PROCEEDINGS OF SCIENCE

Computing Challenges in Oil and Gas Field Simulation

Jeroen C. Vink [1]
Shell International Exploration and Production B.V.
Kesslerpark 1, Rijswijk, The Netherlands
E-mail: Jeroen.Vink@Shell.com

Exploitation of oil and gas fields is becoming increasingly more difficult and expensive. Many existing fields start to decline and enhanced recovery techniques are required to improve production and ultimate recovery and reduce impact on the environment. New fields often are in remote and (environmentally) hostile locations, the reservoir rock is geologically complex (heterogeneous, fractured and faulted) and the hydrocarbon fluids are difficult, for example, oil is ultra heavy, gas is sour or otherwise polluted.

To mitigate the risk of developing such hydrocarbon reservoirs, large scale computer simulations are invaluable. Current simulation techniques, which were primarily developed for the more simple reservoirs of the past, have to be extended and improved: Much larger models, with much more geological detail, have to be simulated; Fluid chemistry and thermodynamics must be captured in more detail; Uncertainty ranges in the simulated results have to be estimated reliably.

This presentation, discusses current simulation challenges and techniques, as well as emerging computing strategies that are being developed in the oil and gas industry.

XI International Workshop on Advanced Computing and Analysis Techniques in Physics Research
Amsterdam, the Netherlands
23-27 April, 2007

[1] Speaker

PoS(ACAT)007

1. Introduction

Currently there is a lot of activity in reservoir simulation research. Easily producible oil and gas is on the decline and remaining hydrocarbon reserves require increasingly more complex production methods. In order to maximize the efficiency of these production methods, new techniques must be developed and a better understanding of the subsurface oil and gas dynamics is required. Reservoir flow simulation is an important tool for understanding the subsurface better: to investigate novel production schemes and technique and to optimize the oil and gas field development plans.

In this brief overview of current trends in reservoir simulation research at Shell and elsewhere, it will not be possible, of course, to aim for completeness. Nor will it be possible to give a comprehensive account of the physics and mathematics of oil field simulation. For introductions to this field, see e.g. [1-4]. Here I will focus on new directions and trends, which I will elaborate on in section 4, after briefly discussing some characteristics of oil and gas fields and the reasons for doing reservoir simulations in section 2, and giving a flavor of the reservoir flow equations and current computing efforts in section 3. Section 5 contains some concluding remarks and speculations for the future.

2. What are oil and gas fields and why simulation?

Oil and gas are hydrocarbon fluid accumulations in subsurface rock. They are the remains of organic material that was deposited millions of years ago in swamps, river deltas, sea lagoons etc. and subsequently covered by sediments. In the course of time, the organic material transformed ('cooked') into oil and gas and became the source material for our present oil and gas fields. Since oil and gas is lighter than water, these fluids have a tendency to seep up through the rock towards the surface. This migration process is possible because rock is actually permeable to fluid flow: it consists of solid grains with space in between. This space, called rock porosity, in some cases can be as large as 20-30% of the bulk rock volume (see fig. 1).

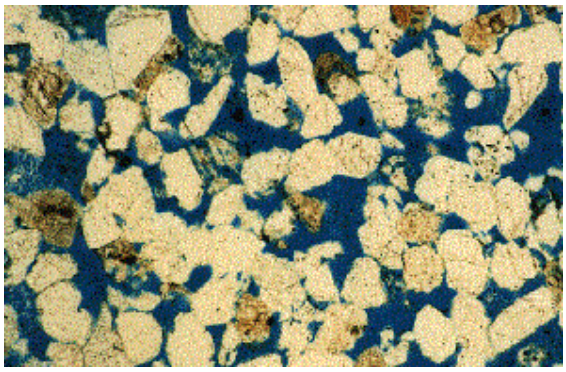

***Fig. 1*: Close-up of reservoir rock; the dark bluish color is the pore space between the grains.**

Oil and gas (but also water) can flow through this connected network of inter-grain pores in the rock. However, the process is slow, because most rock is not very permeable, since the pore space is small or the pores are badly connected. Even with a very slow migration upward, there would be no hydrocarbons in our present time, if there would not be traps for the upward moving fluids. Such traps typically consist of layers of impermeable dense and compact rock. Fig. 2 shows two types of such traps: The first is an anticline, a dome-shaped sealing layer with permeable rock underneath. The second trap is slightly more complex, and requires a fault, where the rock layers are broken and shifted.

