	45
4.1.5	Wells	46
4.2	The Two-Dimensional Model	47
5	Simulation Results	51
5.1	A Note on the ECLIPSE Simulation Results	54
5.2	Simulation A: Black-Oil Model of Stenlille	57
5.3	Simulation B.1: VE Model of Stenlille	65
5.4	Simulation B.2: VE Simulations on Refined Grid	69
6	Conclusions	77
A	Results of Simulation A	83
A.1	Well Bottom-Hole Pressures	83
A.2	Water Production	87
B	Results of Simulation B	90
B.1	Bottom hole pressures	90
B.2	Water Production	95
C	Other Plots	98
C.1	Individual Well Schedules	99
C.2	Comparison of ECLIPSE Simulation Results and Historical Values .	103
C.2.1	Tubing Head Pressures	103
C.2.2	Water Production Rate	107

Chapter 1

Introduction

Natural gas is a vital fuel source in the world economy. The gas itself is a hydrocarbon mixture that consists primarily of methane. When burned, it gives off considerable amounts of energy, and can therefore be efficiently used for electricity generation, heating and cooking. In 2013, it supplied an estimated 15.1 % of the worlds total energy consumption, a percentage that is close to that of electricity and surpasses coal [1].

The gas is primarily created deep underground, where the pressure and earth temperature is high enough to cook organic matter. The gas created in this process has low density compared to other fluids, and will therefore slowly rise towards the surface. Its upwards migration will continue until it eventually reaches the earths surface and dissipates, unless it meets a geological formation capable of trapping it along the way. Such a formation would consist of a layer of porous rock where the gas can reside, that is topped by a denser, impermeable rock in the shape of an inverted saucer.

If undisturbed, natural gas can reside can be trapped in such a formation for millions and millions of years. When such a reservoir is discovered, wells can be drilled into it to bring the gas to the surface, and the gas can then be directed via pipelines to the place where it will be used as fuel. This process of transporting natural gas is difficult and costly compared to other fossil fuels [5]. Furthermore, the consumption rate often exhibits strong seasonal variations, and it is not always possible to match the production rate with the consumption rate. This immediately raise the question - how can one store gas in a manner that is both efficient and safe?

The answer comes as a surprise to many people: The earth itself can be used efficiently for storage. Gas has already been stored for millions of years in under-




Figure 1.1: A schematic representation of the Stenlille gas storage facility. Figure from [11]

ground aquifers, and the same method of storage will do equally well in the short term. Perhaps even more surprising, this is by no means a new invention; it was utilized for the first time over a 100 years ago, with major developments happening since the 1940s [8].

The Stenlille gas storage facility is located in Denmark, and is the largest storage facility in Scandinavia. The gas is stored roughly 1500 meters below the surface, in a sandstone aquifer that is separated into distinct storage zones by thin layers of impermeable claystone.

A total of thirteen different wells have been drilled down into the aquifer, of which three are observation wells that flank it. The main operational restraint during injection is that the injection rate and pressure must be kept beneath the overburden pressure of the rock, and low enough to not push gas out of the structure. During withdrawal, the withdrawal rate per well must be kept low due to risk of *water coning* and sand production.

So far, the injection and production rates have been set based on operational experience, reservoir simulations in ECLIPSE and comparisons with other aquifer gas storage facilities. The numerical simulations have been able to predict the well pressures with reasonable success, along with the large-scale migration of gas in the aquifer. However, simulations have so far been unsuccessful in predicting when large production rates lead to water being produced instead of gas. This brings us to the topic of this thesis - numerical simulations of natural gas storage. With the hope of discovering improvements that can be made upon the simulations, the

ECLIPSE model of the aquifer has been transferred into the Matlab Reservoir Simulation Toolbox (MRST) framework, where it is easier to experiment with the model itself.

The MRST model will seek to improve both the accuracy of the simulation predictions, and reduce the run-time of the simulations themselves. The latter will be done by using Vertical Equilibrium (VE) models, in which case the dimension of the model equations can be reduced from three to two, and the simulation run-time is dramatically cut.

On one hand, the work done for this thesis is motivated by the goal of improving the predictive success of simulations, so that they can be more aptly used to make engineering decisions. On the other hand, they are also motivated by the fact that there is valuable scientific knowledge to be gained from simulations of gas storage. The wells in Stenlille provide high precision, high frequency data, gathered over the course of several years, where the cycle of injection and production has been repeated several times. This data can be used to experiment with and validate models of fluid flow in porous media.

The future of natural gas as a fuel source, and thereby also the relevance of aquifer gas storage, is difficult to predict. In the hope of combatting climate change, large efforts are currently being made to phase out fossil fuels completely. However, this process is proving to be slow and arduous, and will likely take decades to complete [19]. Furthermore, natural gas is the "cleanest" of all fossil fuels in the sense that it emits the least CO₂ when burned. This has in several cases motivated a switch from coal to natural gas, as a first step in reducing CO₂ emissions [21].

It is also possible that natural gas will continue as an important co-star on the world energy stage, also after a shift has been made to renewable energy sources. Solar and wind generators will naturally yield variable energy outputs, and so far there are no good options for full-scale storage of this energy. Natural gas on the other hand can already be efficiently stored, and can therefore serve as a flexible back-up in the future energy system [28].

Chapter 2

Aquifer Flow Modeling

2.1 An Introduction to Aquifers

All rocks and soils will consist of a certain percentage of solid and a certain percent of void space. Dry rocks on land will have this pore space filled with air, whereas aquifers will have it filled by water. The "voidness" of a rock is specified by its **porosity**, which is defined by the relation

$$\text{porosity } \phi = \frac{\text{void volume}}{\text{total volume}}$$

As was mentioned in the introduction, the Stenlille aquifer consists mainly of sandstone, with thin layers of silt and clay dividing it into distinct layers. Sandstone has a relatively large porosity, somewhere in the 15-40 % range, which explains the large storage potential of the Stenlille aquifer.

Claystone is composed of very fine particles and has small pores, which results in lower porosity, usually in the 6-12 % range. However, we know that the silt and clay layers are semi-impermeable or fully impermeable to the gas, and that some of the folds and domes in these layers hence are capable of serving as a caprock against the gas plume. Given the non-zero porosity of the caprock, it is safe to assume that porosity can not be the only factor affecting the flow of gas through a porous medium. This brings us to the next important concept, that of **permeability**.

Permeability describes the ease with which fluid can flow through a rock. There are two factors that affect it, namely the porosity of the rock, and the specific layout of the pore networks spread through it. As an illustrative example, imagine

a rock that has large pores, but these form fluid pathways that wind and curve strongly. This rock will have high porosity, but not high permeability.

Going back to using Stenlille as an example, we will see in Chapter 4 that the claystone caprocks do indeed have lower permeability than the sandstone layers in the aquifer. However, the permeability is also non-zero in most areas, so this still does not suggest that the caprock can act as an effective seal against the gas plume. To explain why this is so, one more concept needs to be introduced.

Imagine you are adding a cube of sugar to your morning cup of coffee, and decide to carefully touch a corner of it to the surface of the hot fluid. If you watch carefully, you will see the fluid being drawn up and into the sugar cube. This motion is driven by the stronger attraction the surface of the sugar cube pores have to water compared to the attraction they have to air. The same mechanism is at work in aquifers; water tends to be more strongly attracted to the rock than gas is. The fluid with the strongest attraction is therefore commonly called the **wetting fluid**, while the one with the weakest attraction is denoted the **non-wetting fluid**.

With these concepts defined, the caprock can now be seen as a low porosity and low permeability rock, where the gas pressing upwards from below is unable to penetrate because of the caprock's preference for water. Keeping the mental image of how fluids flow through the pore space in rocks, we are now ready to move on to give an introduction to the mathematical model of fluid flow in porous media.

2.2 Porous Media Flow

2.2.1 Law of Mass Conservation

Aquifer modelling is characterized by vast differences in scale. Our discussion of the concepts of porosity, permeability and wetting/non-wetting fluids was based on considerations of how fluid flows through the pore networks in a rock. While a specific rock pore might be a few micrometers thick, the aquifers we seek to model often span a square kilometer or more in the horizontal plane. Considering individual pore networks in the aquifer would therefore quickly prove to be an exercise in futility. Instead, we shift focus to continuum mechanics, where the model is based instead on representative elementary volumes (REVs) of the rock.

A REV is often defined to be the smallest volume over which measurements will yield values that are representative of the whole. E.g., lab measurements of the porosity of a rock would oscillate strongly with small samples. These oscillations

would dampen out as the sample size became larger and larger, and a representative elementary volume could be defined when the measurements gave consistent readings.

The law of mass conservation is based on the fact that for any REV we must have that

$$\left\{ \begin{array}{l} \text{Change of} \\ \text{mass in the RVE} \end{array} \right\} = \left\{ \begin{array}{l} \text{Mass flowing} \\ \text{into the RVE} \end{array} \right\} - \left\{ \begin{array}{l} \text{Mass flowing} \\ \text{out of the RVE} \end{array} \right\} + \left\{ \begin{array}{l} \text{Mass created} \\ \text{in the RVE} \end{array} \right\}.$$

If we denote the region contained in the REV by Ω , and the boundary of Ω as $\partial\Omega$, this is expressed mathematically by the relation

$$\frac{d}{dt} \int_{\Omega} \rho\phi \, dV + \int_{\partial\Omega} \rho\phi\vec{u} \cdot \vec{n} \, dA = \int_{\Omega} q \, dV, \quad (2.1)$$

where m is the mass of the substance in the region Ω , \vec{u} is the fluid velocity, ρ is the density of the substance, q is a source or sink, and \vec{n} is the unit vector normal to the surface $\partial\Omega$ in the outward direction. The two terms on the left hand side are multiplied with ϕ because the fluids only reside in the pore space.

If the functions involved are bounded and sufficiently smooth, the derivative can be transferred into the integral,

$$\int_{\Omega} \frac{\partial(\phi\rho)}{\partial t} \, dV + \int_{\partial\Omega} \rho\phi\vec{u} \cdot \vec{n} \, dA = \int_{\Omega} q \, dV \quad (2.2)$$

and by using the divergence theorem this simplifies to

$$\int_{\Omega} \left[\frac{\partial(\rho\phi)}{\partial t} + \nabla \cdot (\rho\phi\vec{u}) - q \right] \, dV = 0. \quad (2.3)$$

Since this equation holds for every sufficiently smooth and bounded region Ω , the integral content must necessarily be zero almost everywhere. The conservation of fluid mass in porous rocks can then be succinctly stated by the relation

$$\frac{\partial}{\partial t}(\rho\phi) + \nabla \cdot (\rho\vec{v}) = q, \quad (2.4)$$

where \vec{v} is the Darcy velocity, defined as $\vec{v} = \vec{u}/\phi$. The term velocity is perhaps a bit misleading as \vec{v} is actually a weighted flux, where the division by ϕ is motivated by the fact that fluid only flows through the pore space. The term fluid velocity can instead be interpreted to mean the velocity at which a tracker would be carried along with the flow.

2.2.2 Darcy's Law

The next step in developing the aquifer model is to find an expression for the Darcy flow \vec{v} . Given the name, it will likely not come as a surprise that it was the French engineer Henry Darcy who first succeeded in doing so, based on his experimental observations of water flowing through sand. Today, it is usually stated on the form

$$\vec{v} = \frac{\mathbf{k}}{\mu}(\nabla p - \rho g \nabla z), \quad (2.5)$$

where \mathbf{k} denotes the rock permeability, μ the viscosity of the fluid and p pressure. If the rock is anisotropic, \mathbf{k} will be a tensor. Pressure can be applied in three directions, resulting in a three dimensional permeability, and the tensor is therefore realized using a 3×3 matrix.

This relation can be interpreted as describing the forces that act on the fluid, and how the result in fluid flux. The first term shows that fluid flux can be driven by pressure differences across the aquifer, and the second that it can be driven by gravitational forces acting on it. The flux is impeded by the \mathbf{k}/μ term, which accounts for "friction" in the system, where friction is interpreted as the resistance the fluid and rock put up against motion.

2.2.3 Two-Phase Flow in Porous Media

Two different fluids, namely brine and natural gas, reside in the Stenlille aquifer. Since Darcy's law only takes into account one fluid phase, a natural next step is now to extend it to take into account two-phase flow.

Before stating the extended Darcy's law, it is necessary to define some new concepts. The first is the concept of saturation s_α , defined to be the fraction of total pore space occupied by fluid phase α . The second is that of relative permeability; if pore space is already occupied by one fluid, it will inhibit the other fluid from flowing through that pore space. This reduction is accounted for using the relative permeability $k_{r,\alpha}$, which will naturally lie in the range between 0 and 1.

The multiphase version of Darcy's law is now ready to be stated. It is given by the relation

$$\vec{v}_\alpha = -\mathbf{k} \lambda_\alpha (\nabla p_\alpha - \rho_\alpha \vec{g}), \quad \alpha = \{n, w\}, \quad (2.6)$$

where λ_α is the phase mobility, defined as

$$\lambda_\alpha = \frac{k_{r,\alpha}}{\mu_\alpha}, \quad (2.7)$$

and the subscripts n and w denote the non-wetting and wetting fluids respectively.

Next, we need to consider that capillary forces between the two fluid phases can support non-zero stresses. This allows the two fluids to coexist at different pressures. The difference in phase pressures is given by the capillary pressure, i.e.

$$p_{\text{cap}} := p_n - p_w. \quad (2.8)$$

Lastly, the saturations must of course add up to one, so we have the relation

$$s_w + s_g = 1. \quad (2.9)$$

Now, using mass conservation given for each phase exactly as before, i.e.,

$$\frac{\partial}{\partial t}(\phi \rho_\alpha s_\alpha) + \nabla \cdot (\rho_\alpha \vec{v}_\alpha) = 0, \quad \alpha = \{n, w\} \quad (2.10)$$

we have a complete set of equations governing two-phase flow in porous media.

However, a quick inventory shows that the four equations (2.6),(2.8), (2.9) and (2.10) have introduced a total of ten unknown variables. The final step will therefore be to introduce the relations necessary to close the system. This is commonly done by choosing the saturations s_α and phase pressures p_α as the primary variables, and assuming that the capillary pressure p_{cap} , densities ρ_α , viscosities μ_α and relative permeability functions $k_{r,\alpha}$ are all known functions of either one of these.

More specifically, we will close the system by introducing the relations $\rho_\alpha = \rho_\alpha(p_\alpha)$, $p_{\text{cap}} = p_{\text{cap}}(s_w)$ and $k_{r,\alpha} = k_{r,\alpha}(s_\alpha)$. A short discussion of each of these is now in its place, as they touch on several important concepts that we have not yet discussed.


The first relation is simply a statement of fluid compressibility, where only density changes due to pressure have been taken into account, and the effect of temperature ignored. In many cases where the density changes are relatively small, it is sufficient to use the fluid compressibility factor

$$c_f = \frac{1}{\rho} \frac{d\rho}{dp}, \quad (2.11)$$

which may easily be obtained from laboratory experiments.

The rock volume similarly tends to vary with pressure, which motivates the use of an analogous rock compressibility factor

$$c_r = \frac{1}{\phi} \frac{d\phi}{dp}. \quad (2.12)$$



(a) An idealized example of relative permeability curves, using a Corey correlation with power 1, where the residual saturation of both fluids has been assumed to equal zero.

(b) Relative permeability curves from a more realistic scenario where residual saturation is taken into account. The thick and dotted curves apply to drainage and imbibition scenarios respectively.

Figure 2.1: Relative permeability curves, for an idealized case and for a more realistic case.

The equations (2.11) and (2.12) can easily be solved to yield the relations

$$\rho(p) = \rho_0 e^{c_f(p-p_{\text{ref}})} \quad (2.13)$$

$$V_{\text{pore}}(p) = V_0 e^{c_r(p-p_{\text{ref}})} \quad (2.14)$$

where p_{ref} is a reference pressure the compressibility factor is determined for, and V_0 and ρ_0 are the volume and density values at the same reference pressure.

The relation $k_{r,\alpha} = k_{r,\alpha}(s_\alpha)$ is commonly determined using a *relative permeability curve* for the fluid. These curves can be obtained from analytic expressions, or extrapolated from the results of laboratory measurement on core samples from the aquifer. Most analytic expressions have similarly been constructed using simplified models obtained from experimental data.

Figure 2.1a shows a typical example of a curve obtained from the first method, using the simplest form of the Corey correlation, which is a power law in the water saturation. Figure 2.1b shows a more realistic example that has been constructed using core samples from Stenlille. By comparing the two, we can observe several interesting discrepancies between them, both in the shape of each curve, and in their endpoints. To understand why the endpoints of an idealized curve and a realistic curve can differ, we need first to define the concept of *residual saturation*. If one were to observe gas being injected into an aquifer at pore-scale, one would observe gas intruding into pore networks and displacing parts of the resident brine.

At a certain saturation, the residual water will exist as small immobile droplets trapped in the rock. This leads to the concept of a residual saturation, which is defined as the highest saturation value at which the fluid is immobile.

Figure 2.1a uses a Corey correlation where the residual saturation of each fluid is assumed zero, which is why the relative permeability start or end in either (0,1) or (1,0). We can easily use the endpoints in Figure 2.1b to note the residual saturations for this case, but first we should explain why this figure contains two sets of curves, one dotted and one thick.

The thick curves are for the case of *imbibition*, where the wetting fluid displaces the non-wetting fluid. The opposite case, that of *drainage*, takes place when the wetting fluid is being displaced by the non-wetting fluid. The two cases are governed by different relationships between saturation and relative permeability, where the difference is caused by the different affinity that each fluid has for the rock.

Now that we have a clear understanding what case each of the curves in Figure 2.1b represents, we can note the residual saturations for each case. In case of imbibition, we see that wetting phase has a residual of 0.25, and the non-wetting phase a residual saturation of $1-0.7=0.3$. In the case of drainage, the same values are given by 0.25 and 0 respectively.

The last closing relation, $p_{\text{cap}} = p_{\text{cap}}(s_w)$, gives the relationship between the capillary pressure and saturation. Due to surface tension, there will generally be a pressure difference between two fluid phases in contact with each other. The exact relation is specified similarly as to that of relative permeability, this time using capillary pressure curves.

2.2.4 The Black Oil Model

The standard model used by the petroleum industry for oil and gas simulations is called the black-oil model. In essence, it describes conservation laws for a three-phase, multicomponent system, where dissolution of gas into oil is also taken into consideration. The model is given by the equations

$$\partial_t(\phi b_\alpha \rho_\alpha^s s_\alpha) + \nabla \cdot (b_\alpha \rho_\alpha^s \vec{v}_\alpha) = q_\alpha, \quad \alpha = \{o, w\} \quad (2.15)$$

$$\partial_t(\phi(b_g S_g - b_o r_{so} S_o)) + \nabla \cdot (b_g \vec{v}_g + b_o r_{so} \vec{v}_o) - (b_g q_g + b_o r_{so}) q_o = q_o, \quad (2.16)$$

where the subscript *o* refers to the oleic phase, *g* to the gas phase and *w* to the water phase. The density variable $\rho(p)$ used in (2.10) has been abandoned in

favour of using a constant surface density ρ^s for each fluids along with an reciprocal formation volume factor $b = b(p)$. The inverse formation volume factor is simply given by $b = V^s/V(p)$, and a table relating it with different pressure values can be found easily with laboratory experiments.

The variable $r_{so} = V_g^s/V_o^s$ relates the volume of gas dissolved into oil with the volume of oil at surface conditions, and is used to take into account that gas can dissolve into oil. For the case of natural gas stored in an aquifer, only the two phases $\alpha = \{g, w\}$ will be needed, and the oleic phase ignored completely. The equations then simplify to essentially the same set of transport equations as the ones we have already introduced, i.e.,

$$\partial_t(\phi b_w \rho_w^s s_w) + \nabla \cdot (\rho_w \vec{v}_w) = q_w, \quad (2.17)$$

$$\partial_t(\phi b_g \rho_g^s s_g) + \nabla \cdot (\rho_n \vec{v}_g) = q_g, \quad (2.18)$$

except that the density variable has been exchanged in favor of a reference density and the formation value factor.