The first task for any (successful) oil company is to find such reservoirs of trapped oil and gas. To some extent this relies on a combination of clever guessing, delicate probing techniques and trail and error. The probing one can do into the subsurface is fairly limited, because it should be quick, reasonably cheap and penetrate very deeply. Using acoustic waves, seismic data acquisition, is currently the most

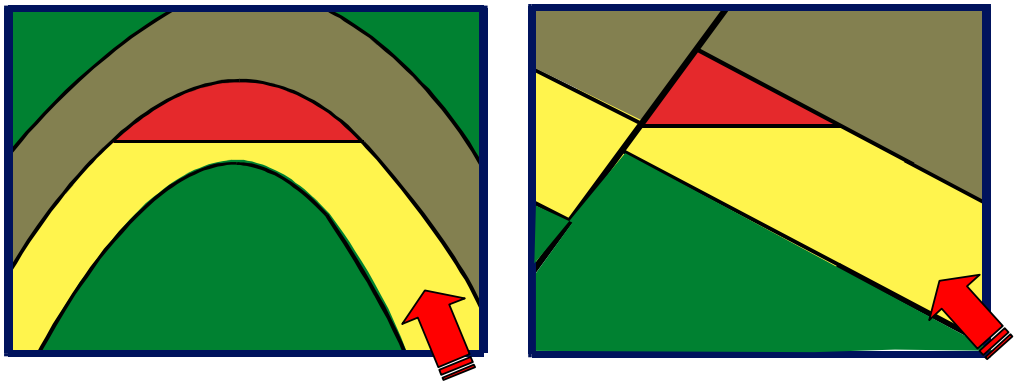

Fig.2: **Two types of oil trap: Left is an anticline, right a fault. The grayish layer is the sealing cap rock, yellow the permeable sandstone and red is the trapped oil.**

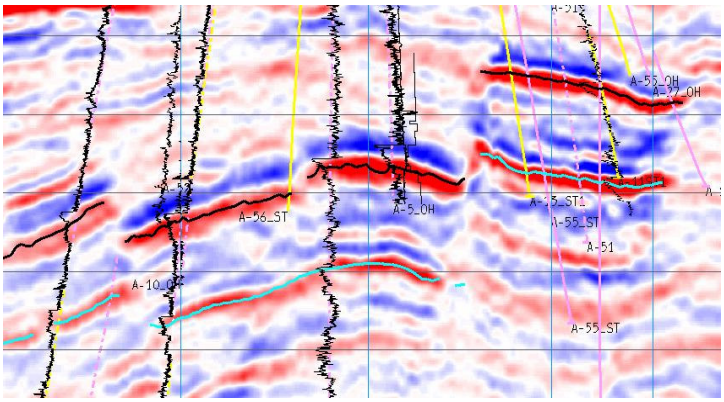

Fig.3: **Seismic image of the subsurface. The colors indicate layers with different (acoustic) properties.**

widespread method. Here one sends sound blasts into the rock and records the echoes in arrays of sensitive microphones. These seismic reflections contain information about the layering (rock porosity, density, elasticity) in the subsurface, but it is a formidable task to process the recorded acoustic data into meaningful information that could be used to decide if a certain location could contain trapped hydrocarbon accumulations. Fig. 3 shows an example of the type of images that can be produces through seismic surveys. For the non-specialist such pictures at best give a hint of layering and fault structures, but specialists have learned to interpret these data and are able to identify possible or likely hydrocarbon reservoirs.

After this stage of "exploration" one must confront the seismic interpreters with their predictions and drill exploration wells. If these wells are on-shore, then the cost can be modest, but if the prospected reservoir is off-shore in (ultra) deep water, drilling a well is very expensive (tens of millions of dollars) and it becomes an interesting strategy game to balance the risk of drilling a "dry" well against the risk of missing a "big cat". Seismic data gives the broad contours of the reservoir but with low data resolution. Near the (exploration) wells one can extract a very detailed picture of the reservoir rock and fluids, using down-hole logging tools that use quite advanced techniques (gamma-rays, NMR, electrical resistivity, etc.) to map-out the reservoir properties (very) close to the well. Obviously there remains a lot of uncertainty in the reservoir properties even after combining seismic data, well-log data and educated guesses from experienced geophysicists and geologists.

The activities described so far are not in the main domain or reservoir engineers. Reservoir engineers become involved if a reservoir has been found and its location (depths, extent, layering and fault structure) has been roughly mapped out using the seismic data and data from the exploration wells. The task of reservoir engineers is to use this information to make a "field development plan", which describes in suitable detail where future production wells must be drilled and what type of production strategy will be used. In the early days of oil and gas recovery, reservoir engineering was simple: One would simple drill a hole and at some point there would be oil gushing out. If that did not happen, one would try again a bit farther out. These wells typically were in easily accessible locations and drilling depths were very modest so the cost of drilling such wells was low. Nowadays much of the oil is offshore and drilling depths can extend to extreme depths (the current depth record is close to 8 kilometers). Even more importantly, it is increasingly more a requirement to produce as much as (technically) possible from the subsurface oil and gas. With the simple strategy sketched above, called "primary depletion", it is only the intrinsic reservoir pressure that squeezes the oil to the surface, but this pressure declines rapidly and only a small fraction of the oil can be recovered (5-15%, say). In order to reach a higher "ultimate recovery (UR)", one must re-pressurize the reservoir, for example by injecting water or gas.