2.2.5 Model Reformulations

This section is devoted to showing two manipulations of the set of equations (2.6)-(2.9). Both of these rewrites are commonly referred to as the *fractional flow formulation* in phase pressures, one for a simplified incompressible case, and one for the more complicated compressible case.

The first rewrite will be useful for later discussions of the vertical equilibrium model, where we will perform a similar rewrite and make use of the similarities between the two. The second rewrite, for the compressible case, serves the purpose of clarifying the nature of the system (2.6)-(2.9). The observations made will come in handy in later discussions of the numerical solution, which will exploit the nature of the system to develop faster solution methods.

Incompressible case:

One starts by defining total velocity $\vec{v} = \vec{v}_w + \vec{v}_n$, total source term $q = q_w + q_n$ and total mobility $\lambda = \lambda_w + \lambda_n$. Adding together (2.10) for $\alpha = \{w, n\}$ and assuming incompressibility yields an elliptic Poisson-type equation

$$\nabla \cdot \vec{v} = q, \quad (2.19)$$

where the total velocity is given by

$$\vec{v} = -(\lambda \mathbf{k}) \nabla p + \lambda_w \nabla p_c + (\lambda_n \rho_n + \lambda_w \rho_w) \vec{g} \nabla z. \quad (2.20)$$

The time derivative vanished since the saturations of each phase add up to one, and incompressibility implies that the fluid densities are constant.

Next, by using Darcy's law for $\lambda_n \vec{v}_w - \lambda_w \vec{v}_n$ and simplifying, we have the relation

$$\lambda_n \vec{v}_w - \lambda_w \vec{v}_n = \lambda_w \lambda_n \mathbf{k} [\nabla p_c + \Delta \rho g \nabla z], \quad (2.21)$$

where $\Delta \rho = \rho_w - \rho_n$. Solving for \vec{v}_w and inserting this back into (2.10) yields the so-called saturation equation

$$\phi \frac{\partial s_w}{\partial t} + \nabla \cdot [f_w (\vec{v} + \lambda_n \Delta \rho g \nabla z + \lambda_n P'_{\text{cap}} \nabla s_w)] = q_z, \quad (2.22)$$

where f_w is given by $f_w = \lambda_w / (\lambda_n + \lambda_w)$ and $P'_{\text{cap}} = \nabla (p_{\text{cap}} / s_w)$.

Compressible case: The first step is to expand the accumulation term in (2.10) by using the chain rule, and then rewriting using the expression for fluid compressibility (2.11), which yields

$$\frac{\partial}{\partial t} (\phi \rho_\alpha S_\alpha) = \rho_\alpha S_\alpha \frac{\partial \phi}{\partial t} + \phi S_\alpha \frac{\partial \rho_\alpha}{\partial t} + \phi \rho_\alpha \frac{\partial S_\alpha}{\partial t} \quad (2.23)$$

$$= \rho_\alpha S_\alpha \frac{\partial \phi}{\partial t} + \phi S_\alpha \frac{d\rho_\alpha}{dp_\alpha} \frac{\partial p_\alpha}{\partial t} + \phi \rho_\alpha \frac{\partial S_\alpha}{\partial t} \quad (2.24)$$

$$= \rho_\alpha S_\alpha \frac{\partial \phi}{\partial t} + \phi S_\alpha c_\alpha \frac{\partial \rho_\alpha}{\partial t} + \phi \rho_\alpha \frac{\partial S_\alpha}{\partial t}. \quad (2.25)$$

By inserting this into (2.10), then summing the equations for $\alpha = \{n, w\}$, and then simplifying, we get

$$\frac{\partial \phi}{\partial t} + \phi c_n S_n \frac{\partial p_n}{\partial t} + \phi c_w S_w \frac{\partial p_w}{\partial t} + \frac{1}{\rho_n} \nabla \cdot (\rho_n \vec{v}_n) + \frac{1}{\rho_w} \nabla \cdot (\rho_w \vec{v}_w) = q_n + q_w \quad (2.26)$$

where some of the accumulation terms have disappeared or been simplified by using the relation $s_n + s_w = 1$.

For aquifer simulations in general, including simulations of Stenlille, the capillary pressure will be small compared to the reservoir pressure. The dubious reader is invited to compare the capillary pressure shown in Figure 4.9 and the initial conditions shown in Chapter 4.1.4 to see that this is in fact the case. When examining the nature of the pressure equation, it is therefore natural to ignore capillary forces and simply set $p = p_n = p_w$.

By using the chain rule $\nabla(\rho_\alpha \vec{v}_\alpha) = \rho_\alpha \nabla \vec{v}_\alpha + \vec{v}_\alpha \nabla \rho_\alpha$, and the fact that the spatial density variations will normally be small, we can make the additional assumption that $\nabla \rho = 0$. Equation (2.26) can then be simplified to

$$\phi c \frac{\partial p}{\partial t} + \nabla \cdot (\mathbf{K} \lambda \nabla p - \rho \vec{g}) = q \quad (2.27)$$

where λ is the total mobility $\lambda = \lambda_w + \lambda_n$, q is the total source term $q = q_w + q_n$, ρ is the total density $\rho = \rho_n + \rho_w$, and the rock compressibility from (2.12) has been used when defining the total compressibility $c = (c_r + c_n S_n + c_w S_w)$.

Furthermore, the spatial density variations will normally be small, in which case we can set $\nabla \rho_\alpha = 0$. This then simplifies to

$$\phi c \frac{\partial p}{\partial t} - \nabla \cdot (\mathbf{K} \nabla p) = \hat{q}, \quad (2.28)$$

\hat{q} is a combined source term that accounts for both volumetric source terms and pressure variations with depth.

The equation (2.28) is parabolic. However, compressibility tends to decrease with increasing pressure. At the conditions found in a typical aquifer, it will therefore have a strong elliptic character.

2.3 A Vertical Equilibrium Model for Aquifer Flow

Long simulation run-times are a common issue when working with aquifers, caused by the large number of grid cells and time steps needed to resolve the effect of great differences in spatial and temporal scales. The horizontal extent of a reservoir is often somewhere in the 0.1-1 km² range, while an aquifer can potentially extend several kilometres in the same directions. The properties of the REV's however, arise from considering physical processes that occur on a scale of micrometers. The result of this duality is a challenging multiscale problem that requires both clever upscaling methods and a large amount of grid cells in the discretization.

The assumption of vertical equilibrium (VE), i.e., no vertical flow makes it possible to reduce the three-dimensional black-oil equations to a set of two-dimensional equations, which drastically cuts simulation run times. The dimension reduction is done by integrating the set of equations (2.10) and (2.6), and then using the VE assumption to reformulate the governing equations into their two-dimensional version. This type of model has been successfully used to model long-term storage of CO₂ in aquifers [19], for which the three-dimensional black-oil model is rendered an intractable option by the large time-scales involved.

Stenlille has the potential to be a good candidate for a VE model, as the horizontal scale of the aquifer is large compared to the vertical scale, and the density difference between the two fluid phases is quite large. Due to the differences in the horizontal and vertical scale, the vertical flows will in many cases be negligible compared

to the horizontal flows, which motivates the simplifying assumption of vertical equilibrium. The density of natural gas at 1500 m depth will be about 115 kg/m^3 , which is small compared to the density of brine at around 1100 kg/m^3 . Buoyancy forces will therefore rapidly cause natural gas and brine to separate into two layers, separated by a thin transition zone. If the interface is thin enough, it can be approximated well by a *sharp interface* between the two fluids.

2.3.1 Deriving the Vertically Integrated Model




Figure 2.2: The coordinate system used when setting up the vertical-equilibrium formulation of the model. The interface between brine and CO_2 is given by $h(x, y)$. The region between $h(x, y)$ and $h_{\max}(x, y)$ contains residually trapped CO_2 , that left behind when the CO_2 plume travelled through it. Figure taken from [15].

The derivation given here will follow roughly the same steps as the derivation given in [15].

Assume the spatial coordinates of an aquifer are known. The aquifer might be tilted at an angle, so we define a coordinate system (x, y, z) that is aligned so that the plane spanned by x and y is roughly tangential with the main flow direction. The aquifer is described by a top surface $Z(x, y)$, and the height in the z direction given by $H(x, y)$.

We start by simply integrating (2.10) vertically, from the top to the bottom of the aquifer, which yields the relation

$$\int_H^{Z+H} \frac{\partial}{\partial t} (\phi b_\alpha s_\alpha) dz + \int_H^{Z+H} \nabla \cdot b_\alpha \vec{v}_\alpha dz = \int_H^{Z+H} q_\alpha dz. \quad (2.29)$$

If ϕs_α is sufficiently smooth, the time derivative and the integral on the left hand side can be interchanged. The spatial derivatives require using Leibnitz's rule, which yields

$$\frac{\partial}{\partial t} \int_H^{Z+H} \phi b_\alpha s_\alpha dz + \nabla_{\parallel} \cdot \int_H^{Z+H} b_\alpha \vec{v}_\alpha dz = \Upsilon, \quad (2.30)$$

where the new source term Υ on the right hand side is given by

$$\Upsilon = \int_H^{Z+H} q_\alpha dz. \quad (2.31)$$

The "parallel to" subscript signifies a reduction to the plane spanned by x and y , so for example $\nabla_{\parallel} = \frac{\partial}{\partial x} \vec{e}_x + \frac{\partial}{\partial y} \vec{e}_y$. Any terms multiplied by v_z disappeared since the vertical flow was assumed to be zero.

The next step is to utilize the VE and sharp interface assumption in computing these integrals, in order to put them in a form that can be easily evaluated at each time step. The compressibility term b_α poses some difficulty, in that it significantly complicates the computation of the integrals in (2.30).

For a full derivation of the VE equations with variable density, the reader is referred to [18]. This article examined the effect compressibility had on the results CO₂ storage simulations was examined, and concluded that it was small in most cases. It further found that the semi-incompressible models, i.e., models where only lateral variations in density are allowed, were sufficient to capture most of the effects from compressibility. This is partly due to the small height of the a typical CO₂ plume, which makes the density differences over each column quite small. Another factor is that temperature increases with depth below the earth surface, and increasing temperature decreases density, thus offsetting some of density increase that is caused by an increase in pressure.

It is assumed that the same holds true for Stenlille, i.e. that a simplified semi-incompressible model will capture the most prominent effects of compressibility. This was justified by noting that the density differences in a column of gas in Stenlille were in fact small, usually less than one percent. Furthermore, the effects of temperature on the density of methane were found to be smaller than the effect on CO₂.

Using now the assumption of a sharp interface between natural gas and water, the aquifer will be cleanly split into three sections as shown in Figure 2.2, given by:

- The natural gas plume between Z and $Z + h$, where the brine saturation is equal to the residual brine saturation $s_{w,r}$, the natural gas saturation is given by $1 - s_{w,r}$ and the natural gas mobility is given by $\lambda_{n,e} = \lambda_n(1 - s_{w,r})$.
- The residual natural gas plume between $Z + h$ and $Z + h_{\max}$, where the natural gas saturation is equal to the residual natural gas saturation $s_{n,r}$, the brine saturation is given by $1 - s_{n,r}$ and the brine mobility is given by $\lambda_{w,e} = \lambda_n(1 - s_{n,r})$. The natural gas mobility is of course equal to zero since there are only residual drops left in this region.
- The region between $Z + h_{\max}$ and $Z + H$ completely filled with brine.

Armed with this knowledge it is now possible to find an easily evaluable expression for each term in Equation (2.30). Using the assumption of semi-incompressibility, we will have $b_\alpha = b_\alpha(x, y)$. The difference in porosity will similarly be ignored in the vertical direction, so that $\phi = \phi(x, y)$. For the first term it is then easy to see that

$$\frac{\partial}{\partial t} \Phi_n = \frac{\partial}{\partial t} \int_H^{Z+H} \phi b_n s_n dz = \frac{\partial}{\partial t} \phi b_n [h(1 - s_{w,r}) + (h_{\max} - h)s_{n,r}], \quad (2.32)$$

$$\frac{\partial}{\partial t} \Phi_w = \frac{\partial}{\partial t} \int_H^{Z+H} \phi b_w s_w dz = \frac{\partial}{\partial t} \phi b_w [s_{w,r}h + (h_{\max} - h)(1 - s_{n,r}) + (H - h_{\max}) \cdot 1]. \quad (2.33)$$

The vertical equilibrium assumption makes it possible to calculate the pressure at any given point. If the pressure at the top surface $Z(x_1, x_2)$ is denoted P_Z , the pressure will be given by

$$p(z) = \begin{cases} P_Z + \rho_n g_z (z - Z), & \text{if } Z \leq z \leq Z + h, \\ P_Z + \rho_n g_z h + \rho_w g_z (z - Z - h), & \text{if } Z + h \leq z \leq Z + H. \end{cases} \quad (2.34)$$

V_n and V_w are defined to be the vertically averaged Darcy velocities, i.e.

$$V_\alpha = \int_H^{Z+H} \vec{v}_\alpha dz, \quad \alpha = \{w, n\}. \quad (2.35)$$

Integrating Darcy's law vertically yields

$$\int_H^{Z+H} \vec{v}_\alpha dz = \int_H^{Z+H} -\mathbf{k} \lambda_\alpha (\nabla p - \rho_\alpha \vec{g}) dz, \quad \alpha = \{w, n\}. \quad (2.36)$$

In order to formulate Darcy's law in two dimensions, the quantities

$$\mathbf{K} = \frac{1}{H} \int_H^{Z+H} \mathbf{k}_{\parallel} dz, \quad (2.37)$$

$$\Lambda_n(h) = \frac{1}{H} \left[\int_Z^{Z+h} \lambda_{n,e} \mathbf{k}_{\parallel} dz \right] \mathbf{K}^{-1}, \quad (2.38)$$

$$\Lambda_w(h, h_{\max}) = \frac{1}{H} \left[\int_{Z+h}^{Z+h_{\max}} \lambda_{w,e} \mathbf{k}_{\parallel} dz + \int_{Z+h_{\max}}^{Z+H} \lambda_w(1) \mathbf{k}_{\parallel} dz \right] \mathbf{K}^{-1}, \quad (2.39)$$

are defined. This yields expressions for the vertically integrated Darcy velocities given by

$$\vec{V}_n = -H \Lambda_n \mathbf{K} (\nabla_{\parallel} (P_Z - \rho_n g_z Z) - \rho_n \vec{g}_{\parallel}), \quad (2.40)$$

$$\vec{V}_w = -H \Lambda_w \mathbf{K} (\nabla_{\parallel} (P_Z - \rho_w g_z Z) - g_z \Delta \rho \nabla_{\parallel} h - \rho_w \vec{g}_{\parallel}), \quad (2.41)$$

where $\Delta \rho = \rho_w - \rho_n$. If ϕ and \mathbf{K} are constant in the z-direction, the upscaled mobilities can be calculated as

$$\Lambda_n(h) = h \lambda_{n,e}, \quad (2.42)$$

$$\Lambda_w(h, h_{\max}) = (h_{\max} - h) \lambda_{w,e} + (H - h_{\max}) \lambda_w(1). \quad (2.43)$$

The final upscaled equation is on the form

$$\frac{\partial}{\partial t} \Phi_{\alpha} + \nabla_{\parallel} \cdot V_{\alpha} = \Upsilon_{\alpha}. \quad (2.44)$$

2.3.2 Rewriting Into a Fractional Flow Formulation

Rewriting the VE model equations into the fractional flow formulation in phase pressure will allow us to see some similarities they bear with the standard black-oil model equations. To this end, first define the upscaled total Darcy flux $\vec{V} = \vec{V}_n + \vec{V}_w$. Adding together (2.44) for $\alpha = n$ and $\alpha = w$ yields

$$\nabla_{\parallel} \cdot \vec{V} = Q, \quad (2.45)$$

$$\vec{V} = -H \Lambda \mathbf{K} \left[\nabla_{\parallel} P_z - (\rho_n \mathbf{F}_n + \rho_w \mathbf{F}_w) (\vec{g}_{\parallel} + g_z \nabla_{\parallel} Z) - \mathbf{F}_w g_z \Delta \rho \nabla_{\parallel} h \right], \quad (2.46)$$

where $\Lambda(h, h_{\max}) = \Lambda_w(h, h_{\max}) + \Lambda_n(h)$ and $\mathbf{F}_\alpha(h, h_{\max}) = \Lambda_\alpha \Lambda^{-1}$. The time derivative vanished since $s_n + s_w = 1$ always.

The transport equation is formed by using (2.44) again for $\alpha = n$ and manipulating the second term to be a function of \vec{V} . This yields

$$\frac{\partial}{\partial t} \Phi + \nabla_{\parallel} \cdot [\mathbf{F}_n \vec{V} + \Delta \rho \mathbf{K} \Lambda_w \mathbf{F}_n (g_{\parallel} + g_z \nabla_{\parallel} (Z + h))] = Q, \quad (2.47)$$

where

$$\Phi(h) = \phi(1 - s_{w,r})h + \phi(1 - s_{n,r})(h_{\max} - h), \quad (2.48)$$

$$\Lambda_n(h) = h \lambda_{n,e}, \quad (2.49)$$

$$\Lambda_w(h, h_{\max}) = (h - h_{\max}) \lambda_{w,e} + (H - h_{\max}) \lambda_w(1). \quad (2.50)$$

This is often called the h -formulation of the model.

One can further derive the so called S -formulation by introducing the fractional content $S = h/H$ of gas in a column. S then serves the role of saturation. The fractional content of gas in a column is given by

$$S(h, h_{\max}) = \frac{h(1 - s_{w,r}) + (h_{\max} - h)s_{n,r}}{H}, \quad (2.51)$$

$$S_{\max}(h_{\max}) = \frac{h_{\max}(1 - s_{w,r})}{H}. \quad (2.52)$$

A simple inversion yields

$$h(S, S_{\max}) = H \frac{S(1 - s_{w,r}) - S_{\max} s_{n,r}}{(1 - s_{w,r})(1 - s_{w,r} - s_{n,r})}, \quad (2.53)$$

$$h_{\max} = \frac{HS_{\max}}{(1 - s_{w,r})}. \quad (2.54)$$

Finally, shifting the coordinate system so that $Z = 0$ lets one evaluate the integrals describing the mobility, i.e.

$$\Lambda_n(S, S_{\max}) = \frac{1}{H} \left[\int_0^h \lambda_{n,e} k_{\parallel} dz \right] \mathbf{K}^{-1} = \frac{h(S, S_{\max})}{H} \lambda_{n,e}, \quad (2.55)$$

$$\Lambda_w(S, S_{\max}) = \frac{1}{H} \left[\int_h^{h_{\max}} \lambda_{w,e} k_{\parallel} dz + \int_H^{h_{\max}} \lambda_w(1) k_{\parallel} dz \right] \mathbf{K}^{-1}, \quad (2.56)$$

$$= \frac{h_{\max}(S_{\max}) - h(S, S_{\max})}{H} \lambda_{w,e} + \frac{H - h_{\max}(S, S_{\max})}{H} \lambda_w(1). \quad (2.57)$$

The equations (2.45) and (2.47) are now formally similar to the standard flow formulation as stated in (2.19) and (2.22) when S plays the role of saturation. The only difference between them is the formulation of the mobility function and capillary pressure function.

Chapter 3

Numerical Methods and Implementation

The previous chapter we gave an introduction to aquifer flow modelling, and developed two different models for two-phase flow in porous media. With this concluded, the question of finding a solution follows naturally - how can we best solve the model equations numerically, and how should we implement the solution method? This chapter is devoted to answering those questions.

The first section in this chapter will introduce the discretization of the model equations, which uses a standard two point finite-volume approximation. The next section then describes the Newton-Raphson method used to solve the system of discretized equations, and gives a short introduction to the concept of automatic differentiation, which is used for the construction of the Jacobian of the system. This is then followed by an introduction to well equations used to model the relationship between well pressure and well rate. Finally, an overview is given of the solution strategy used for solving the final system of linear equations.

The final section will then be spent introducing the MATLAB Reservoir Simulation Toolbox (MRST) [2], which serves as the framework where the numerical techniques and methods are implemented.

3.1 Discretization

The model equations will be discretized using finite-volume methods. In finite volume methods, the partial differential equations are integrated, and any terms containing a divergence operator are converted to surface integrals by use of the

divergence theorem. This integration has the advantage of making the numerical schemes conservative, in that the physical quantities are conserved over the cell volumes. Furthermore, it shifts the focus to the cell averages of each quantity, which removes most problems with discontinuity [14].

It will be convenient to start simple, in order to motivate and define some of the central concepts, before stating the full discretization of the model equations.

In the case of a one-phase, incompressible fluid model, the model equations will be given by

$$\nabla \cdot \vec{v} = q, \quad (3.1)$$

$$\vec{v} = -\mathbf{k}\nabla p, \quad (3.2)$$

which are assumed to hold on the domain $\Omega \subset \mathbb{R}^d$. We start by simply integrating this equation over a cell i , and manipulating the left hand side using the divergence theorem, which yields

$$\int_{\partial\Omega_i} \vec{v} \cdot \vec{n} ds = \int_{\Omega_i} q d\vec{x}. \quad (3.3)$$

The left hand side can then be split up and evaluated on each edge of the cell individually. Let us denote the flux across the face shared by the cells i and k , in the direction from cell i to cell k , as $v_{i,k}$. Using the midpoint rule, the flux across this face can be approximated by

$$v_{i,k} = \int_{\Gamma_{ik}} \vec{v} \cdot \vec{n} ds \approx A_{i,k} \vec{v}(\vec{x}_{i,k}) \cdot \vec{n}_{i,k}, \quad (3.4)$$

where Γ_{ik} is the face between the two cells, $\vec{x}_{i,k}$ this face's centroid, $\vec{n}_{i,k}$ its normal vector and $A_{i,k}$ its area.

This expression can be further evaluated using Darcy's law, $\vec{v} = -\mathbf{k}\nabla p$. The next step is now to express the gradient of p , by using a one-sided finite difference between the pressure at the centroid of the cell and the pressure at the centroid of the face. To do this, we make the simplifying assumption that the pressure is linear in each cell. In this case, the average pressure will equal the pressure at the centroid, and we will denote both by p_i . Using the notation shown in the figure, the flux $v_{i,k}$ is now given by

$$v_{i,k} = A_{i,k} \mathbf{k}_i \frac{(p_i - \pi_{i,k}) \vec{c}_{i,k}}{|\vec{c}_{i,k}|^2} \cdot \vec{n}_{i,k} = T_{i,k} (p_i - \pi_{i,k}), \quad (3.5)$$

Here we have introduced the direction-specific transmissibility $T_{i,k}$. Next, we will define the (direction-independent) transmissibility T_{ik} , which is interpreted as the

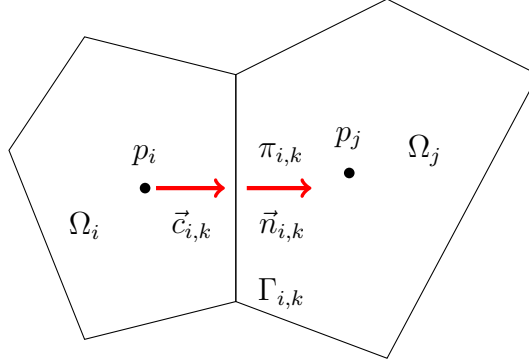


Figure 3.1: An illustration of two cells used to define the TPFA of the Laplace operator.

effective communication between to cells, and given as the harmonic average of the direction-specific transmissibilities, i.e.,

$$T_{ik} = [T_{i,k}^{-1} + T_{k,i}^{-1}]^{-1}. \quad (3.6)$$

Using then that the fluxes across the face must be continuous, so that $\vec{v}_{i,k} = -\vec{v}_{k,i}$, and that the face pressures must be equal, $\pi_{i,k} = \pi_{k,i} = \pi_{ik}$, we finally have

$$\vec{v}_{ik} = T_{ik}(p_i - p_k). \quad (3.7)$$

Inserting this back into (3.3), we get the two-point flux approximation (TPFA) scheme

$$\sum_k T_{ik}(p_i - p_k) = q_i, \quad \forall \Omega_i \subset \Omega. \quad (3.8)$$

Next, let us introduce the concepts of discrete flux and divergence operators. Writing a variable in bold font will now be used to denote values per cell or flux, and a square bracket is used to evaluate it on a specific cell or face.

In this spirit, let $\mathbf{v} \in \mathbb{R}^{n_f}$ denote a discrete flux and $\mathbf{v}[f]$ its restriction onto face f with orientation from $N_1(f)$ to $N_2(f)$. The divergence operator, denoted DIV , maps from faces to cells [16]. The divergence of a flux \mathbf{v} from cell c is given by

$$\text{DIV}(\mathbf{v})[c] = \sum_{f \in F(c)} \mathbf{v}[f] \mathbf{1}_{c=N_1(f)} - \sum_{f \in F(c)} \mathbf{v}[f] \mathbf{1}_{c=N_2(f)}, \quad (3.9)$$

where $\mathbf{1}_{\text{expr}}$ is the indicator function.

The gradient operator, similarly denoted GRAD , maps from cell pairs to faces. Restricted to one face f it is defined as

$$\text{GRAD}(\mathbf{p})[f] = \mathbf{p}[N_2(f)] - \mathbf{p}[N_1(f)], \quad (3.10)$$

for any $\mathbf{p} \in \mathbb{R}_c^n$.

The value in introducing these operators becomes apparent when we can succinctly rewrite the discretization (3.8) as

$$\text{DIV}(\mathbf{v}) = q, \quad (3.11)$$

$$\mathbf{v} = -T \text{GRAD}(\mathbf{p}). \quad (3.12)$$

Furthermore, these operators are well suited for implementation in Matlab, as they can each be represented as a sparse matrices. In the case of e.g. a variable vector \mathbf{p} with a value for each grid cell, its gradient can be found by multiplying \mathbf{p} with a matrix C . The matrix C will then be composed of ones and zeros, with a one in row i and column k if the cells i and k share a face, and zero if they do not.

Moreover, the GRAD operator is in fact the negative adjoint of the divergence operator, as

$$\int_{\Omega} p \nabla \cdot \vec{v} d\Omega = - \int_{\Omega} \vec{v} \cdot \nabla p d\Omega. \quad (3.13)$$

The proof that this holds in a discrete setting is given in [14]. Following this relation, the matrix C used to for the gradient operator will in fact be the transpose of the divergence operator C^T .

We are now ready to state the discretization of the black-oil model equations. Let us first simply recall that they are given by

$$\partial_t(\phi b_{\alpha} s_{\alpha}) + \nabla \cdot (b_{\alpha} \vec{v}_{\alpha}) = b_{\alpha} q_{\alpha} \quad , \alpha \in \{n, w\}, \quad (3.14)$$

$$\vec{v}_{\alpha} = -\mathbf{k} \lambda_{\alpha} (\nabla p_{\alpha} - \rho_{\alpha} \vec{g}) \quad , \alpha \in \{n, w\}. \quad (3.15)$$

The discretization itself is constructed using the same method as the one described for the single-phase, incompressible model. Using a first order implicit scheme, each phase $\alpha \in \{n, w\}$ is discretized as

$$\frac{(\phi(\mathbf{p}[c]) \mathbf{b}(\mathbf{p}[c]) \mathbf{s}[c])^{n+1} - (\phi(\mathbf{p}[c]) \mathbf{b}(\mathbf{p}[c]) \mathbf{s}[c])^n}{\Delta t} + \text{DIV}(\mathbf{b}(\mathbf{p}[c]) \mathbf{v}) [c]^{n+1} - (\mathbf{b}\mathbf{q}) [c]^{n+1} = 0, \quad (3.16)$$

where

$$\mathbf{v}[f] = -\text{UPW}(\boldsymbol{\lambda}) [f] \mathbf{T}[f] (\text{GRAD}(\mathbf{p} - \mathbf{p}_c^{nw})[f] - g \text{AVG}(\boldsymbol{\rho}) [f] \text{GRAD}(\mathbf{z}) [f]), \quad (3.17)$$

$\lambda = bk_r/\mu$ and $p_c^{nw} = p_w - p_n$ is the capillary pressure between the gas and water phase.

Two new discrete operators were used here that have not yet been defined, namely an averaging function and an upwind function. The first of these, $\text{AVG}(\rho)[f]$, is the face-valued density computed as an arithmetic average over the pressures in the two adjacent cells. The second of them is UPW, the upwind function, given by

$$\text{UPW}(\lambda[f]) = \begin{cases} \lambda[N_1(f)], & \text{if } \text{GRAD}(\mathbf{p} - \mathbf{p}_c^{\text{ow}})[f] - g \text{AVG}(\rho)[f] \text{GRAD}(\mathbf{z})[f] > 0, \\ \lambda[N_2(f)], & \text{otherwise.} \end{cases}$$

3.2 Newtons Method and Automatic Differentiation

The system of equations made up by (3.16) and (3.17) will be highly non-linear, for which Newton's method is often the chosen solution method.

Newton's method itself is fairly straight-forward. First, let the system of equations to be solved at each time step be denoted \mathbf{G} , so that we want to solve $\mathbf{G}(\mathbf{p}^{i+1}; \mathbf{p}^i) = 0$. Newton's method is then to iterate through

$$\frac{\partial \mathbf{G}(\mathbf{p}^i)}{\partial \mathbf{p}^i} \delta \mathbf{p}^{i+1} = -\mathbf{G}(\mathbf{p}^{i+1}), \quad \mathbf{p}^{i+1} \leftarrow \mathbf{p}^i + \delta \mathbf{p}^{i+1}, \quad (3.18)$$

for $i = 1, 2, 3, \dots$, where $\mathbf{J}(\mathbf{p}^i) = \partial \mathbf{G}(\mathbf{p}^i) / \partial \mathbf{p}^i$ is the Jacobian matrix of \mathbf{G} . The $\delta \mathbf{p}^{i+1}$ term is often referred to as the Newton update, as the new solution approximation in each time step is constructed by adding the Newton update to the approximation found in the previous time step. The iterations are stopped when the residual is found to be sufficiently small.

The Jacobian has to be constructed based on each model, and will depend on both the coupling of variables and the choice of linearizations, i.e. which quantities are evaluated at step i and which are evaluated at step $i + 1$. Therefore, if a change is made to either the model or the linearization, the Jacobian would have to be recalculated. The process of deriving and implementing the Jacobian is often a bottleneck in the simulation work flow, as it is both time-consuming and highly prone to errors.

This is where the beauty of automatic differentiation enters the stage. In MRST, automatic differentiation is implemented by the use of an ADI class that keeps track of both a variable and its derivatives [16]. Whenever an operation is applied to an ADI variable, the corresponding operation is applied to its derivative. For example, if we denote by $\langle f, f_x \rangle$ an ADI variable f and its derivative f_x , addition

and multiplication would be defined as

$$\langle f, f_x \rangle + \langle g, g_x \rangle = \langle f + g, f_x + g_x \rangle \quad (3.19)$$

$$\langle f, f_x \rangle * \langle g, g_x \rangle = \langle gf, fg_x + f_xg \rangle \quad (3.20)$$

Operator overloading is used to implement this functionality. Whenever Matlab encounters an expression on the form $f+g$, the correct operator is chosen according to the type of f and g .

3.3 Well Equations

In real life, wells are normally controlled by specifying either a surface rate or a bottom-hole pressure at which they should they be operating. Here, the bottom-hole pressure is defined to be the pressure at the bottom-most perforation made in the well, and a perforation is a hole made in the well casing, through which the well can inject or produce fluid.

A straightforward way to model the well injection and extraction of fluids from the aquifer would be to simply assign the desired well rates as a source term q in the model equations. However, the matter is complicated by the fact that the size of a typical grid cell will most often dwarf the size of the well itself. The localized pressure differences around the wells will therefore not be picked up upon, in which case it will be troublesome to assign the well pressure to equal the cell pressure, or the other way around.

This problem is often solved by utilizing a well model, where Peaceman's model is the standard choice [19]. Let us for simplicity assume that only one fluid is flowing in the immediate area around the well. Peaceman used Darcy's law for radial flow to model the source term

$$q = \frac{2\pi h \mathbf{k} b}{\mu} \frac{p_c - p_{\text{well}}}{\ln(r_o/r_{\text{well}})}, \quad (3.21)$$

where h is the formation thickness, p_{well} is the bottom-hole well pressure, r_{well} is the wellbore radius and p_c is the cell radius. The variable r_o is defined to be the radius at which the analytic steady state pressure p in the well is equal to the pressure p_c for the well cell. For anisotropic permeabilities, it is often defined as [22]

$$r_o = \frac{(\sqrt{K_y/K_x} \Delta x^2 + \sqrt{K_y/K_x} \Delta y^2)^{\frac{1}{2}}}{(K_y/K_x)^{\frac{1}{4}} + (K_y/K_x)^{\frac{1}{4}}}. \quad (3.22)$$

If we now define the well index to be

$$W_i = \frac{2\pi h}{\ln(r_o/r_w)} \quad (3.23)$$

and extend the model to two-phase flow, we have following equation to model the relationship between well pressure and well flow,

$$q_\alpha = W_i \lambda_\alpha (p_{c,\alpha} - p_{\text{well}}). \quad (3.24)$$

3.4 The Solver

Given an aquifer with two fluid phases and n_{wells} wells, the final system of discretized equations will consist $2+3n_{\text{wells}}$ equations, namely

- A water equation and a gas equation, given by (3.16),
- $2n_{\text{wells}}$ equations that for the two phases set the well surface rate to the sum of the perforation contributions,
- n_{wells} well control equations that ensures that the well operates at the prescribed rate,

where a well perforation is a hole punched in the casing of the well to connect it to the reservoir. The primary variables are chosen to be the wetting phase pressure \mathbf{p}_s , the wetting phase saturation \mathbf{s}_w , the well source terms for the wetting fluid \mathbf{q}_{Ws} , the well source terms for the non-wetting fluid \mathbf{q}_{Gs} , and the bottom hole pressures \mathbf{p}_{bh} .

Using the discretization outlined in Section 3.1, and Newtons method as proposed in Section 3.2, this results in a linear system on the form

$$Ax = b \quad (3.25)$$

that needs to be solved. The backslash operator in Matlab has been found to provide satisfactory efficiency for small systems. For larger systems however, as the one that arises when modeling Stenlille, a Constrained Pressure Residual (CPR) preconditioned Generalized Minimal Residual (GMRES) method is more efficient [16].

3.4.1 The GMRES Method

The GMRES method is an example of a Krylov subspace method, meaning that it works by looking for solution approximations that are constricted to lie in the

n th Krylov subspace of the problem. If we denote the n th Krylov subspace by K_n , which is given by

$$K_n = \text{span}\{b, Ab, A^2, \dots, A^{n-1}b\}, \quad (3.26)$$

we can more succinctly state that the GMRES method works by minimizing $r_n = Ax_n - b$ when $x_n \in K_n$.

There are several different Krylov subspace methods for solving linear equations, the difference between them lying in the way that the basis for the approximation space is built and the preconditioners used when solving the system. The vectors $b, Ab, A^2, \dots, A^{n-1}b$ cannot be used as a basis, as they may not be linearly independent.

The GMRES method uses an Arnoldi iteration to find such a basis q_1, q_2, \dots, q_n for the n th Krylov subspace. The basic idea of the algorithm is that beginning with an arbitrary vector v_0 , at each step the previous Arnoldi vector v_j is multiplied by A , before being orthonormalized against all previous basis vectors v_i 's using the Gram-Schmidt procedure. For full details the reader is referred to [23].

After a basis has been found using Arnoldi iterations, we must have that any vector in the space $x_0 + K_n$ can be written as $x = x_0 + V_m y$, where the matrix V_m is composed of the basis vectors v_i . A matrix referred to as the Hessenberg matrix, composed of certain variables from the Gram-Schmidt orthonormalization in each step, has several interesting properties. Using these it is possible to show that this y also minimizes the Euclidean norm of the residual

$$r_n = \tilde{H}_n y_n - \beta e_1, \quad (3.27)$$

where $\beta = \|b - Ax_0\|$ and Hessenberg matrix \tilde{H} is known from the Arnoldi iteration [23]. This is a linear least squares problem, for which the solution method is well known.

3.4.2 The CPR Preconditioner

Using a Constrained Pressure Residual (CPR) preconditioner improves the efficiency and stability of the iterative solution method described above [16]. By definition, a preconditioner P of a matrix A has the quality that $P^{-1}A$ has a smaller condition number than A . In other words, $P^{-1}A$ is less sensitive to small variations in the input variable. This improves both stability and convergence rate, as convergence rate tends to increase as the condition number of A decreases.

The CPR preconditioner is a two-step preconditioning method. The first step is to extract a pressure equation from the implicit matrix, which is then iteratively

solved for the pressure correction. This pressure correction is then used to form a new residual. The next step is to apply an inexpensive implicit preconditioner to this residual. Finally, the sum of the two steps is used as the approximate solution.

In more detail, the CPR method implemented in MRST follows the following steps [10]:

1. Eliminate well rate and bottom hole-hole pressure variables. This results in a system $J\mathbf{x} = \mathbf{b}$, where the Jacobian J consists of a large 2 x 2 block system and \mathbf{x} consists of the two primary variables, i.e. $\mathbf{x} = (\mathbf{p}, \mathbf{s}_w)$.
2. Set the first block-row in the system equal to the sum of the two block-rows, i.e., $J(1, m) = \sum_n J(n, m)$. $J(1, 1)$ will then resemble the pressure equation (2.26), which we have seen is close to elliptic. Any equations that harm desired diagonal dominance in $J(1, 1)$ are left out.
3. A two-stage pre-conditioner $M_2^{-1}M_1^{-1}$ is then set up.
 - (a) M_1^{-1} is set up to solve the near elliptic subsystem $J(1, 1)\delta\mathbf{p}^{i+1} = -r_p^i$ to obtain the pressure update δp^{i+1} . For large systems an algebraic multi-grid solver is preferable.
 - (b) M_2^{-1} is based on an incomplete LU-decomposition $LU \approx J$ of the full system, and is set up to perform a variable update $\delta\mathbf{x}^{i+1}$ on the full set of variables by solving $LU\delta\mathbf{x}^{i+1} = -\mathbf{r}_x^i$.
4. Solve the full system with GMRES using $M_2^{-1}M_1^{-1}$ as a pre-conditioner.
5. Recover rate and bottom-hole pressure variables.

Some terms were used here that deserve further explanation. *Multigrid methods* are useful in solving elliptic PDEs, by using a hierarchy of different discretizations. The main idea of is that the convergence of the basic iterative method can be improved by a global correction of the fine grid solution, which is found by solving the coarser grid. The iterative method for solving on the coarser grid can be similarly obtained by using a global correction, that is again found from a coarser grid. This recursive process is repeated until the cost of direct solution on the coarsest grid is negligible compared to doing another relaxation sweep.

For *algebraic multigrid methods*, the hierarchy of discretizations is found solely from the system matrix [17]. It therefore does not need any information about the system of the geometrical problem, which makes it highly useful as a black-box solver for certain classes of sparse matrices. Furthermore, for an *aggregation-based algebraic multigrid method*, the coarsening is obtained by aggregation of the unknowns [20].

3.5 The MATLAB Reservoir Simulation Toolbox

The MATLAB Reservoir Simulation Toolbox (MRST) is an open source toolbox developed by SINTEF Applied Mathematics. The motivation behind its creation was to enable rapid prototyping and testing of new methods in reservoir simulation. It contains routines for grid processing, rock and fluid modeling, numerical discretization, solvers and visualization.