Water and gas injection to re-pressurize the reservoir and push the oil towards producing wells are examples of "secondary recovery" (fig. 4). With secondary recovery the UR can be significantly higher, 30-60%. However, the ambition nowadays is to reach ultimate recoveries of 70-80% and this requires even more enhanced oil recovery (EOR) techniques. For example, one can inject chemicals that dissolve oil and "wash-out" the rock much more effectively than plain (sea) water does. Or one can use chemicals that make the oil less viscous so that it flows more easily to the producing

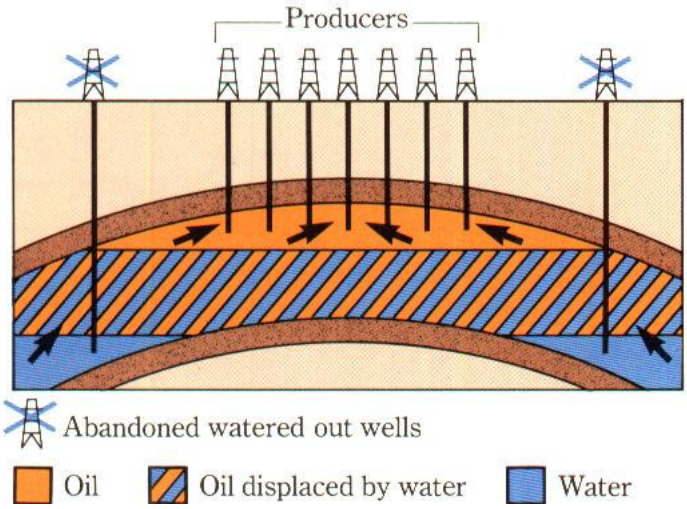

***Fig.4*: An example of secondary oil recovery: water is injected to keep the reservoir at a high pressure and push the oil towards the producing wells.**

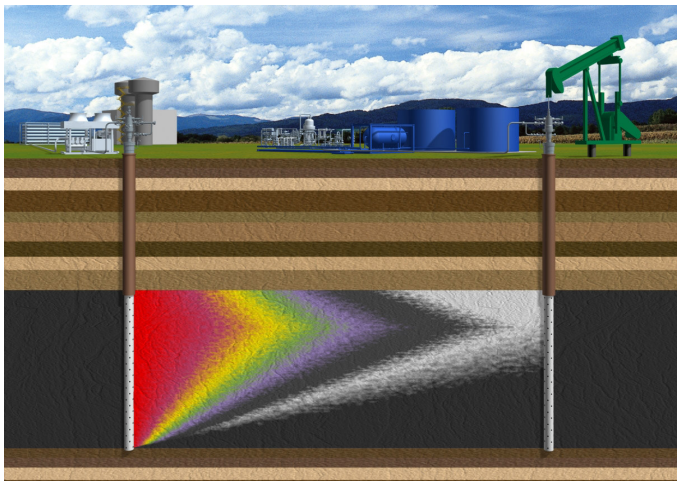

***Fig5*: An example of tertiary oil recovery, In Situ Combustion: ail is injected to "burn" the heavy oil, producing lighter, more mobile hydrocarbons and gases that help pushing the oil to the producer.**

wells. This viscosity reduction is mandatory when attempting to recover very heavy oil or bitumen. This type of hydrocarbon looks more like the stuff one uses to make a hockey puck: it essentially is a solid unless one heats it up significantly. Yet another quite advanced recovery technique uses air injection (Fig. 5). The oxygen in the air reacts with the (heavy) oil, and this "burning" produces heat and gases that help to push the oil forward. At the same time, the burning converts heavy hydrocarbons into lighter ones (and a small percentage of coal-type residue). If such a process can be controlled at the field scale, the UR could be as high as 80-90%.

With the need for these more advanced field development concepts, the role of the reservoir engineer has become more important and also the need grows for tools to help developing such plans, preferable finding the options with the largest chance of a high ultimate recovery on an economically attractive time scale and with the least environmental impact. Reservoir flow simulation is the main quantitative tool that allows exploring alternative development concepts and can give forecasts with uncertainty ranges for the various options. Also the characteristics of new or complex EOR methods can be investigated, for example the effect of injecting steam, or polymers that can dissolve oil, or other chemicals or even bacteria. By combining lab-scale experiments with field-scale reservoir simulations, the margins of uncertainty around applying such novel and usually expensive enhanced recovery methods can be reduced.

.

3. Numerical simulation of flow through porous media

To simulate the oil and gas flow on a computer, one needs a model of the subsurface rock, fluids and their interaction. Roughly speaking, this amounts to finding models that tell how the fluid moves when subject to the three relevant forces: External pressure differentials (or viscous forces), gravitation and capillary forces. In addition one must be able to compute fluid properties like viscosity, density, the phase-split (how much of the hydrocarbons are in the liquid and gas phase) as a function of pressure, temperature and fluid composition. Also the rock reacts to pressure changes and this geomechanical effect on pore volumes and possibly permeability must be modeled to some approximation. In addition to the subsurface modeling, there are the wells and usually some kind of surface facility network (pipes, gathering stations, separators, compressors etc.).

To give a flavor of the equations, one can look at the basic equation for fluid flow in porous media, the Darcy equation,

$$\mathbf{v} = -\frac{1}{\mu} K \nabla P \qquad (1)$$

This equation expresses the empirical law, that the fluid velocity **v** is proportional to the pressure gradient ∇P. The proportionally constant is the ratio of *K*, the rock permeability and μ, the fluid viscosity. If this relation is combined with the usual mass conservation equations, one ends up with one of the basic equations for fluid flow in porous media, applicable for single-phase fluids (e.g. only water or only oil):

$$\frac{\partial(\rho\phi)}{\partial t} - \nabla \cdot \left(\frac{\rho}{\mu} K \nabla P \right) = 0.$$

In this equation ρ is the fluid density, ϕ the rock porosity and *t* is the time. This simple equation has to be extended to describe multi-phase fluids: fluids that can be in the liquid or gas phase depending on the pressure, temperature and fluid composition. Besides the viscous forces shown in the Darcy equation (1) above, gravitational or

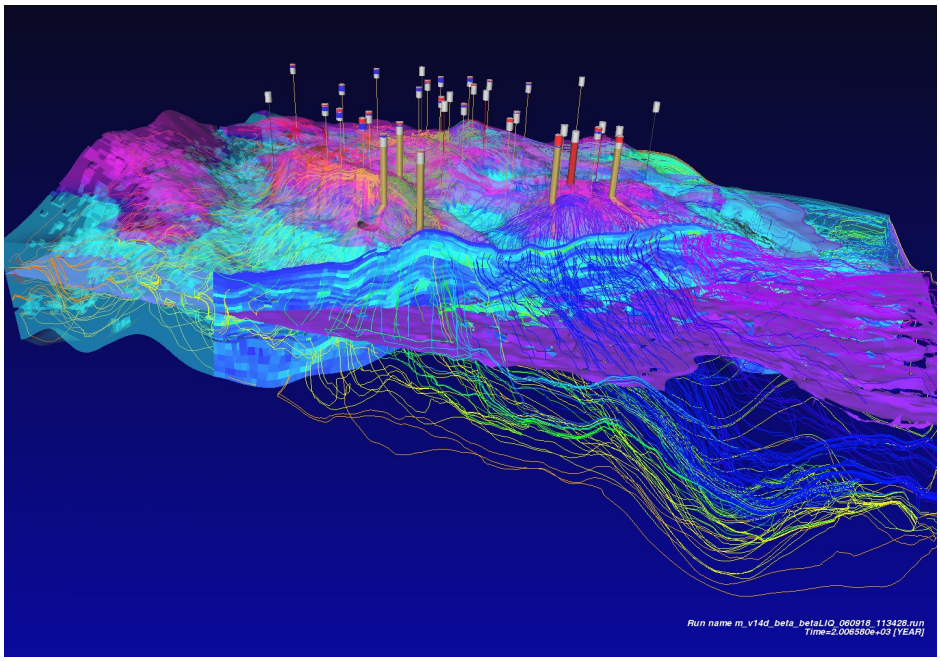

***Fig.6*: A 3D view of a MoReS reservoir model. The bumpy solid body is the actual reservoir, with oil-rich regions colored purple-red. The cut-plane shows the layer structure (permeability). The wavy lines indicate the streamlines of the fluid flow and the purple sheet is a pressure isosurface.**

buoyancy forces have to be added, which become relevant with multiple phases, because water, oil and gas have different densities. Interfacial tension between different fluids leads to capillary forces that must be accounted for as well. And finally, the simple proportionality constant K (the rock permeability) in the Darcy equation must be modified, using an empirically determined "relative" permeability factor that depends on the saturations of the fluids. The saturation of a fluid is the volume fraction of the pore space that is occupied by that fluid.

In order to solve the resulting equations numerically, they have to be discretized on a grid and in time. Fig. 6 shows a 3D view of such a "numerical" reservoir model. The grid must be chosen such that it captures the structure (layering, faults) and rock properties of the hydrocarbon reservoir. The equations for the fluid movement and properties are discretized using a finite volume method, which rigorously preserves mass conservation for the discrete equations. In this fashion one ends up with a set of coupled nonlinear equations for the amount of mass per grid block of each of the fluid (chemical) components one cares to model. For an accurate representation of the fluid properties, one would need to take a large number of (hydrocarbon) fluid constituents into account (CH4, C2H6, etc.) and compute the fluid properties (density, viscosity, phase-split) using some kind of "Equation Of State" description. But this obviously increase the computational cost and approximations are available the represent the fluid properties with only a few "effective" hydrocarbon components. For further details of the discretized, multi-phase fluid flow equations one is invited to consult references like [1-4, 5].