MRST consists of two main parts; a core that offers basic functionality and a set of add-on modules. The MRST core uses a set of common components, implemented using simple structs, that together define a full model:

- *Grid*: Defines the cells, faces and nodes in the grid. The struct also contains reference mappings between a cell and its faces, a cells neighbours, face normals, cell volumes and so on.
- *Rock*: Defines the rock porosity and permeability in each cell.
- *Fluid*: Contains the fluid properties, such as relative permeability, density, viscosity and compressibility.
- *Boundary conditions*: Contains the types (rates/pressure), values and cells each boundary condition applies to.
- *Sources*: The location and value of a source.
- *Wells*: Defines the physical properties of a well and also its control, i.e. the prescribed rate or bottom hole pressure that controls injection and production at each time step.
- *State(s)*: Contains the state variables, such as pressure, saturation and flux values.

The basic structure of an MRST simulation script can be seen in Figure 3.2. The grid, fluid struct and rock struct are passed into a function that creates the desired model of the problem. Most models in MRST inherit the properties of a general reservoir model, but define their own set of constitutive equations in order to implement the desired model.

The sources, wells and boundary conditions are passed into a struct denoted as the schedule. It is named so after the similar data structure in Eclipse, which contains the "schedule" by which the system should be operated. The schedule contains the value of each time steps, along with information about which driving forces are at work during each time step.

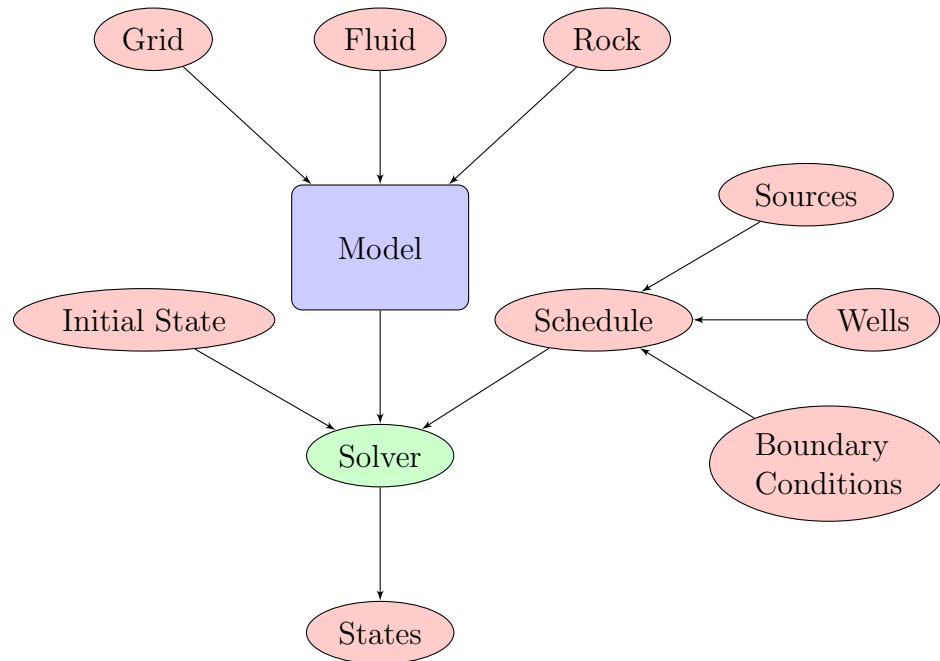


Figure 3.2: The basic structure of an MRST simulation script. A red ellipse indicates a simple struct, a green ellipse indicates a function, and a blue rectangle indicates an instance of a class.

The reservoir model is complete when the initial state, model and schedule have been fully specified. The next step of the simulation is to pass it on to a solver function, which will pass back the results in the form of a *states* struct when it is done

The solver is an example of functions that are normally not contained in the core, but chosen from an appropriate module for the problem. MRST contains a number of such modules, that add to or override functions and methods contained in the core. Several modules were extensively used in this thesis, most notably:

- The *ad-core* module, which contains the framework for fully implicit simulations using ADI
- The *black-oil* modules, that contain the necessary routines to create black-oil models
- The *deckformat* module, used to transfer Eclipse models into the MRST framework, see 3.5.1
- The *co2lab* module that supports the use of VE models, see 3.5.2.
- *AD core* module [10], short for automatic differentiation, ensures that no analytical derivatives have to be calculated or programmed explicitly, see 3.2

- The AGMG module used to increase the speed of the linear solver, using the method explained in 3.4

3.5.1 Integration of Eclipse Models

ECLIPSE constitutes the industry-standard simulator for oil and gas reservoirs, developed by Schlumberger and in wide use across the petroleum industry. Energinet.dk have developed an extensive ECLIPSE model of Stenlille, which has been used with relative success to predict pressure distribution in the aquifer. Since their model will form the basis for the model used in this thesis, an introduction to ECLIPSE input files will be useful here for later reference.

A reservoir model is constructed by setting up the ECLIPSE data file, which comprises of the following seven sections [25]

- *Runspec*: Contains specifications for the simulation, such as the name of the reservoir, the dates for which the simulation is run, which phases are involved, the dimensions of the grid and so forth.
- *Grid*: Contains information about the grid and petrophysics, e.g. the porosity and permeability.
- *Props*: The properties of the rock and fluids, e.g. relative permeabilities, density and compressibility.
- *Regions*: Defines regions of the grid, for example so that different rel.perm. curves can be applied in different regions.
- *Solution*: Defines initial conditions under the assumption of vertical equilibrium.
- *Summary*: What results should be output after running a simulation.
- *Schedule*: Defines how and when the different wells should inject, produce or observe.

The deckformat module contains routines for transferring an ECLIPSE model into the MRST framework.

3.5.2 The co2lab module

The co2lab module [4] contains the functions and routines necessary to perform a VE simulation in MRST.




Figure 3.3: The basic structure used when converting a traditional black-oil model into a VE model.

The basic routine used for constructing a VE model is shown in Figure 3.3. First, the function `topSurfaceGrid()` in this module is used to construct the two-dimensional version of the grid that is utilized in the VE model. This new grid describes the top surface $Z(x, y)$ as shown in Figure 2.2, and also stores information about the height H of each column in the three-dimensional grid.

The function `averageRock()` is then used to construct the two-dimensional version of the rock model, by simply found by taking averages of the rock properties in each column. Finally, the VE fluid can be constructed using one of the models implemented in the module, and a full VE model in MRST constructed.

Chapter 4

Description of the Stenlille Model

Recall our discussion of aquifer flow modelling in Chapter 2, where we derived two sets of model equations for two-phase flow in porous media. In Chapter 3 we then outlined specific numerical methods capable of solving these equations, and introduced the MRST framework in which both model and methods will be implemented. With these two chapters to build on, we are ready to give a description of the model of Stenlille that will be used in later simulations.

The simulations performed for this thesis used MRST, while earlier simulations were run primarily using ECLIPSE. Substantial work has gone into developing the ECLIPSE model of Stenlille, where we now understand the term ECLIPSE model to mean the specific grid, rock properties, fluid properties and well models that have been created for the Stenlille aquifer. The MRST model will be constructed by transferring each of these elements into the MRST framework.

The MRST model will be constructed so as to mimic the ECLIPSE model as closely as possible. This was highly useful for the process of debugging and understanding the limitations of the MRST model. When the two models were similarly constructed, their simulation results should also be similar, and the two can be compared to validate the implementation of the MRST model.

The first section of this chapter is devoted to describing the three-dimensional model that will be utilized for black-oil simulations. This section is further split into five subsections, that each correspond with one of the basic structures of an MRST model, as they were shown in Figure 3.2. The second section will then describe the results of transforming this model into its two-dimensional VE counterpart.

4.1 The Three-Dimensional Model

4.1.1 Grid and Boundary Conditions

Before introducing the grid and boundary conditions we will use for the Stenlille model, we should first understand the situation we are aiming to model. It will therefore be useful to give some general remarks on the structure of the Stenlille aquifer.

Figure 4.1 shows a schematic cross section view of the different layers the aquifer is made up of. Impermeable or semi-impermeable layers of silt and clay split the aquifer into six different storage zones, of which Zone 5 has the largest storage potential. The simulations in this thesis have for this reason been chosen to focus on this storage zone.

Zone 5 is separated from the zone above by an impermeable layer of clay, capable of acting as a caprock that prevents the gas plume from escaping. This layer is also thought to separate Zone 5 hydraulically from the zones above. Conversely, the layer that separates Zone 5 from the storage zone beneath it is understood to be semi-permeable, and is thought to allow hydraulic contact between Zones 5 and 6.

Figure 4.2 shows the grid used in ECLIPSE simulations, that will also be utilized in MRST. Storage zone 6 is included in the grid since it is in hydraulic contact with zone 5, and they are indicated using blue and red cells respectively. The cells belonging to the layer that separates them are indicated in green.

Although it is impossible to see from the figure, the three regions of the grid are separated by a small but non-zero distance that disconnects them from each other. Hydraulic connectivity is instead taken into account in the ECLIPSE model by the inclusion of *aquifer connections*.

An aquifer connection in ECLIPSE connects selected cells in the grid with analytic aquifers, which are used to model the complex pressure changes that occur along the boundaries of the grid. The areas right above and below the grid are made up of semi- or impermeable layers, so it makes sense to have no-flow boundary conditions on the top and bottom boundary faces. The boundary faces on the east, west, north and south side, however, are in contact with the rest of the aquifer. When gas is produced, the pressure in the gridded portion will sink, driving an influx of water into it from the rest of the aquifer. Conversely, when gas is injected, the increase in pressure will drive the water back out. This is the effect that is modelled by using analytic aquifer connections along the boundary. In the case of Stenlille, the aquifer connections also connect storage zones 5 and 6




Figure 4.1: A schematic cross section view of the different layers of the Stenlille aquifer. Layer 5 has the largest storage potential and will therefore be the object of study in this thesis. It is sealed from the layer above it by an impermeable layer of clay. However, the layer of clay that separated it from the zone below is semi-permeable, which allows for hydraulic contact between the two. Figure from [12].




Figure 4.2: The grid used in black-oil simulations of Stenlille. Storage zones 5 and 6 are indicated in red and blue respectively, and the layer that separates them is indicated in green.

hydraulically.

The ECLIPSE model of Stenlille connects the cells along the east, west, south and west boundary of the grid with Carter-Tracy aquifers. The Carter-Tracy model expresses the pressure drop at the aquifer boundary as

$$p_a - \bar{p} = \frac{Q_a}{\beta} \text{PI}_D(t_D), \quad (4.1)$$

where Q_a is the aquifer inflow rate, \bar{p} is the average water pressure on the aquifer boundary, β is an aquifer influx constant, $\text{PI}_D(t_D)$ is interpolated from a table, and t_D is given by $t_D = t/T_c$, where T_c is a time constant [26]. The time and

aquifer influx constants are more specifically given by

$$T_c = \frac{\mu_w \phi_a c_{t,a} r_{o,a}^2}{k_a c_1}, \quad (4.2)$$

$$\beta = c_2 h_a \theta c_t r_o^2, \quad (4.3)$$

where k_a is permeability of the analytic aquifer, ϕ_a is its porosity, $r_{o,a}$ is its inner radius, h_a is its height, $c_{t,a}$ its total compressibility, and c_1 and c_2 are constants. The average inflow rate from the analytic aquifer to a specific grid block i over the time interval $(t, t + \Delta t)$ is further given by

$$\bar{Q}_{ai} = \alpha_i (a - b(p_i(t + \Delta t) - p_i)) \quad (4.4)$$

where α_i is the area fraction for each connection and the variables a and b are given by

$$a = \frac{1}{T_c} \frac{\beta \Delta p_{ai} - W_\alpha(t) \text{PI}'_D(t + \Delta t)_D}{\text{PI}_D(t + \Delta t)_D - t_D \text{PI}'_D(t + \Delta t)_D}, \quad (4.5)$$

$$b = \frac{\beta}{T_c (\text{PI}_D(t + \Delta t)_D - t_D \text{PI}'_D(t + \Delta t)_D)}. \quad (4.6)$$

Here, Δp_{ai} is the pressure drop $p_{a0} + \rho g(d_i - d_a) - p_i(t)$, where d_i is the grid block depth, and d_a is the datum depth of the aquifer.

MRST does not currently support the addition of analytic aquifers to the model, so some changes were made to the grid in an attempt to imitate the effects they will have on simulations. The volume and rock porosity of the cells shown in yellow in Figure 4.3 were multiplied by a factor of a 1000, allowing them to provide pressure and flow support to the rest of the grid. Additionally, they were defined to be neighbours of the cells indicated in blue, so that the two zones were in hydraulic contact with each other.

Examining the grid in Figure 4.3 again, it is possible to see darker columns towards the top. These are caused by the local grid refinements made around each well in the model, which make the cells so small that the area appears black when viewing the grid as a whole.

The refinements are seen more clearly in the closeup in Figure 4.10. They were added to increase the resolution of the solution in the area around each well. Each local grid refinement was constructed using a column of 11 cells, that included the well cells themselves, as host cells. The cells in this column were then refined, with a refinement factor of 11 in the horizontal directions, and 2 in the vertical direction.

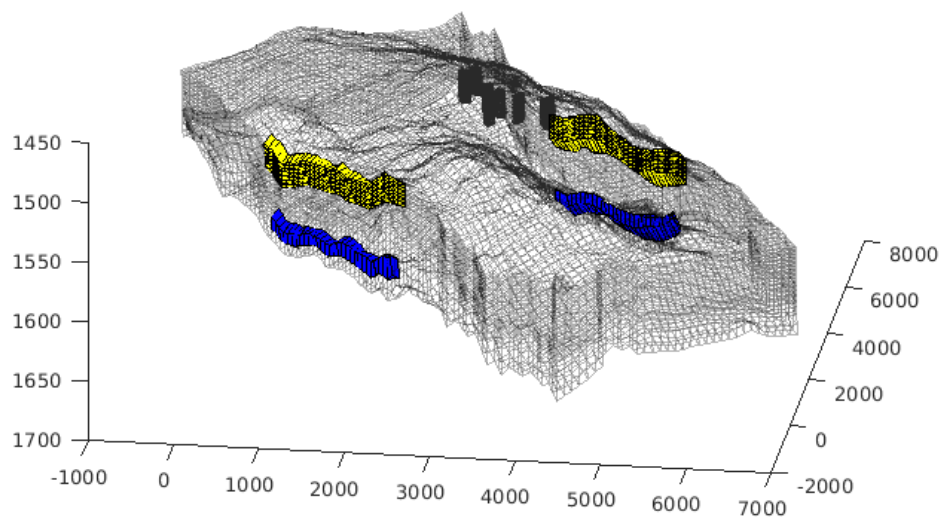


Figure 4.3: The grid used in MRST simulations. The cells indicated in yellow have had their volume blown up 10.000 times, in an effort to imitate the effect from the analytic aquifers used in ECLIPSE simulations. They have also been defined to neighbour the cells shown in blue, with the purpose of enabling hydraulic connectivity between storage zones 5 and 6.




Figure 4.4: A closeup view of the local grid refinements made around each well.


4.1.2 Rock Properties

The process of transferring the rock model from ECLIPSE to MRST was straightforward; the porosity and permeability value for each cell in the MRST grid were simply assigned the corresponding value in ECLIPSE.

Two scatter-plots, showing the porosity and permeability values for all cells in the grid, can be seen in Figure 4.5. The plots are color coded, using the same color convention as in Figure 4.2. The average value porosity or permeability for each layer is also indicated with a triangle on the right-hand side axis. Figure 4.5b shows the values for the horizontal permeabilities. The vertical permeabilities have been set to equal 10% of the horizontal permeability.

From Figure 4.5a we can see that the cells in the grid typically have porosity values ranging between 0.15 and 0.3, with no major distinctions between each zone. From Figure 4.5b however it is clear that the permeability values differ greatly from zone to zone, with the semi-impermeable layer between them having the lowest values, which is as expected.

Rock compressibility was also included in the model, by using Equation (2.14) with a reference pressure given by $p_{\text{ref}} = 158$ bars and a rock compressibility given by $c_r = 10^{-14}$.



(a) A scatterplot of the porosity value in each cell.

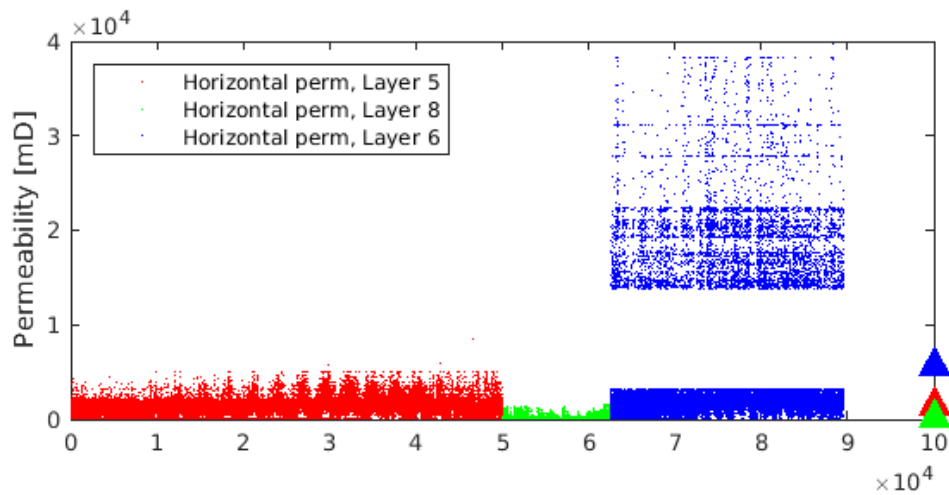
(b) A scatterplot of the permeability values k_x values in each cell.

Figure 4.5: Scatter-plots showing the porosity and permeability for each cell in the grid. The values indicated in red and blue belong to Zones 5 and 6 respectively, and the values indicated in green belong to the semi-impermeable layer between them. The average porosity and permeability values for each are indicated, by the use of triangles on the right-hand axis. The horizontal permeabilities were set to equal each other in the model, so that $k_y = k_x$. The vertical permeability was set to equal 10% of the horizontal permeability, to give a better model fit.

Table 4.1: The fluid properties used in the Stenlille model. Some properties are assigned using analytic functions, others are assigned by tabulating from a table. If the latter is the case, the entry for that property contains a reference to a figure showing a plot of the table data.

	Brine	Natural gas
Density	$\rho_w(p_w) = \rho_{w,s} + [1 + c_w(p - p_{\text{ref}})]$ $\rho_{n,s} = 0.7681 \text{ kg/mm}^3$ $c_w = 1.92 \cdot 10^{-5}$ $p_{\text{ref}} = 155.1 \text{ bars}$	$\rho_n(p_n) = b(p_n)\rho_{n,s}$, $\rho_{w,s} = 1116 \text{ kg/m}^3$ $b(p_n)$ tab. from Figure 4.6
Viscosity	$\mu_w = \mu_{w,\text{ref}} + c_v(p - p_{\text{ref}})$ $c_v = 1.87 \cdot 10^{-5}$ $p_{\text{ref}} = 155.1 \text{ bars}$	Tab. from Figure 4.7
Rel. permeability	Tab. from Figure 4.8a	Tab. from Figure 4.8b
Cap. pressure	Tab. from Figure 4.9	Tab. from Figure 4.9

4.1.3 Fluid Properties

The fluid properties were taken directly from the ECLIPSE model with no alterations. The ECLIPSE model utilizes PVT functions for the fluid properties, i.e., functions that are assumed to depend solely on the pressure, temperature and volume of the fluid. Assuming further that the temperature in the aquifer is constant, and that volume can be modelled as a function of pressure, this reduces to the assumption stated in Section 2 that all fluid properties depend solely on pressure (or equivalently on saturation).

The fluid properties are summarized in Table 4.1. Some properties are modelled using an analytic function, and the function constants are then given in the table as well. If this is not the case, the table entry contains instead a reference to a figure showing the values for which it was tabulated from.