To finish this brief introduction to numerical flow simulation, it may be worthwhile to sketch the size of the current computational effort. Typical reservoirs, if they contain economically significant amounts of oil or gas, are large but thin; they are a bit like pancakes. The thickness of a reservoir usually is in the range of 10-100 meters, but laterally they can extend over tens of kilometers. Often the reservoir rock consists of layers, which have the same characteristics: thin but laterally extended. In order to capture this vertical heterogeneity, one needs a reasonably large number of grid blocks vertically; say between 10 and 100. Hence the thickness of grid blocks in a reservoir model is typically 0.1-10 meters. If the same grid block size would be used horizontally, in a 10 by 10 kilometer reservoir, this would easily lead to models with hundreds of millions of grid blocks. Currently such model sizes are beyond what can be simulated within reasonable times. Hence, the grid resolution horizontally typically has to betaken a factor 10-100 less than vertically. The simulation period required for a field development study, is typically some 30-60 years. Because fluids move slowly, the time step size can be as large as a week or even month, but still one needs to solve some 1000 times a very large set of highly non-linear equations. Since the reservoir engineer must always perform a large number of simulations, with different well configurations, the model sizes in practice are in the range 10,000 – 500,000 grid blocks, with simulation run times (well) below 10 hours.

The code for a full-fledged reservoir simulation tool, like the proprietary simulator MoReS used in Shell, is quite substantial. A general purpose, commercial grade simulator has to offer a lot of extra functionality besides solving the flow equations:

- An extensive intuitive graphical user interface (GUI). This must enable easy model building and results analysis.
- Connectivity to data bases and external packages to compute fluid or rock properties.
- Grid property manipulations such as coarsening and refining the grid and its properties.
- Well flow and surface facility flow modeling.

Such functionality, and more, is offered by the integrated dynamic modeling platform "Dynamo", which was developed in Shell during the last 15 years or so. Fig. 7 shows some screen shots of the Dynamo GUI. Dynamo has an object oriented data organization and separated functional modules to allow extensions of the modeling capabilities as required by the evolving business needs. It combines upscaling, reservoir flow simulation and surface facility modeling in a single integrated tool. Most of the code is written in C++, but certain compute intensive parts, for example in the linear equation solver, are written in FORTRAN. One of Dynamo's unique capabilities is its well-integrated scriptable interfaces to the data. This allows (advanced) users to tailor the simulator to solve quite complex, non-standard problems. Dynamo runs on Linux and Windows PCs and on IBM workstations. It can run parallel simulations on distributed memory platforms using MPI.

Developing such a simulation platform is a big effort and involves many people working on code development, code quality and speed improvements, release management, testing, deployment and support.