Examining this table, it is clear that water properties are calculated using analytic expressions more often than gas properties are. This is easily explained by noting that water is nearly incompressible, which implies that the water volume does not change much with pressure. In this case, many properties can be more accurately modelled using a constant rate of change.

Figure 4.8 shows the relative permeability curves for water and gas. Two different curves are given for each fluid phase, where one applies in cases of imbibition,

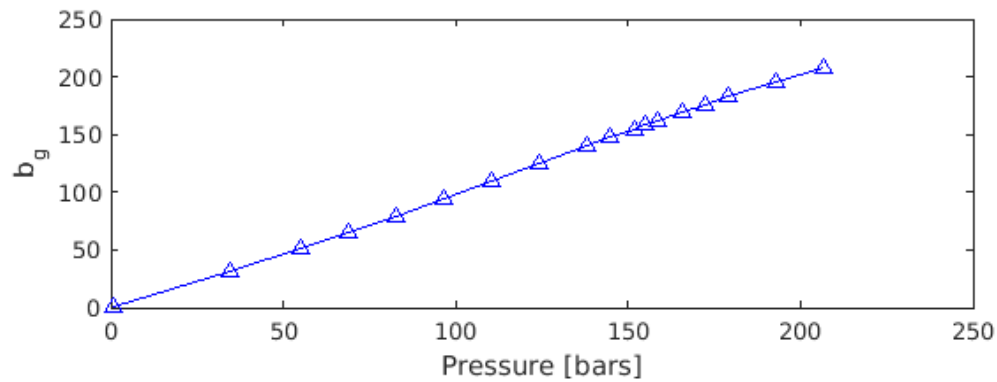


Figure 4.6: A plot of the entries in the tabulation table used to calculate the inverse formation volume factor b_g for gas as a function of phase pressure p_n .


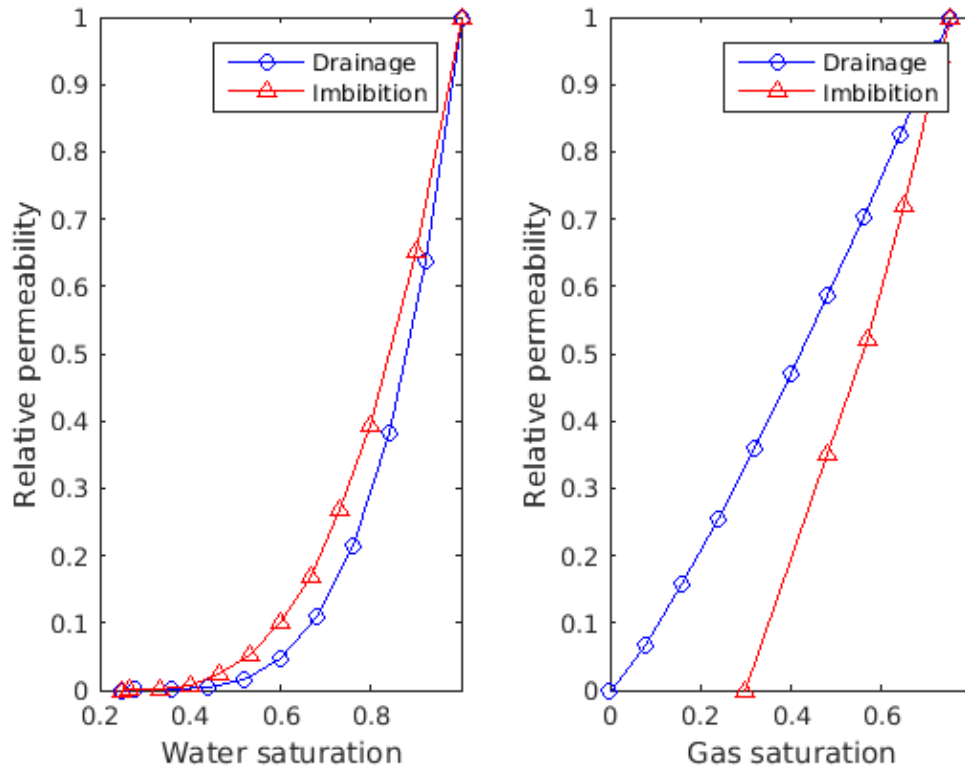


Figure 4.7: A plot of the entries in the tabulation table used to calculate the viscosity of natural gas μ_n as a function of phase pressure p_n .

while the other applies in the case of drainage.



(a) A plot of the relative permeability curve for water, $k_{r,w}$ in Stenlille (b) A plot of the relative permeability curve for gas, $k_{r,n}$ in Stenlille

Figure 4.8: Plots showing the relative permeability curves for water and curves. The curves applying in the case of imbibition are shown in red, and the curves applying to the case of drainage is indicated in blue.

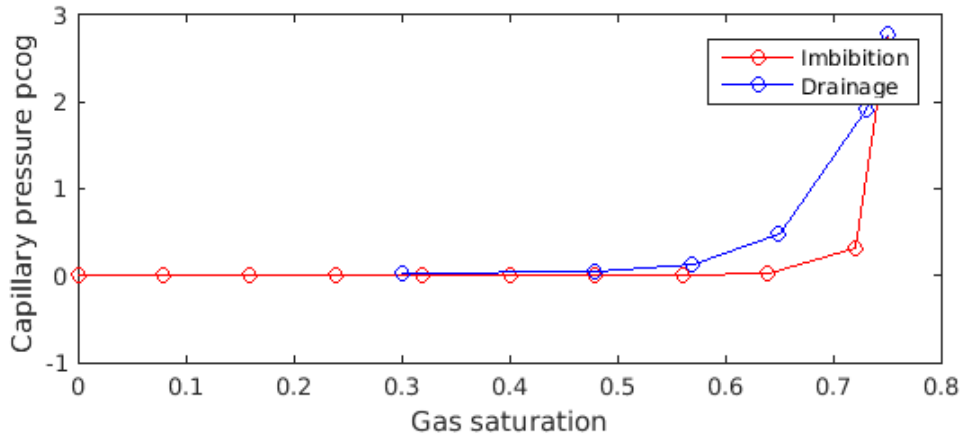


Figure 4.9: A plot of the capillary pressure function p_{cap} , as a function of gas saturation s_n .

4.1.4 Initial Conditions

For the ECLIPSE simulations, initial conditions were set based on the results of a prior simulation covering the time span from the very beginning of gas injection into Stenlille, up to the start date of the simulation. These results were not available from the summary file, and could therefore not be used in the MRST simulations.

Instead, it was assumed that the initial conditions were static, so that the flux is zero in all cells. In this case, Darcy's law simplifies to

$$\frac{dp_\alpha}{dz} = \rho_\alpha(p_\alpha)g, \quad (4.7)$$

for each phase $\alpha = \{n, w\}$. The solution section of the ECLIPSE data file contained information on appropriate initial conditions for this case; the reference was assumed to be $p_{ref} = 158.795$ bar at a reference depth $z_{ref} = 1500$ meters, and the gas-water contact located at 1510.95 m below the surface.

Since the reference depth is above the gas-water contact, this equation should be first solved for all cells containing gas. This is easily done using an ODE solver in Matlab. Thereafter, assuming water to be incompressible, the pressure can be easily evaluated for cells containing water, by the relation

$$p_w = \rho_w g(z - z_{ref}) + p_{ref}. \quad (4.8)$$

4.1.5 Wells

While Figure 4.4 gave a sneak peek of the wells used in the Stenlille model, we will now take a moment to discuss them in more detail.

Figure 4.10 shows an overview of the wells that extend down into storage zone 5, with the observation wells indicated in blue, and injection/production wells indicated in red. All wells have a tube diameter equaling 0.216 m. The three observation wells have a reference depth of 1610 m, while the injection/production wells have their reference depth set as the height of the bottom-most perforation in the well. The location of the cells that contain a well perforation can be seen more clearly in Figure 4.10, where they are indicated by the use of a thicker line on the well illustration.

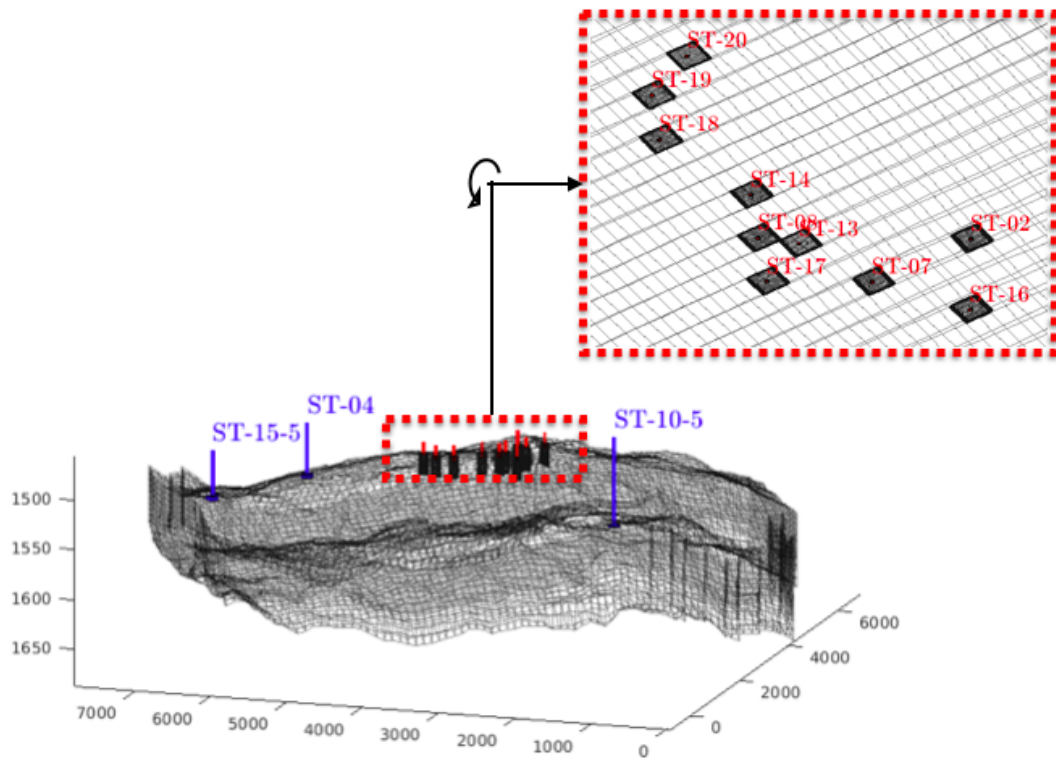


Figure 4.10: An overview of the wells that extend down into storage zone 5. Injection/production wells and observation wells are indicated in red and blue respectively.




Figure 4.11: The top-surface grid used in the VE model.

4.2 The Two-Dimensional Model

The purpose of this section is to show the results of transforming the three-dimensional model of Stenlille into its two-dimensional counterpart, which will be employed when running VE simulations.

First, the three-dimensional grid has to be transformed into a two-dimensional grid suitable for the VE model. This is done by constructing a top-surface grid $z(x, y)$ that describes the caprock surface, and also contains information about the vertical height of each column. The results of doing this for the Stenlille grid can be seen in Figure 4.11. The figure shows a three-dimensional plot of the Stenlille top-surface grid, along with an indication of the column height H of each cell. It was necessary to remove the local grid refinements before constructing the top-surface grid, as the `topSurfaceGrid()` function assumes a regular cartesian grid as input.

The two-dimensional rock structure was constructed by simply taking a harmonic average the rock properties in each column of the original grid. The vertically averaged porosity this resulted in is shown in Figure 4.12, and the vertically averaged permeability is shown in Figure 4.13. The cells that appear bright yellow in the porosity figure have been chosen as aquifer support cells, and had their porosity multiplied by a factor of a 1000.

The initial pressure was set similarly as in the three-dimensional model, using




Figure 4.12: The vertically averaged porosity. The cells appearing in bright yellow have had their original porosity value multiplied by a factor of 1000, in order to imitate the pressure- and flowsupport from the analytic aquifers used in the ECLIPSE model.

the tactic explained in Chapter 4.1.4. The only alteration made was in the interpretation of z in Equations (4.7) and (4.8), which was now interpreted to be the depth $z(x, y)$ of cells in the top-surface grid, instead of the vertical coordinate of the cell centroid. The reference pressure was also increased slightly, to a value of $p_{\text{ref}} = 160.9$ bar at $z_{\text{ref}} = 1500$ m.

The construction of initial saturation values needed more thought. In the case of no residually trapped water or gas, the saturation terms in the VE model are simply interpreted to be the fractional height of the gas or water column. More specifically, the gas saturation will then be given by $s_n = h/H$, and the water saturation will be given by $s_w = (H - h)/H$. The initial saturations used in the three-dimensional model, however, accounted for residual water trapped in the gas column. In this case, the saturations will be given by

$$s_n = \frac{h}{H} s_{n,r}, \quad (4.9)$$

$$s_w = \frac{1-h}{H} + \frac{h}{H} (1 - s_{n,r}). \quad (4.10)$$

The initial saturations could then be computed, using the same location of the gas-water interface as before.




Figure 4.13: The vertically averaged permeability. The darker blue rectangles were caused by changes made to the permeability values in the original ECLIPSE model, made with the intent of improving the history-matching capabilities of the model. These changes made the vertically averaged permeability appear peculiar when plotted, but was not found to be problematic.

The state struct used in VE simulations also contains a variable denoted $s_{n,\max}$, which was not needed in the black-oil model. This variable stores the maximal gas saturation ever experienced by each cells, and is used to model hysteresis. If the gas saturation value in a cell is smaller than the historical maximum $s_{n,\max}$, the difference is assumed to be residually trapped. For the VE model, it was initially set to equal $s_{n,\max} = s_{n,\text{init}}$.

Next, the VE fluid model of Stenlille needs to be constructed. Recall the observation made in Chapter 2.3.2, where we saw that the fractional flow formulation in phase pressure of the VE model was formally similar to the same reformulation of the three-dimensional black-oil model, given by (2.45) and (2.47). The difference between the two lay in the formulations of the mobility and capillary pressure; the VE model used the upscaled versions of the mobility term, given by (2.57), and a modified capillary pressure term.

The VE fluid model for Stenlille can therefore use the same gas inverse formation volume factors $b_n(p)$, the same compressibility of water c_w and the same viscosities $\mu_n(p)$ and $\mu_w(p)$ as the ones for the black-oil fluid model, as they were summarized in Table 4.1. Then, the VE fluid model was simplified to assume no capillary pressure, so that the modified capillary pressure term found in Equation (2.47)

could be used directly. The relative permeability curves were also simplified. The residual saturation was assumed to be the same as before, but the relative permeability curve of each fluid was assumed linear for the saturation values where the fluid is mobile.

Lastly, the well models had to be adapted to fit the VE grid, a process that was largely straightforward. Each well was placed in the correct cell in the top-surface grid, more specifically the cell which in its column contained the cells where the three-dimensional well was placed. The radius was kept the same, but the well indexes found for the three-dimensional well were summed together, to account for the effect of each perforation.

Chapter 5

Simulation Results

Recall our derivation of two different sets of constitutive equations to model aquifer flow in porous media, which was followed by an introduction to a numerical strategy capable of solving them. With this mathematical background to build upon, we then introduced the Stenlille implementation of these models specifically. At this point, the stage is finally set to show and discuss the simulation results obtained from utilizing these two models.

We will examine the results of three different simulations, each run over the same simulation scenario. The first simulation will use the three-dimensional black-oil model we described in Chapter 2.2.4, and the second and third will use its two-dimensional VE counterpart, as we derived it in Chapter 2.3.1. The results will be analyzed by comparing them with either historical data, other simulation results, or a combination of the two.

The Simulation Scenario:

The simulations will be run covering the time period 01.09.11-01.06.12, with the goal of successfully simulating the fluid flow caused by injection and production of gas in this time period.

The wells in Stenlille are rate-driven, meaning that they are controlled by specifying rates at which they should inject or produce gas. It was therefore natural to use rate-controlled well schedules in the simulation as well, so that the well rates were treated as a given at each time step. The well schedule rates were set using historical data, and are plotted individually for the time period 31.07.11-01.06.12 in Figure C.1 in Appendix C.1.

The simulation start date of 01.09.11 was not set arbitrarily. Recall our discussion of initial conditions in Chapter 4.1.4, where discussed how the simulations would

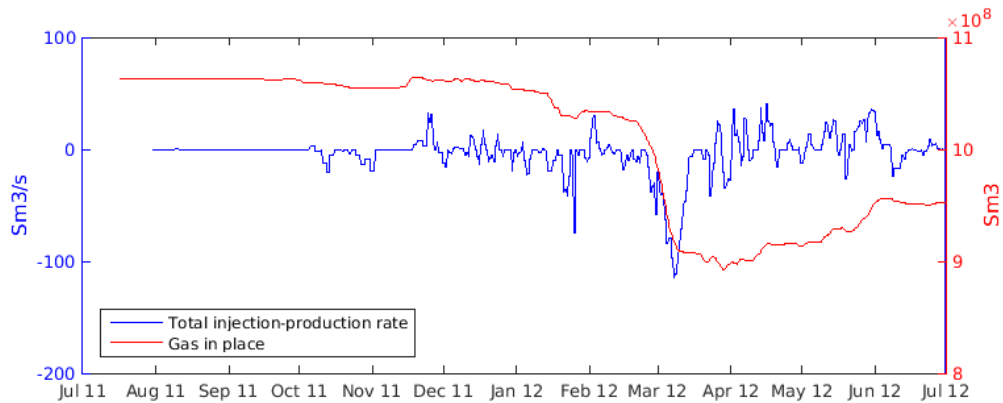


Figure 5.1: The total injection/production rate for Stenlille in the time period 31. July 2011 to 30. June 2012. Positive rates are interpreted as injection and negative rate interpreted as production.

need to be initialized using the assumption of a static starting state. Figure 5.1 shows the total injection and production rate of gas in Stenlille, over the time period 31.07.11-01.06.12., indicated in blue. The figure also shows the total volume of gas residing in the aquifer, this time indicated in red. From the figure it is clear that no gas was injected or produced in the time 31.07.11-30.08.11, allowing the aquifer to settle into a state close to no-flow at 01.09.11.

Validation Methods:

We will use comparisons of well pressures and water production rates as the primary way to judge the success of each model, along with observations of how the gas-water interface moves over the course of each simulation.

The comparison of well pressures is motivated by the fact that the wells in Stenlille are rate-driven, and that the prescribed flow rates are obtained by introducing pressure changes in each well. The specific magnitude of these pressure changes will depend on both the well model and the aquifer model as a whole. Comparing simulation well pressures, against both historical values and values obtained from other simulations, will therefore be a powerful tool in judging the validity of each model.

For Simulation A, we will primarily compare the results with those from a similar ECLIPSE simulation. As they both utilize the black-oil model, the results will be expected to correlate well, but not perfectly as some aspects of the MRST model have been changed compared to the ECLIPSE model of Stenlille. To further judge to what degree these simplifications pose a problem for simulations of Stenlille, it will be useful to compare them against historical values.

Unfortunately, the comparison of well pressure results from MRST simulations against their historical counterpart is not a straightforward process. The wells in Stenlille record values for the Tubing Head Pressure (THP), defined as the pressure measured at the tubing head, which is located at the surface end of the well. In the context of simulation it is more natural to be concerned with the Bottom-Hole Pressure (BHP), defined as the pressure measured at the bottom-most perforation made in the well. Since these measurements are much less accessible, they are in practice often not measured at all. In Stenlille, the observation wells ST 10-5 and ST 15-5 are the only wells that record its value.

The lack of BHP measurements is a common issue in reservoir engineering, and is dealt with in ECLIPSE by employing Vertical Flow Performance (VFP) tables to convert between it and the THP. However, MRST does not currently support the use of VFP tables. Before presenting the simulation results, we will therefore stop and give a short discussion on the validity of ECLIPSE simulation results, and compare the results for THP and BHP. The observations made here will then be made use of later in the discussions of the MRST simulation results.