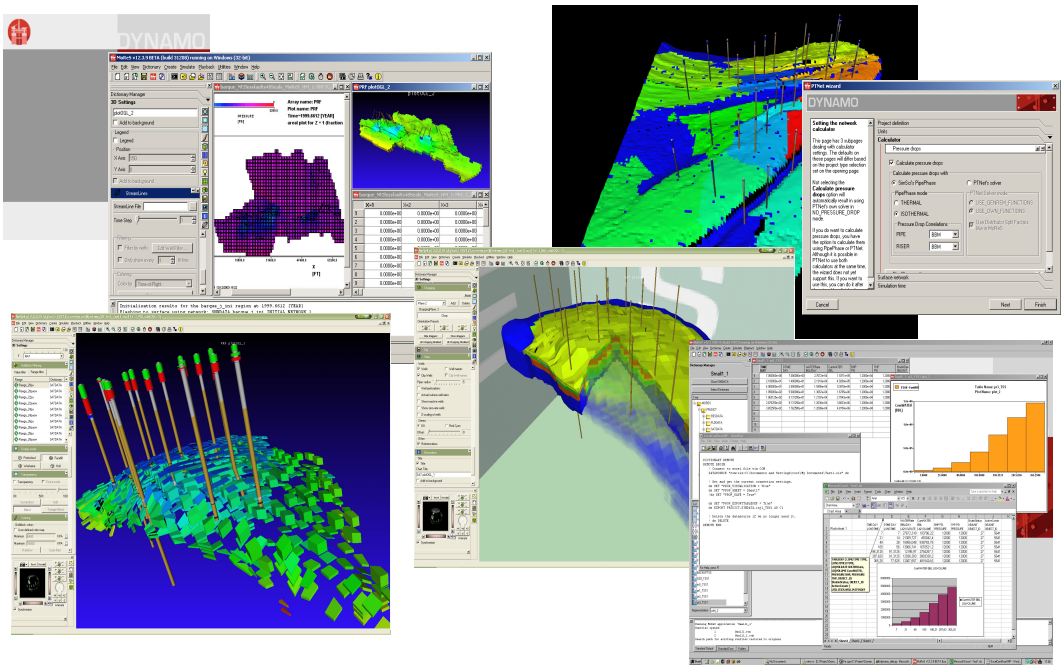

***Fig7*: Screen shots from Dynamo, the integrated reservoir simulation platform used in Shell.**

4. Computing challenges and new directions

Currently, simulation tools like Dynamo/MoReS or commercial simulators like Eclipse, are well suited for field development planning of easy to medium difficult reservoirs. However, as mentioned in the introduction, the easily accessible oil and gas is dwindling and there is a growing need for clever development options applicable to (very) difficult cases. For example, there is still a huge hydrocarbon reserve in Canada, however this is in the form of very heavy oils and bitumen that cannot be recovered using conventional reservoir engineering methods. Also the simulation tools must be extended to make them usable for these and other EOR cases. In the industry one can recognize a number of directions along which simulation tools are being developed or extended [6].

1. **Integration**. This involves connecting data and functionality along the full modeling workflow. The (static) reservoir properties are addressed with seismic tools and geostatistical packages; reservoir dynamics is modeled in fluid flow, well and facility simulators; the results are processed in economical and risk-assessment packages. It is a formidable task to connect all data and tools in a smoothly integrated workflow.
2. **Ease-of-use**. Reservoir simulation involves a lot of (input) data and produces large amounts of output. This data must be checked and validated and the reservoir engineer must keep a good overview of what the simulator is actually computing. This requires an extensive, intuitive graphical user interface.

3. **Optimization under uncertainty**. In essence the aim of reservoir simulation is to find an optimal field development plan. However, much of the input data is only roughly known and there is (hopefully quantifiable) uncertainty in the model. Hence the task of a simulation tool suite is to help the reservoir engineer in optimizing models in the presence of uncertainty.
4. **Speed and efficiency**. Improving the model quality by increasing grid resolution increases CPU times and memory requirements. Also the need to capture model uncertainty be simulating a large number of "equally probable" model realizations increases the computational burden. Hence, there is an urgent drive to find faster numerical methods, to apply more efficient parallelization methods etc.
5. **Advanced physics and processes**. Many of the EOR techniques requires a much more detailed understanding of the processes that take place in the reservoir. Such processes involve more than fluid displacement and thermodynamics: (exothermal) chemical reactions, diffusion and adsorption processes, rock geomechanics etc.

In the remainder of this section I will discuss each of these five directions in more detail and give examples of ongoing research. The examples are based on work done in Shell Dynamic Modeling R&D in Rijswijk, but also elsewhere around the world, in other oil companies and universities similar developments are ongoing.

4.1 Integration

Truly seamless tools and data integration along the full subsurface modeling workflow has been the Holy Grail for the last decade or so. Unfortunately this workflow extends across many disciplines: Seismologists and geophysicists model the static reservoir features (layer and fault structure, rock property distributions); reservoir engineers model the fluid properties, rock-fluid interactions (relative permeabilities, capillary pressure); production technologists are responsible for the well bore flow and facility engineers take care of the surface networks. Finally economists and managers need to process the data and turn the numbers into (economically and environmentally) sound strategies and take development decisions. This hybrid collection of workflows, each with their own tools, is too large to capture in a single, well-integrated software platform. Hence one must find ways to link the tools and transfer data from one tool to the next, without loosing information or accuracy. Often dedicated "middleware" software, like CORBA, plays an important role here, but also efficient usage of databases is essential.

Much of this effort heavily depends on fruitfully applying new IT techniques and since Shell is not an IT company, such integration activities are not the main focus of the Shell R&D group in Rijswijk. However, we work extensively with IT specialists on tool and data integration, and Dynamo offers an integration platform with a very smooth workflow that extends from importing static data from static reservoir modeling packages such as Petrel, using industry standard data exchange methods (RESCUE-B), through upscaling and preprocessing the data into a manageable flow simulation model. Fluid data, well production data and other input data can all be imported reasonably easily through generic or dedicated links to databases. Also the link with surface

facilities is easy within Dynamo, either using the Shell proprietary network simulator PTNet, or connecting to third party packages like GAP or HYSYS.

However, the ultimate prize of a truly integrated workflow, "from seismic inversion to point of sale" is still a desirable treasure that has not been found yet.

4.2 Ease-of-use

In some sense the power and effectiveness of a simulation tool is not only limited by the algorithms and methods that have been implemented in the software, it is also limited by the extent to which the reservoir engineers are able to use the tool. Obviously the Holy Grail here is to build a tool with very advanced capabilities that can be used effectively by a broad group of reservoir engineers. The tool should enable novice users to soundly perform basic reservoir simulation studies and enable power-users to address highly complex EOR type reservoir engineering problems and challenges.