On a final note, recall that there are 13 wells with perforations in storage zone 5 in Stenlille, of which ten are active and can produce water. The number of wells poses a slight difficulty when it comes to presenting the results in a space efficient manner, as each simulation will result in ten plots of water production rates, and thirteen plots of the well pressures. In order to not disrupt the flow of discussion, the full set of results from each simulation will therefore be given in the appendix. The results section this chapter will contain a reference to the appendix where the full set of results is given, along with a set of representative figures from it.

Table 5.1: A table showing the percentage of successful simulation prediction by ECLIPSE of the THP for each well. A simulation value was seen as successful if it lay within 10% of the historically recorded value.

Well	Successful time-steps
ST-02	98.40%
ST-07	98.86%
ST-08	98.86%
ST-13	99.08%
ST-14	97.94%
ST-16	98.86%
ST-17	87.99%
ST-18	97.83%
ST-19	96.91%
ST-20	93.36%

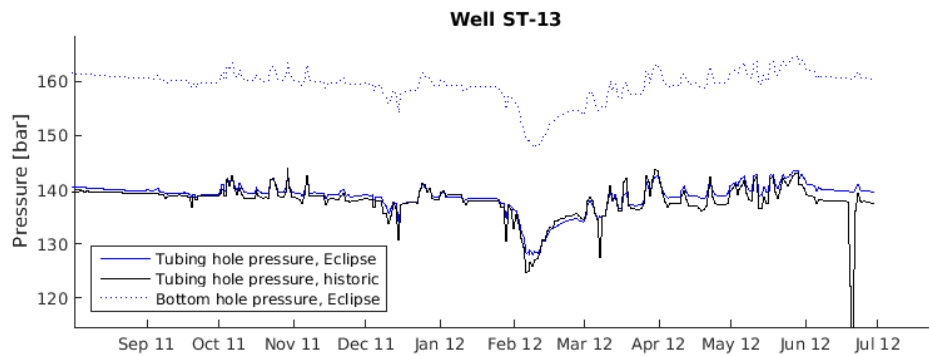
5.1 A Note on the ECLIPSE Simulation Results

At this point a comparison of ECLIPSE simulation results and historical data will be given, which will be useful when later seeking to determine the success of the MRST simulation results.


The full set of the ECLIPSE results is given in Appendix C.2.1. Figure C.2 shows the ECLIPSE simulation results for the tubing head pressure and bottom-hole pressure for each individual well, along with the historical tubing head pressure. The three curves are plotted in solid blue, dashed blue and solid black respectively. The simulation prediction and historically recorded values for water production rates are plotted in Figure C.3, this time using dashed blue and solid black respectively. Please note that the simulated and historical water production rates have been plotted on different y -axes.

Wells ST-13 and ST-20 were chosen as representative examples for this discussion, and their well pressures are plotted again in Figure 5.2 on the next page. From the figure it seems that the simulation results correspond reasonably well with the historical values, at least in the sense that that the relative error stays small for most time steps. If we further define a simulation value for the THP to be successful if it is within 10% of the historically measured value, we can see from Table 5.1 that the majority of time steps were successful in predicting well pressures.

Going back to examining the plots for the simulated and historically measured values for tubing head pressures, it is clear that the simulation pressures tend to lie slightly higher than the historically measured value. Examining e.g. the



(a)




(b)

Figure 5.2: Well pressure results for wells ST-13 and ST-20, as they were found in the ECLIPSE simulation. The tubing hole and bottom-hole well pressures from ECLIPSE are plotted in solid blue and dashed blue respectively. The historically measured tubing head pressure is plotted in black.


well pressure results for well ST-13, we see that it experiences consistent pressure differences as high as 2 bar. However, considering that this still yields a small relative error, this was not found to be of great importance.

Examining again the plot for well ST-13, we see that the simulated and historical well pressures for this well consistently follow the same trend with few major discrepancies. Nonetheless, there appears to be a factor difference between the pressure changes needed to achieve certain flow rates; from the figure it appears that the simulation well pressures "underreact" with a factor of two compared with the historical results.

Examining now the plot for well ST-20, it is easy to see that there are several



(a)



(b)

Figure 5.3: The water production rates predicted for wells ST-13 and ST-20. The rates predicted by the ECLIPSE simulation are shown in dashed blue, and the historically recorder water production is shown in black.

places where the historically recorded pressure drops suddenly. If the plots were zoomed out, the reader would see that the pressure drops extremely low, often down to 1 bar. It is not entirely clear what causes this behaviour, but it is highly possible these sudden dips occur whenever the well has been closed off. This would cause the well pressure to be isolated from the rest of aquifer pressure, in which case it would quickly decrease. The ECLIPSE simulation results do not mirror this behaviour, perhaps because the simulation well pressure is always measured as if the well were open.

The simulation predictions for water production rates by wells ST-13 and ST-20 are plotted in Figure 5.3, together with the water rate that was observed historically. From the figure, it is clear that the two bear close to no resemblance, neither in shape nor the scale of the production rate. The results of Simulation A will shed more light on this discrepancy, and we will therefore postpone the discussion of these results until the next section.




Figure 5.4: The initial gas saturation used in Simulation A. The gas-water interface was assumed to be flat and located at a height of $z_{\text{water}} = 1510.95$. The pore space of the cells above this point were assumed filled with gas and residual water, yielding a saturation of $s_n = 0.75$. The cells below this point were assumed completely filled with water.

5.2 Simulation A: Black-Oil Model of Stenlille

The goal of Simulation A is simply to recreate the results of an ECLIPSE simulation, covering the time period 01.09.11-15.07.12, in order to validate the the MRST model of Stenlille has been implemented correctly. The term MRST model refers here to the model implemented in MRST as was described in the previous chapter.

Simulation Set Up:

The MRST model was set up as described in Chapter 4, with the grid shown in Figure 4.2, the rock properties shown in Figure 4.5 and the wells shown in Figure 4.10 and the fluid properties given in Table 4.1.

The initial saturation was set as shown in Figure 5.4, with a flat gas-water interface located at 1510.95 meters below the surface. The pore space of the cells above this point was assumed filled with gas and residual water from the drainage process, yielding a gas saturation of $s_n = 0.75$. The location of the gas-water interface was chosen so that the total volume of gas in place would equal 10.74 Sm^3 , matching




Figure 5.5: The initial pressure used in Simulation A. The aquifer was assumed to be in vertical equilibrium at the start of the simulation, with a reference pressure $p_{\text{ref}} = 159.9$ bar at a reference depth of $z_{\text{ref}} = 1500$ m.

the resident gas volume given in Figure 5.1 for the starting date of the simulation.

The initial pressure was calculated using the method outlined in Chapter 4.1.4, with a reference pressure of $p_{\text{ref}} = 159.5$ bars at a reference height of $z_{\text{ref}} = 1500$ m. The initial pressure this resulted in is plotted in Figure 5.5.

The fluid model was set up as described in Chapter 4.1.3. The MRST fluid model only allowed one relative permeability curve to be defined for each fluid phase, which made it necessary to decide ahead of time if the curves for imbibition or drainage should be used. Examining Figure 5.1 again, it is clear that there are rough seasonal variations in the well rates. Production was most common in the time period from simulation start to mid March, after which injection was most common until the end time of the simulation. Therefore, the MRST simulations will use the imbibition curves from Figures 4.8a and 4.8b in the first time period, and the drainage curves in the second time period. This implementation of the relative permeability curves differs from the one used in the ECLIPSE model, and is a possible cause of errors.

Comparison and Results:

The full set of results for this simulation is given in Appendix A. Figures A.1 and A.2 show the simulated bottom-hole pressures at each time step, together with the simulated pressure found in the ECLIPSE simulation. Figure A.3 shows the water production rate predicted for each well, also together with the values prediction found in the ECLIPSE simulation. The simulation had a total run time of 7 hours




Figure 5.6: The bottom-hole well pressures calculated in Simulation A for wells ST-07 and ST-20. The results from Simulation A are plotted in red, and their ECLIPSE counterpart is plotted in blue.

and 34 minutes.

Wells ST-07 and ST-20 were chosen as representative examples for this simulation, and their well pressures are plotted on the next page in Figure 5.6. The simulated well pressures were found to agree reasonably well, at least in the sense that they consistently followed the same trend. However, it seemed for several of the wells that the pressure found in Simulation A "overreacted" compared to the simulated pressure in ECLIPSE. More specifically, the MRST simulation wells often required a bigger pressure difference than the ECLIPSE simulation, in order to inject or produce gas at the prescribed rate.

There are three aspects we know have been implemented differently in the MRST model of Stenlille; it utilized static initial conditions, a simplified version of the analytic aquifer connections, and a simplified method for picking the relevant relative permeability curve for each time step. Judging from the figures showing the well pressure results, it does not seem that this first change had a significant impact, as the relative error in the first time step is quite small for all wells. The same is

thought of the use of simplified aquifer connections, as each of the wells follow the correct trend, which indicates that the aquifer connections are providing a suitable pressure and flow support.

This last statement is supported by observations made during the debugging process, where it was seen that intentionally or unintentionally altering the aquifer support did not have much impact on the results. More specifically, it was found that different cells could be chosen to have their volume increased, and this factor could be made smaller or larger, without observing noticeable changes in the results. Furthermore, it was observed that the link between the two storage zones could be forgotten completely, without having a significant impact on the well pressure curves. The most important thing seemed to be to provide some form pressure- and flow support to the grid. The specifics of this support were found to be of little importance.

It is possible that this last statement holds true only if the simulations span a time period of less than a year. It seems reasonable that the hydraulic connectivity between the two storage zones, and the specific way the rest of the aquifer provides pressure and flow support, is of much higher importance when running simulations that span the entire operational lifetime of Stenlille.

In any case, the observation made about the irrelevance hydraulic connectivity between the two storage zones for shorter simulations, proved to be quite useful in the debugging process. The grid cells not belonging to storage zone 5 could then be removed, reducing the number of grid cells by nearly half. The simulation run times using this reduced grid were found to be less than half of the simulation run times experienced when using the full grid.

Returning to our discussion of the discrepancy between the MRST simulation results and ECLIPSE simulation results, we believe at this point that the simplified initial conditions and simplified aquifer connections are not the main culprits behind them. This leaves us with the possibility that it was the simplified method for choosing the relative permeability curves that causes the difference between the two simulation well pressure results.

Returning back to Figure 5.6, it is clear that while the well pressures simulated by MRST follow the correct trend, they appear to either underreact or overreact compared with the equivalent ECLIPSE results. According to Darcy's law, as it was stated in Equation 2.6, the flow caused by pressure differences is proportional to the permeability tensor and the relative permeability of the fluid. It is therefore likely that a factor error is caused by differing values for either of these two.

The permeability values in each cell are believed to have been set exactly as in




Figure 5.7: Typical relative permeability curves for a non-wetting phase, with the scanning curve also indicated. $s_{n,max}$ is the maximum gas saturation, $s_{n,cri}$ is the critical gas saturation, $s_{n,crt}$ is the trapped critical saturation. Illustration from [26].

the ECLIPSE simulation, which leaves us with the possibility that the relative permeability curves are causing the discrepancy between the simulation results. As has already been mentioned, the MRST simulations of Stenlille used a quite simple method to calculate the relative permeability values, where the relevant curve was chosen according to if the time step belonged to production season or injection season. As some of the wells inject during production season, or vice versa, this method will undoubtedly cause errors.

The ECLIPSE simulations employed a much more sophisticated scheme for determining the relative permeability curve. First of all, the ECLIPSE simulations chose the correct flow scenario by examining the flow direction in each cell. Secondly, the relative permeability curves shown in Figure 4.8 are interpreted by ECLIPSE to be the *primary drainage curve* and the *pendular imbibition curve* for the model.

The primary drainage curve applies to a drainage process which starts at the maximum wetting phase saturation and decreases all the way down to the minimum wetting phase saturation. Similarly, the pendular imbibition curve is used if the initial wetting saturation is at the minimum possible value and increases upwards to the maximum value [26]. If either the drainage process or the imbibition process is reversed at some point, ECLIPSE does not simply run back over its previous values, but uses a *scanning curve* to make the switch [27]. A typical example of such a scanning curve is shown in Figure 5.7. The ECLIPSE model of Stenlille utilizes a Killough scanning curve specifically [9]. For more information on the

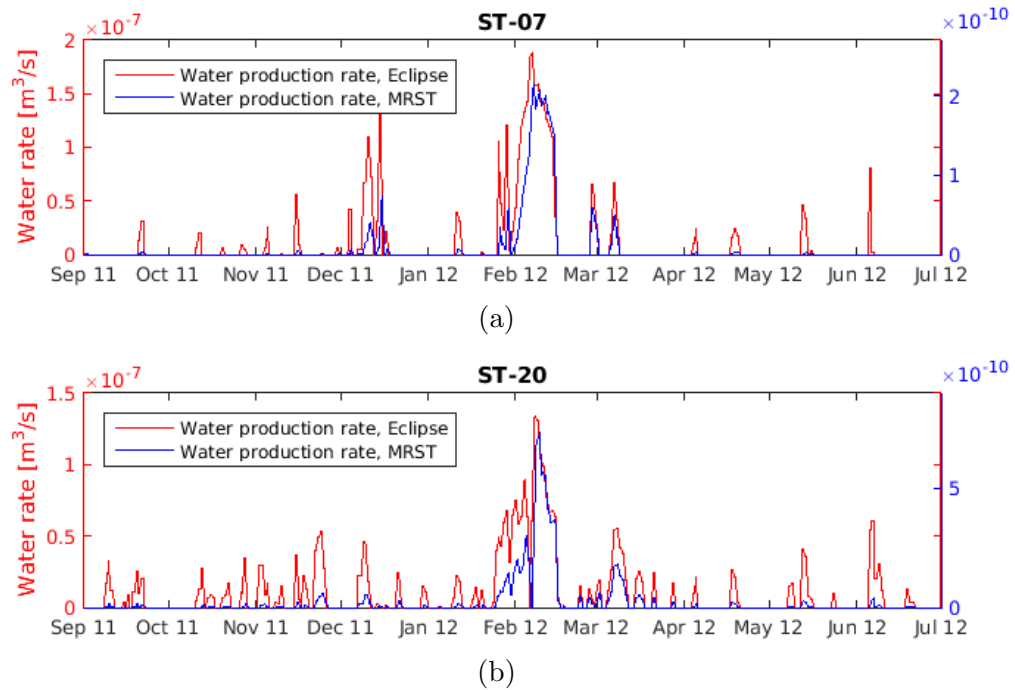


Figure 5.8: The water production rates calculated for wells ST-07 and ST-20 in Simulation A. The results from Simulation A are plotted in blue, and their ECLIPSE counterpart is plotted in red.

implementation of scanning curves in ECLIPSE, the reader is referenced to the ECLIPSE Technical Description.

To conclude, the MRST simulations used a highly simplified scheme to calculate the correct relative permeability curves, and it is highly likely that this simplification is one of the major causes of the discrepancy between the two simulation results. However, if we go back and look at the accuracy of the ECLIPSE simulation results in predicting the historic pressure, this error does not seem so troublesome after all. All things considered, it was therefore concluded that although the MRST simulation results were not perfect, large portions have been implemented correctly, and the results were good enough.

Figure A.3 show the predictions of water production from each well, as they were found in Simulation A, plotted against the corresponding water production rate in ECLIPSE. Wells ST-07 and ST-20 were again chosen as representative examples, and were therefore also plotted in Figure 5.8.

The overall validity of these results was drawn into question after examining the saturation around the well perforations for different time steps. Figure 5.9 shows

the saturation in a cross section of the local grid refinement around Well ST-20 for assorted time steps, along with the location of the well and its perforation. From the figure, it seems highly improbable that well ST-20 should be producing water at all, as the location of the gas-water interface is always a minimum of five meters away from the well perforations.

It is possible that the predicted water production rate is simply caused by the residual error allowed by the solver. This theory is strengthened by noticing the very small scales of the water rate found in the MRST simulation results. There large factor discrepancy between the ECLIPSE and MRST simulation results can possibly be explained by noting that the VFP table can alter the water production rates as along with the well pressures. It is feasible that this table has been set up so as to increase the scale of the predicted water rates, with the intent of bringing them closer to the historically recorded values.

We noted in the previous section that there was close to no correlation between the historical and simulated water production rates. We have so far gotten no further on the way to achieve successful predictions of water production, only sown doubt about the validity of the results that we already have. We will therefore conclude this section by venturing two theories for why the simulated results are so far off from the historical values. One possibility is that the square, uniform grid refinements made around the wells cause erroneous numerical results. It might improve the situation to replace this grid refinement with one that gets finer and finer the closer the cell is to the well.

Alternatively, it is possible that the water production is due to some process that occurs on a smaller scale than what we use in the Stenlille model. As seen in Figure 5.9, the location of the gas-water interface was found to lie consistently several meters below the well perforations in Simulation A. Imagine now that there exists cracks or other small pathways in the rock where water is allowed to travel nearly uninhibited, so that water could be "smuggled" above the gas-water interface. This effect would not be reproduced in the Stenlille model, where the permeability of is considered constant in each cell, and the cells around the wells have the rough dimension of $10 \times 10 \times 1$ cubic meters.

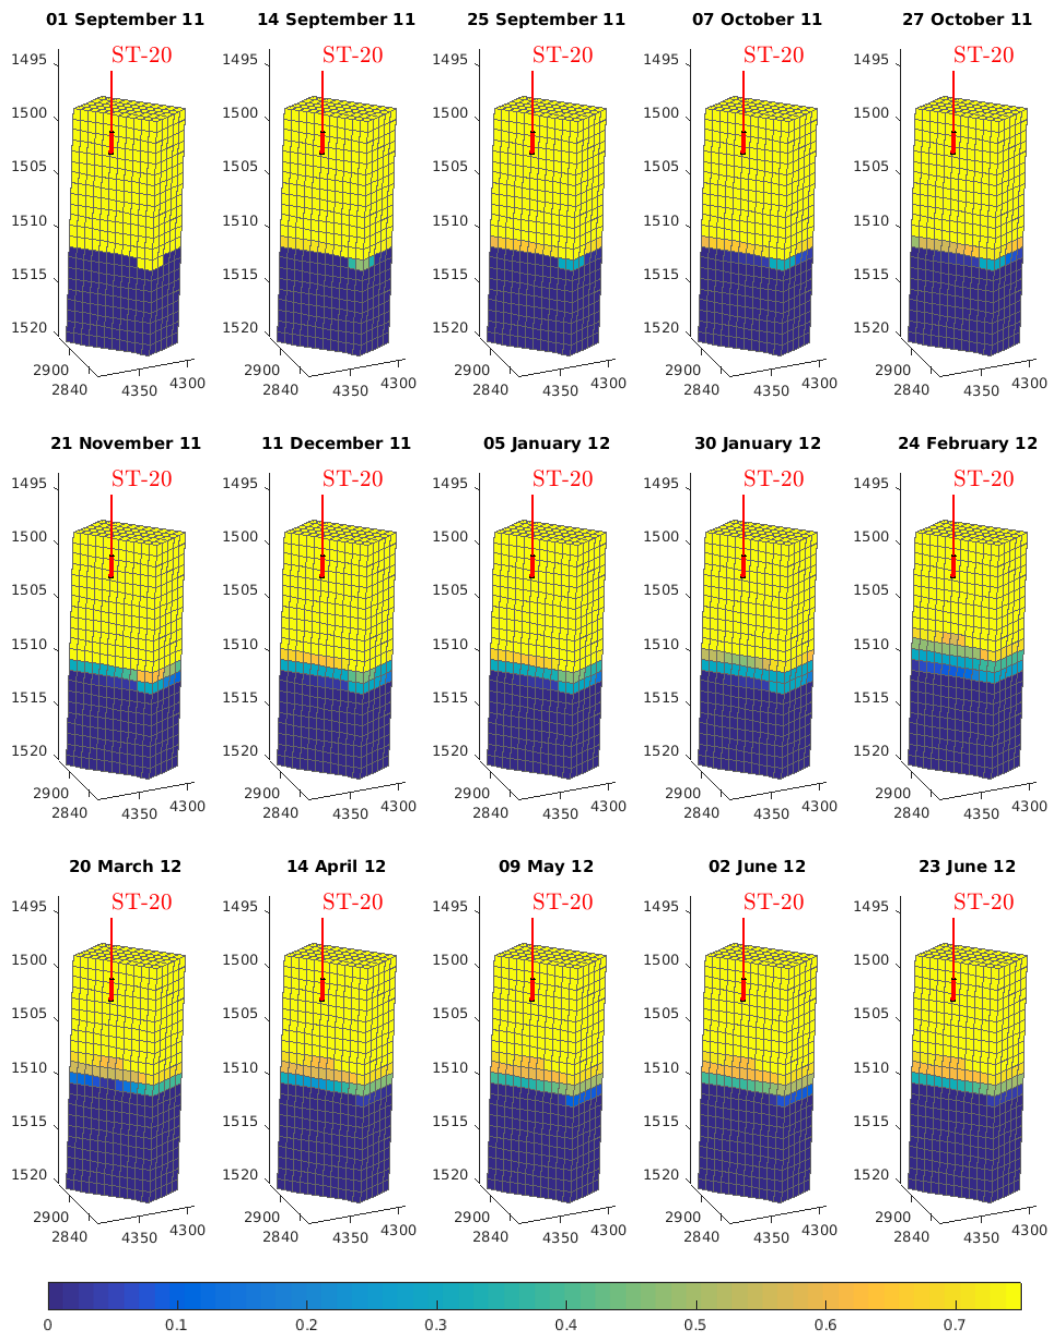


Figure 5.9: The saturation profile around Well ST-20 for different time steps, visualised using a cross section of the local grid refinement around the well. The local grid refinement itself can be seen more clearly in Figure 4.10.




Figure 5.10: The initial gas saturation used in Simulation B.1. The gas-water interface was assumed to be flat and located at a height of $z_{\text{water}} = 1510.95$. Residual water was assumed to be trapped in the gas plume, yielding a gas saturation of $s_n = (1 - s_{w,r})h/H$

5.3 Simulation B.1: VE Model of Stenlille

The purpose of this simulation is to examine the validity of using a VE model for Stenlille, by running the VE equivalent of Simulation A and comparing the results.

Simulation Set Up:

The two-dimensional model of Stenlille was set up as described in Chapter 4.2, with the top-surface grid shown in Figure 4.11 and the rock properties shown in Figure 4.12 and 4.13. The simulations used the fluid properties shown in Figure 4.1, with the alterations described in Chapter 4.2.

The gas-water interface was assumed to be flat and located at a depth of 1510.95 meters, with residual water trapped in the gas column, exactly as in Simulation A. The initial saturation this resulted in for the VE simulation can be seen in Figure 5.10.

The initial pressure was set similarly as in Simulation A, using the method outlined Chapter 4.1.4, with the exception that the reference pressure was assumed to be slightly higher. The reference pressure was in this case set to $p_{\text{ref}} = 160.9$ bar at a reference depth of $z_{\text{ref}} = 1500$ m.

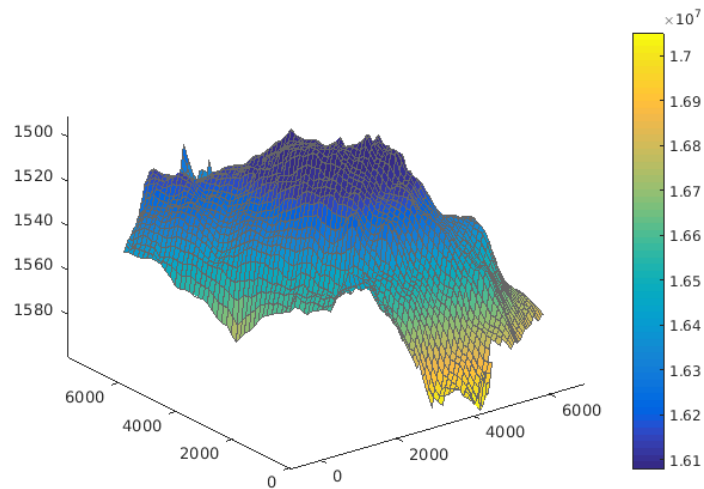


Figure 5.11: The initial pressure used in Simulation B.1. The aquifer was assumed to be in vertical equilibrium at the start of the simulation, with a reference pressure $p_{\text{ref}} = 159.9$ bar at a reference depth of $z_{\text{ref}} = 1500$ m.

Comparison and Results:

The full set of results for this simulation is given in Appendix B. Figures B.1 and B.2 in this appendix show the bottom-hole pressures found for each well, plotted together with the results from running the full three-dimensional model. Figure B.3 shows the predictions for the water production rates found for this well, also plotted together with the results from running the three-dimensional model. The simulation took 52 minutes to run, which is a considerable improvement upon the run-time experienced in Simulation A.

Figure 5.12 shows the well pressure results for wells ST-08 and ST-16, which were chosen as representative examples for this discussion. The simulated well pressures were found to agree very nicely with the ones from the black-oil simulation, especially during production season. In fact, the simulated pressures for well ST-08 made a near perfect match with each other.

In the cases where the well pressures made a less favourable match, the VE simulated pressure generally overreacted compared to the pressure found in Simulation A. Well ST-16 was included to showcase the worst match made by all the wells, its well pressures were found to overreact with a factor of two compared with the non-VE results.

If the reader were to examine the well pressure results for the observation wells, as they are given in Figure B.2 in Appendix B, the results might initially seem




Figure 5.12: The bottom-hole well pressures calculated in Simulation B.1 for wells ST-08 and ST-20. The results from Simulation B.1 are plotted in magenta, and the corresponding results from Simulation A are plotted in blue.

troubling. However, the discrepancy between the VE and non-VE results can be easily explained. They clearly follow the same trend with decent accuracy, but are separated by a near constant difference of roughly 10 bar. This is due to the fact that the observation wells have their reference pressure set at a much deeper reference depth than the other wells. Since the top-surface grid contains only one cell per column, the well model is unable to adjust the pressure to fit the reference depth, leading to a near constant pressure difference between them.

Figure 5.13 shows the water production predicted by each well in the VE simulation, together with the rate that was predicted in Simulation A. From the figure, it is easy to see that they follow the same trend, but that the VE results overreact with a factor of 1000. While this may seem dramatic, it is not unexpected, and is

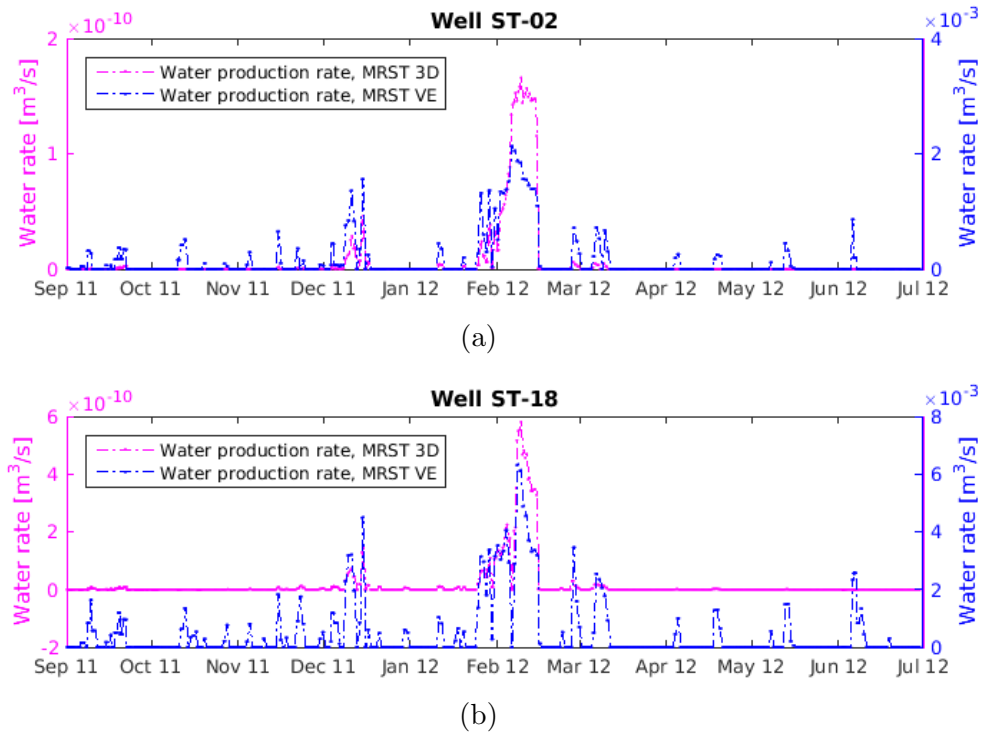


Figure 5.13: The water production rates calculated for wells ST-07 and ST-20 in Simulation B.1. The results from Simulation B.1. are plotted in magenta, and their ECLIPSE counterpart is plotted in blue.

due to a shortcoming in the well model utilized for the VE simulation. The water saturation variable in the VE model is interpreted as the fractional height of the water column, and all columns are at least partially filled with water. From the viewpoint of the well model, all cells in the grid therefore contain water, which leads to water being produced in significant quantities.

5.4 Simulation B.2: VE Simulations on Refined Grid

Recall the simulated water rates found in Simulation B.1. We concluded that these rates were completely off, due to a shortcoming in the well model used for the simulation. Recall also that the shortcoming was specifically caused by the standard well model misinterpreting the meaning of the saturation values in each cell. Since the VE model operates with the fractional water column height as the water saturation variable, and all cells contain at least some water, the well model produces considerable amounts of water along with gas.

A simple workaround for this problem would be to examine the height of the gas plume in each cell, and only allow water production if the gas-water interface is at a certain height. However, the grid used in Simulation B.1 is quite coarse, with the cells each spanning an area of roughly 100×200 square meters. In order to capture the dynamics of the movement of the gas water interface, a higher resolution of the solution is needed. This leads us to the topic of this simulation, which is VE simulations on a refined grid.

The construction of a refined top-surface grid was not a straightforward process, as MRST does not so far support grid refinement, although it contains a whole module for grid coarsening. This ties in nicely with the observation we have made earlier, about how long simulation run-times are a common issue in reservoir simulation. Routines and methods for grid coarsening and upscaling have therefore been given higher priority in MRST, than routines and methods for grid refinement. This problem was solved by creating a custom method for constructing a refined top-surface grid.

Refinement of the Top-Surface Grid

The method created for refining a top-surface grid is outlined in the form of a flow-chart in Figure 5.14. At first glance this method might seem unnecessarily circuitous, but it was nevertheless found to be successful, as it made efficient use of pre-existing functions in both Matlab and MRST.

First, the cell centroids (x,y) , depths z and formation heights H were extracted from the unrefined grid. Then, the `scatteredInterpolant` function in Matlab was used to create interpolations $z(x,y)$ and $H(x,y)$ of the top and bottom surfaces of the aquifer. The centroids of the cells in the refined grid were then calculated, and the interpolations z and H evaluated at these points.

At this point, all the information necessary for an Atlas file has been created. Here, the term Atlas file is used to refer to the file format employed in the CO₂

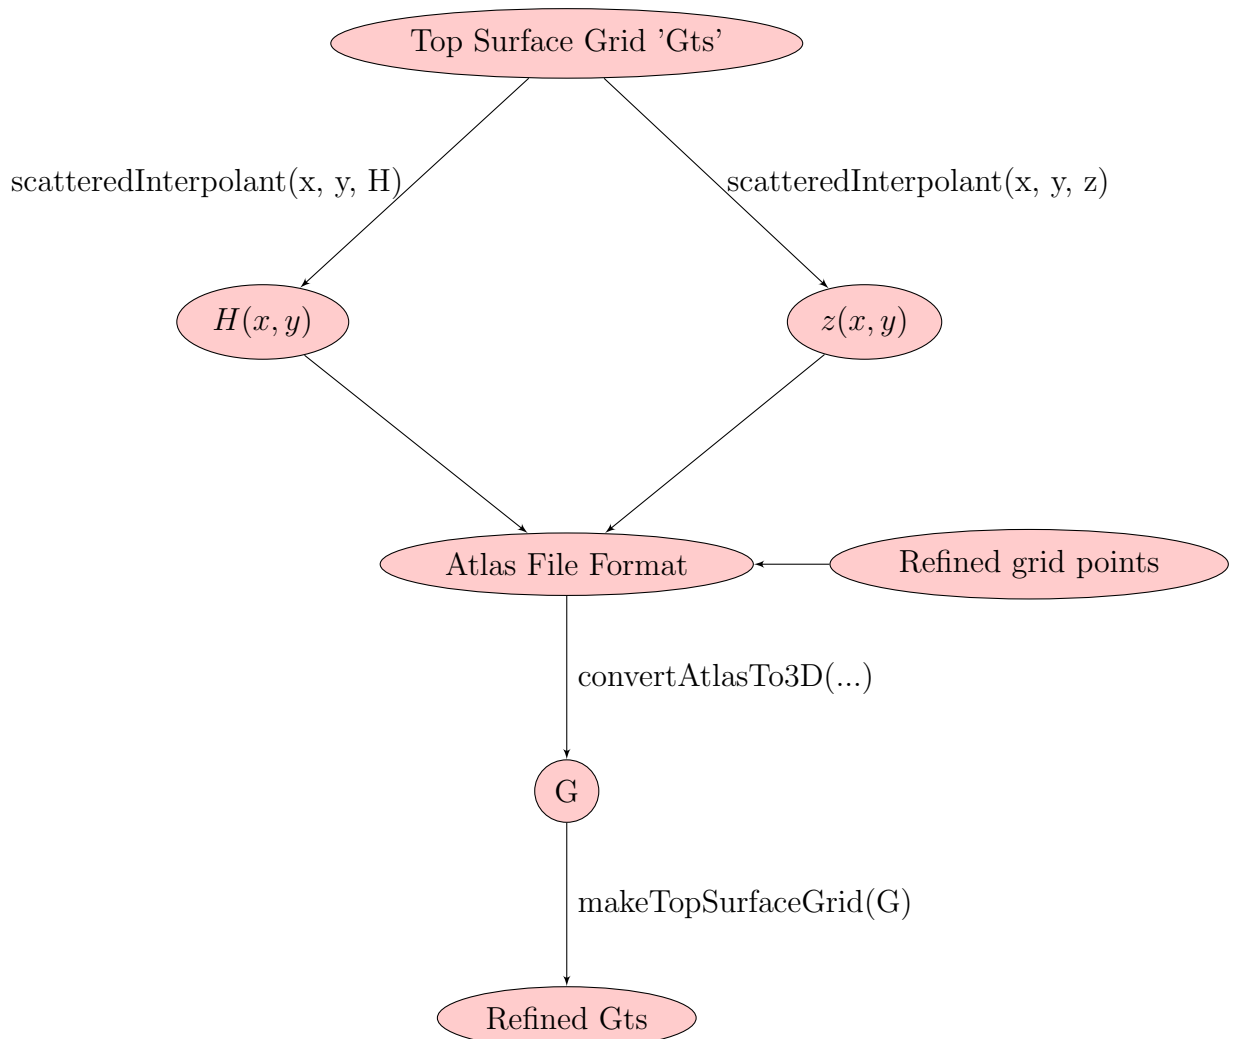


Figure 5.14: A flow-chart showing how a top-surface grid can be refined. The variables (x,y) , H , and z are cell-wise centroid, depths and formation heights respectively. Using a scattered interpolant, interpolations $z(x,y)$ and $H(x,y)$ can be created of the top-surface grid and cell heights. By evaluating the interpolations $H(x,y)$ and $z(x,y)$ at the centroids of the cells in the desired refined grid, all the ingredients necessary for an Atlas file are constructed. The function `convertAtlasTo3D()` can then be used to generate a three-dimensional, refined grid from this information. Finally, a refined top-surface grid is created by from the three-dimensional refined grid.

Storage Atlas created by the Norwegian Petroleum Directorate [6]. It was created with the goal of furthering research on CO₂ storage, and contains an open-source atlas of the height-maps and depths for potential storage sites in the North Sea.

The final step is to construct the grid itself, and this can be done using a function in the `co2lab` module. The function `convertAtlasTo3D(..)` generates a three-dimensional grid from an atlas file, using its data points for the top and bottom surface of the aquifer. Lastly, the refined top-surface grid can be constructed from the refined three-dimensional grid.

If each cell in the top-surface grid of Stenlille is refined say 10 times in each direction, the number of cells in this refined grid will rival that of the three-dimensional grid used in Simulation A. Consequently, the run-times of simulations using this grid will also rival the run-time of Simulation A. The run-time could be reduced again by re-coarsening the refined grid, using the functionality implemented in the grid coarsening module in MRST [3].

The result of refining the Stenlille top-surface grid with a factor of 5 in each direction, and then re-coarsening all cells that are initially filled with water, is shown in Figure 5.15. The three-dimensional grid has to be plotted instead of the top-surface grid, as the the grid visualisation function in MRST does not work well for the very specialized case of a coarsened top-surface grid.

Simulation Set Up:

The simulation was set up exactly as explained and shown for Simulation B.1. The initial conditions and fluid model required no alterations from before. The refined rock model was constructed by assigning the porosity and permeability of each cell the same value assigned to their coarse mother cell. The wells were assigned to the refined cell found in the middle of the coarse mother cell.

Comparison and Results:

Results showing the location of the gas-water interface for different time steps can be seen in Figure 5.16. The results did not turn out as expected. The gas-water interface was expected to undulate up and down as the wells injected or extracted gas from the aquifer. Some undulation was seen, especially when the wells were injecting gas. However, these undulations were small compared to other disturbances of the gas-water interface, which were especially large close to the boundary. It was concluded that these peculiar results must be caused by one or several errors in the simulation.

After some investigation, it was discovered that the initialization of a initial flat gas-water interface was problematic, as the pressures in the VE model equations are defined to be the pressure at the top surface. Since the top-surface itself is not




Figure 5.15: The result of refining the top-surface grid of Stenlille 5 times in each direction, and then re-coarsening the cells that were initially filled with water. The three-dimensional version of the top-surface grid, with only one cell in the vertical direction, has to be used for plotting purposes instead of the top-surface grid. This is because the grid visualisation function in MRST does not work for a top-surface grid that has been coarsened.