The approach taken by Shell to provide these seemingly conflicting requirements is to offer two types of interface: A modern, point-and-click type graphical interface that facilitates a wide range of common and simple tasks. Building and analyzing simple to medium complex reservoir models should be possible entirely within the graphical user interface. Fig. 6 gives an impression of such a GUI. This type of user interface must offer a clear view on the data and available functionality – but only to a certain level of complexity. The second requirement, empowering simulation experts, we satisfy by offering an extensive scripting interface to most (if not all) simulation data. This enables power users to adjust their simulation models according to new insights and address issues that go beyond what was implemented in the standard dialogs in the GUI. Like the integration effort, building and maintaining a top-notch GUI requires special IT skills that are provided by specialist staff form IT companies that work with the simulator R&D development group in Rijswijk.

4.3 Optimization under uncertainty

Current practice in simulation studies for field developments is to run a (limited) number of different field development options. Typically the choice of scenarios is based on the experience of the reservoir engineer and the number of different variations is often limited by the available time for completing the study. Detailed or exhaustive optimizations of well positions or injection strategies are not usually part of the work. Also the effect of uncertainty in the reservoir structural and geophysical properties is investigated in often heuristic ways. A number of systematic methods have been tried on a limited scale and are being further developed to make available in standard reservoir simulation tools:

- Experimental Design is a systematic way to assess the effect of uncertain model parameters. When there are many uncertain parameters (say ten), it quickly becomes impossible to perform model simulation for all feasible parameter combinations. Even if one only considers three values for each parameter (low, medium and high) a model with ten parameters would already require 59049

simulations. Using Experimental Design, a much smaller number of simulations is used to fit the parameters of a linear or quadratic response model for the desired simulation results. This simplified response model, or proxy, is then used to evaluate the probability distribution of the (approximate) simulation results over the full range of parameters.

- Markov Chain Monte Carlo can be used to compute the probability distribution of simulation results that are subject to model parameters with a (known, or assumed) prior probability distribution. These methods are in principle more rigorous than approaches using Experimental Design because no proxy models are involved. However, Monte Carlos is currently hardly used at all, because of the large computational burden of having to run tens of thousands of reservoir simulations. However, with emerging Grid Computing methodologies and growing numbers of fast and powerful PCs rigorous Monte Carlo methods are likely to be used more and more in the near future.
- Direct simulation of uncertainty distribution. Such methods are still in their infancy. Rather than sampling the probability distribution of simulation results by performing a large number of forward simulations, one could try to derive equations for the (low order) moments of the probability distribution. This approach has been used with some success in ground water flow modeling, where the dynamic equations are much more linear than in our field.

An accurate assessment of uncertainty ranges of model predictions is obviously very important and Experimental Design and Monte Carlo methods address this requirement. However, the underlying objective of a simulation study is to find an optimal field development plan. Optimality can mean different things: highest ultimate recovery, most oil in the next 10 years, highest return on invested money, the least amount of side effects like produced (polluted) water etc. A different type of optimization aims at reducing the uncertainty in model parameters. Here the objective is to adjust uncertain parameters such that the model predictions match the actual field data, such as well production rates and pressures.

These optimizations currently are mostly done using a combination of educated guesses and trial and error. However, with the ever-growing compute power, it is also becoming feasible to use automated optimization methods [7]. Such optimization methods usually need gradient information if the number of parameters that must be optimized is large. If the parameter space is low dimensional, gradient free methods (such as genetic algorithms or simulated annealing) are also under investigation. The number of controllable parameters can be quite large: in a field with 20 injectors each of which can inject water of gas or both, with a rate and ratio that can be changed every month or week, there is a huge number of injection patterns. If these injection wells are so-called “smart wells”, then they are equipped with subsurface valves and sleeves that allow the production technologists to selectively inject fluid in certain layers. This further increases the controllability of the field and makes it even more difficult to find an optimal development plan. Also when trying to reduce uncertainty in model parameters (such as fault seal factors, details of the rock permeability and porosity, shape of the relative permeability and capillary pressure curves etc.) the number of adjustable parameters is very large.