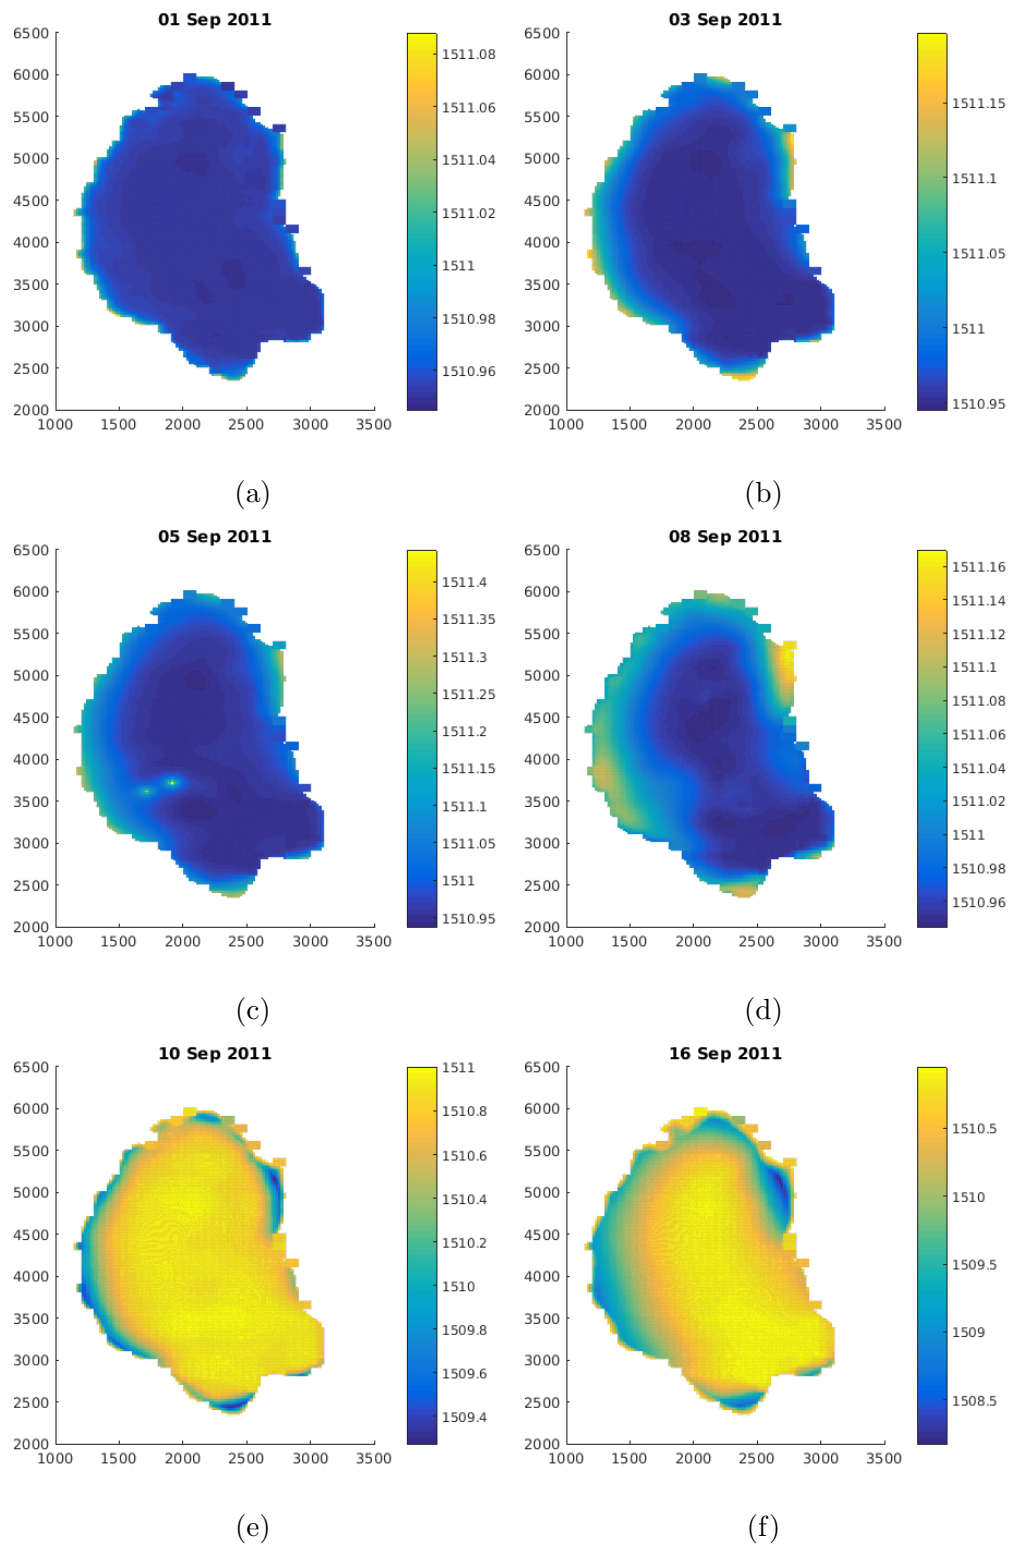


Figure 5.16: The vertical depth of the gas-water interface found by Simulation B.2, shown for assorted time steps.


flat, it was therefore impossible to construct a static, flat gas-water interface.

To verify this, Simulation B.2 was run with all well rates set to zero, to see what equilibrium it would settle into. The pressure this resulted in at the final time step, for all cells that initially contained gas, is shown in Figure 5.17a. The depth of these same cells are shown in Figure 5.17b, and it is clear they correlate perfectly. However, the perturbation of the interface depths did not correlate perfectly with the depth of the top-surface grid, so this cannot be the only issue. It was concluded that there must be another effect at play that warps the results, but it is so far unclear what this effect is.

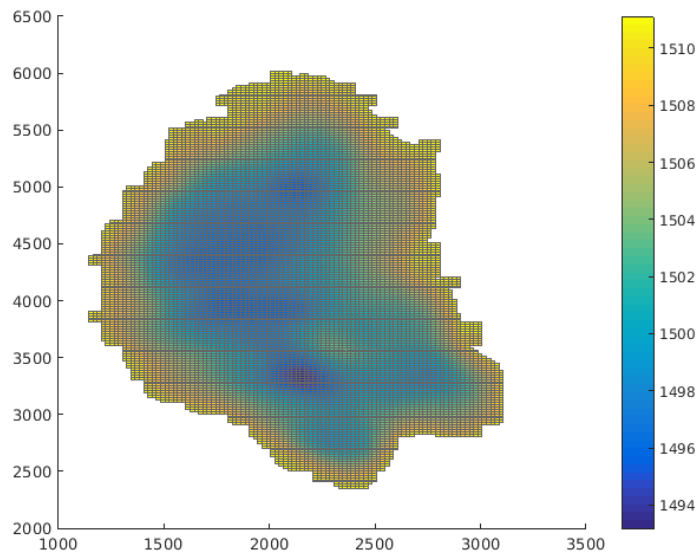
Recall that two assumptions were made when developing the VE equations: First that of vertical equilibrium, and then that of a sharp interface between gas and water. If any of these assumptions can be shown to be wrong, that can go a long way in explaining the peculiar simulation results. Figure 5.18 shows a scatterplot of the fractional value $v_{\text{ver}}/v_{\text{hor}}$ for each time step and each well. Here, v_{ver} was chosen to be the flow velocity through one of the horizontal faces on the cell that contains the bottom-most well perforation, and v_{hor} is the flow velocity through one of the vertical faces on the same cell. The motivation behind comparing the flow velocities through these faces specifically is that the vertical equilibrium assumption will hold the least true for the cells containing well perforations. By examining the figure, it is clear that the vertical flow is generally small compared to the horizontal flow, and that the vertical equilibrium assumption is valid at least for the majority of time steps.

From [7] we have an equation for estimating the segregation time scale, which is necessary for the sharp interface assumption. The time scales for achieving a sharp interface was calculated for the cells containing the wells, and the results are summarized in Table 5.2. The resulting time scales were generally found to be much higher than the time-scales of the simulation steps, which were all smaller than one day. It was concluded that the assumption of a sharp interface is likely invalid, and that this might be the cause of the peculiar simulation results for the location of the gas-water interface.

Both the problems we have found can potentially be fixed. If the pressure term in the VE model equations were taken to be the pressure at the gas-water interface instead, it would be possible to create the flat interface that was desired in the VE simulations. The problem with the assumption of a sharp interface could be remedied by taking into account the effect of a *capillary fringe* between the two fluid phases, i.e., a smooth transition zone between the two [13].



(a) The pressure at the top-surface cells that contain gas, at the end of a simulation that used the same initial conditions but with the well rates set to zero.



(b) The vertical depth z of the cells in the refined top-surface grid that contained gas at the start of the simulation.

Figure 5.17: The result after running a simulation on the refined top-surface grid, that utilized the same initial conditions as before, but had all well rates set to zero. It is clear that the aquifer then settled into an equilibrium where the pressure at the top-surface grid correlated with the vertical depth of each cell, even though it was not initialized with this pressure.




Figure 5.18: The fractional value $v_{\text{hor}}/v_{\text{ver}}$ for each time step and each well, where v_{ver} is the flow velocity through one of the horizontal faces on the cell that contains the bottom-most cell perforation, and v_{hor} is the flow velocity through one of the vertical faces on the same cell.

Table 5.2: A table showing the time scale of fluid segregation t_s , for the cells containing the perforations for each well.

Well	t_s [Days]
ST-02	67
ST-07	117
ST-08	159
ST-13	92
ST-14	92
ST-16	129
ST-17	136
ST-18	147
ST-19	65
ST-20	85

Chapter 6

Conclusions

In this thesis, we have simulated fluid flows in an aquifer used for gas storage, during injection and production of said gas. We have utilized two different models for the simulations; first a standard black-oil model, and then a simplified VE model. The former was constructed by assuming vertical equilibrium and then integrating the black-oil model equations in the vertical direction.

The models were tested by examining their success in simulating gas storage in the Stenlille aquifer, in the time period 31.07.2011-01.06.2012. The goal was to improve upon the results from earlier simulations performed in ECLIPSE. More specifically, the goal was to reduce run times, and improve upon the ability of the simulations to predict well pressures and water production rates.

To this end, the ECLIPSE model of Stenlille was transferred into the MRST framework, where it was easier to experiment with the models themselves. Simulations were then carried out using both the black-oil model and the VE model. The results were validated by comparing them against historical data, the results from other simulations, or a combination of the two.

Results of Using the Black-Oil Model

One simulation was run using this model, with the purpose of verifying that the MRST model of the Stenlille had been implemented correctly.

The simulation results were compared with the results from an ECLIPSE simulation run over the same scenario. The simulated well pressures were found to match reasonably well with each other, with the exception of a factor difference in the well pressure response to flow rates. It was theorized that this factor difference was caused by a simplified method used in the MRST model to determine the relative permeability of each fluid phase. When comparing the discrepancy between

the two simulation results for the well pressure, against the discrepancy between simulated pressures and historical values, the model was nevertheless concluded to have been implemented with sufficient accuracy.

The MRST simulation results for the water production rates of each well were found to match poorly with the corresponding ECLIPSE simulation results. However, after examining the saturation values around the wells, the validity of the ECLIPSE results themselves were drawn into question. The gas-water interface was seen to stay consistently below the well perforations, in which case there should be no mobile water in the well cells. It was therefore suggested that the non-zero water production rates were solely caused by the residual error allowed by the solver.

Furthermore, the simulations results for the water production rates were found to bear no resemblance to the values that have been measured historically. Two theories were ventured to explain this discrepancy; (1) that rectangular grid refinement around the wells produced erroneous numerical results, or (2) that the water production was due to a phenomenon that occurred on a much smaller scale than the grid cells.

Results of Using the VE Model

Two simulations were running use the VE model. The first was run with the intent to verify that a VE model could be successfully used for Stenlille. It used the two-dimensional versions of the grid, rock model, fluid model and wells from the black-oil simulation. The simulation results this yielded for the well pressures were then compared with those found in the black-oil simulation, and they were found to match each other reasonably well. It was concluded that the VE model can be successfully used to simulate the well pressures in Stenlille. This method offered a substantial reduction in the run-time of the simulation, reducing it to less than 10% of the run-time for the black-oil simulation.

The VE simulations were found to give meaningless results for the water production rate of each well, due to a misinterpretation by the well model of what the saturation values signify. A fix for this problem was proposed, that involved checking if the gas-water interface had reached the well before allowing it to produce water.

Since the top-surface grid used in the first simulation was too coarse to capture the movements of the gas-water interface, a method to refine the grid was proposed. The grid was then refined in the areas containing gas at the first time step, and a VE simulation was carried out using this refined grid.

This simulation was deemed unsuccessful, as it produced highly peculiar results

for the depth of the gas-water interface at each time step. One problem was identified, concerning the VE pressure being defined at the top surface, which made it impossible for the gas-water interface to lie flat at equilibrium. This problem can potentially be fixed by changing the pressure used in the VE equations, so that they are defined to be the pressure at the gas-water interface instead.

However, it was then concluded that this problem could not fully account for the peculiar simulation results. Therefore, the assumptions that went into the VE model were examined, in order to say something about their validity. It was found that the assumption of negligible vertical flow was decent, but that the assumption of a sharp interface was problematic. It was suggested that the simulation results could be improved by introducing a capillary fringe between the two fluid phases.

Bibliography

- [1] Key world energy statistics 2015. http://www.iea.org/publications/freepublications/publication/KeyWorld_Statistics_2015.pdf May 2015.
- [2] The matlab reservoir simulation toolbox, version 2015a (july 2016). <http://www.sintef.no/MRST/>.
- [3] Sintef ict, the matlab reservoir simulation toolbox: Grid coarsening module (october 2012). <https://www.sintef.no/projectweb/mrst/modules/agglom/>.
- [4] Sintef ict, the matlab reservoir simulation toolbox: Numerical co2 laboratory (october 2014). <http://www.sintef.no/co2lab>.
- [5] Norwegian Petroleum Directorate. Fakta 2014, norsk petroleumsverksemd. http://www.npd.no/Global/Norsk/3-Publikasjoner/Faktahefter/Fakta2014/Fakta_2014_NO_nett_.pdf May 2016.
- [6] Norwegian Petroleum Directorate. Co2 storage atlas north sea. <http://www.npd.no/en/Publications/Reports/C02-Storage-Atlas-/July2016>, 2006.
- [7] Bo Guo, Karl Bandilla, Florian Doster, Eirik Keilegavlen, and Michael Celia. A vertically integrated model with vertical dybanics for co2 storage. *Water Resources Research*, 2014.
- [8] Donald Katz. *Underground Storage of Fluids*. Ulrich's Books, Inc., 1968.
- [9] John Killough. Reservoir simulation with history-dependent saturation functions. *Trans. AIME*, 1976.
- [10] Stein Krogstad, Knut Andreas Lie, Olav Møyner, Halvor Nilsen, Xavier Raynaud, and Bård Skaffestad. Mrst-ad – an open-source framework for rapid prototyping and evaluation of reservoir simulation problems. *SPE Reservoir Simulation Symposium, 23–25 February, Houston, Texas*, 2015.


- [11] Troels Laier. Perspectives for monitoring fugitive gases – 25 years of experience from stenlille natural gas underground storage. www.atv.dk/category/69-skifergas?download=400 July 2016.
- [12] Troels Laier and Hans Øbro. Environmental and safety monitoring of the natural gas underground storage at stenlille, denmark. *Geological Society, London, Special Publications*, 2009.
- [13] Knut Andreas Lie, Halvor Nilsen, and Odd Andersen. Fully-implicit simulation of vertical-equilibrium models with hysteresis and capillary fringe. *Computational Geosciences*, 2015.
- [14] Knut-Andreas Lie, Halvor Møll Nilsen, and Odd Andersen. *An Introduction to Reservoir Simulation Using MATLAB*. SINTEF ICT, 2015.
- [15] Knut-Andreas Lie, Halvor Møll Nilsen, and Odd Andersen. Robust simulation of sharp-interface models for fast estimation of co2 trapping capacity in large-scale aquifer systems. *Computational Geosciences*, 2015.
- [16] Ingeborg Skjelkvåle Ligaarden, Marit Kjøsnes Natvig, Knut Andreas Lie, Bård Skaflestad, Halvor Nilsen, and Odd Andersen. A open source matlab implementation of consistent discretisations on complex grids. *Computational Geosciences*, 2012.
- [17] Artem Napov and Yvan Notay. An algebraic multigrid method with guaranteed convergence rate. *SIAM Journal on Scientific Computing*, 2012.
- [18] Halvor Møll Nilsen, Odd Andersen, and Sarah Eilee Gasda. Vertically averaged equations with variable density for co2 flow in porous media. *Transport in Porous Media*, 2014.
- [19] Jan Martin Nordbotten and Michael A. Celia. *Geological Storage of CO2: Modeling Approaches for Large-Scale Simulation*. Wiley, 2011.
- [20] Yvan Notay. An aggregation-based algebraic multigrid method. *Electronic Transactions on Numerical Analysis*, 2010.
- [21] Intergovernmental Panel on Climate Change (IPCC). Climate change 2014: Mitigation of climate change. Cambridge University Press, 2014.
- [22] Donald W Peaceman. Interpretation of well-block pressures in numerical reservoir simulation with nonsquare grid blocks and anisotropic permeability. *Society of Petroleum Engineers*, 1983.
- [23] Yousef Saad. *Iterative Methods for Sparse Linear Systems, Second Edition*. Society for Industrial and Applied Mathematics, 2003.

- [24] Schlumberger. Eclipse file formats reference manual, 2007.
- [25] Schlumberger. Eclipse reference manual, 2007.
- [26] Schlumberger. Eclipse technical description, 2007.
- [27] Elizabeth Spiteri, Ruben Juanes, Martin Julian Blunt, and Franklin Orr Jr. A new model of trapping and relative permeability hysteresis for all wettability characteristics. *Electronic Transactions on Numerical Analysis*, 2010.
- [28] An Interdisciplinary MIT Study. The future of natural gas. http://mitei.mit.edu/system/files/NaturalGas_ExecutiveSummary.pdfMay2016, 2011.
- [29] Øystein Pettersen. Lecture notes: Basics of reservoir simulation with the eclipse reservoir simulator. http://folk.uib.no/fciop/index_htm_files/ResSimNotes.pdfMay2016, 2006.


Appendix A

Results of Simulation A


A.1 Well Bottom-Hole Pressures




(a)




(b)




(c)




(d)




(e)




(f)




(g)



(h)



(i)



(j)

Figure A.1: Well pressure results found in Simulation A, for each of the injection/production wells in Stenlille. The bottom-hole pressures found in the Simulation A are plotted in blue, and the bottom-hole pressures found from a similar ECLIPSE simulation are plotted in red.








Figure A.2: Well pressure results found in Simulation A, for each of the observation wells in Stenlille. The bottom-hole pressures found in the Simulation A are indicated using a blue curve, and the bottom-hole pressures found from a similar ECLIPSE simulation are indicated in red.


A.2 Water Production




(a)




(b)




(c)




(d)




(e)




(f)




(g)



(h)



(i)




(j)

Figure A.3: Water production rates from Simulation A. The water production rates found in the Simulation A are plotted in blue, and the water production rates found from a similar ECLIPSE simulation are plotted in red.


Appendix B

Results of Simulation B


B.1 Bottom hole pressures




(a)




(b)




(c)




(d)




(e)




(f)




(g)



(h)



(i)



(j)

Figure B.1: Well pressure results found in Simulation B.1, for each of the injection/production wells in Stenlille. The bottom-hole pressures found in the Simulation B.1 are plotted in magenta, and the bottom-hole pressures found from Simulation A are plotted in blue.








Figure B.2: Well pressure results found in Simulation B.1, for each of the observation wells in Stenlille. The bottom-hole pressures found in the Simulation B.1 are plotted in magenta, and the bottom-hole pressures found from Simulation A are plotted in blue.


B.2 Water Production




(a)




(b)




(c)




(d)




(e)




(f)




(g)



(h)



(i)




(j)

Figure B.3: The predicted water production rates from Simulation B.1. The water production rates found in the Simulation B.1 are plotted in blue, and the water production rates found in Simulation A are plotted in magenta.


Appendix C

Other Plots


C.1 Individual Well Schedules




(a)




(b)




(c)




(d)




(e)




(f)




(g)



(h)



(i)




(j)


Figure C.1: The injection/production rates for the wells in Stenlille, given for the time period 31.07.11-30.06.12

C.2 Comparison of ECLIPSE Simulation Results and Historical Values


C.2.1 Tubing Head Pressures







(a)




(b)




(c)






(h)




(i)




(j)

Figure C.2: Well pressure results found in an ECLIPSE simulation, for each of the injection/production wells in Stenlille. The simulation was over the same scenario as the ones used in Simulations A, B.1 and B.2. The tubing hole and bottom-hole well pressures from ECLIPSE are plotted in solid blue and dashed blue respectively. The historically measured tubing head pressure is plotted in black.


C.2.2 Water Production Rate




(a)




(b)




(c)




(d)




(e)




(f)




(g)



(h)



(i)



(j)

Figure C.3: The water production results for an ECLIPSE simulation, for each of the injection/production wells in Stenlille. The simulation was over the same scenario as the ones used in Simulations A, B.1 and B.2. The predicted water production from the ECLIPSE simulation is shown in dashed blue, and the historically recorder water production is shown in black.