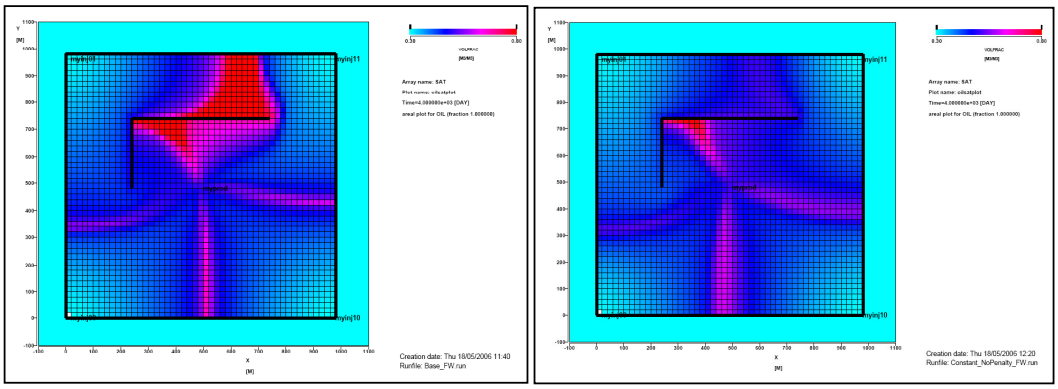

***Fig8*: Automated production optimization with the adjoint method. The two pictures show the remaining oil (red) after 10 years of production. The oil is produced by a well in the center, four water injectors (at the four corners) push the oil to the producer, but the L-shaped barrier (fault) prevents a simple sweep. With adjoint optimization (right) the time-dependent injection rates are adjusted such that the ultimate recovery is significantly higher.**

It is an active and ongoing R&D effort to compute gradients of such (control and uncertain) model parameters using the "adjoint method". This is a very efficient technique to compute gradients when the number of parameters is very large, but the number of objective functions (targets) that must be optimized is low. In combination with steepest descent type optimization algorithms, or other gradient based (non-linear) optimization methods, this can offer automated model optimization. Fig. 8 shows an example in which the adjoint method in MoReS is used to optimize rates of four water injectors [8].

The next step is to combine uncertainty quantification with optimization and perform optimization of field development plans, taking the uncertainty distributions of unknown reservoir parameters properly into account. How this is best done in a practical way is still an open problem.

4.4 Speed and accuracy

Increasing computational speed is obviously a requirement now and in the future. Progress in hardware (higher clock frequencies, multi-core processors etc.) helps, but the need for more compute cycles always appears to grow faster than the supply. In reservoir simulation, as in other high-performance computing areas, three approaches are followed to improve the speed of a simulation run

4.4.1 Low-level optimization

The existing code can always be further optimized. In Shell this is done in close collaboration with hardware vendors. Often this requires platform dependent code modifications and one has to balance the benefit of speed-up against the trouble of maintaining such ad-hoc code variations.

4.4.2 parallelization

Dynamo/MoReS, like most commercial flow simulators, has been adapted to allow simulation on multiple (distributed) processors. For some parts of the code this implies that special algorithms must be used. For example, the incomplete block LU preconditioning in the linear solver has to be replaced by a domain-based preconditioning method. The efficiency of such parallelizable algorithms is often less than the sequential version and this jeopardizes efficiency when using a large number of processors. Hence, finding efficient (linear solver) methods for parallel processing is quite important. Currently, flow simulations cannot yet be done efficiently on very many processors. A speed-up of a factor 2-2.5 on 4 processors and a factor 3.5-4.5 on 8 is state of the art. In order to allow efficient simulations on larger numbers of processors, say 16-128, more efficient and robust linear solver and load balancing methods must be developed. Combining algebraic multi-grid with incomplete LU preconditioning methods is a new development that shows promising results.

4.4.3 Alternative solution methods

The goal here is to find new or modified methods that achieve the desired simulation results with less computational effort. Traditionally a fully implicit (backward Euler) time integration of the finite-volume discretized flow equations is used. This is a rather brute-force approach to solving the problem and there may be other methods that either use fewer degrees of freedom, or break-up the problem in more simple sub-problems.

Researchers from other (simulation) fields often are surprised that we do not use finite-element methods or higher order (Runge-Kutta) integration schemes. These methods have been tried and were not found beneficial. Oil and gas flow in reservoirs has special features: the rock is highly heterogeneous, not only vertically where layers can alternate between very permeable and almost impermeable, but also laterally, where a background sedimentary deposition can be intersected by river depositions or other localized features. At the scale of the simulation grid blocks, this implies an almost complete lack of continuity. In time there are drastic changes in boundary conditions (well rates) at every time step. Both higher-order finite element methods an time integration schemes require some measure of continuity to become efficient. However, this does not mean there is no place for higher-order schemes: reducing numerical dispersion, and achieving sharper fluid fronts is very important in certain types of studies and higher order schemes could be used successfully there.

Two (other) examples of finding efficient alternative solutions methods are multi-scale techniques and streamline methods. Both methods try to exploit the fact that reservoir flow requires computing pressures and saturations, which have quite different characteristics: pressure is smoothly changing, even if the underlying rock properties are heterogeneous. Saturations, i.e. fluid distributions, are very erratic specifically near fronts where two or more fluids meet (water pushing oil and/or gas). To some extend the pressure gradients drive the fluid displacements and if the back-reaction of the fluids on the pressure distribution can be ignored, it is possible to solve pressure first and compute the fluid transport in a separate step.

This is done in streamline methods [9], where a pressure field is computed first, after which the fluid transport is computed in a clever way: By using a local coordinate

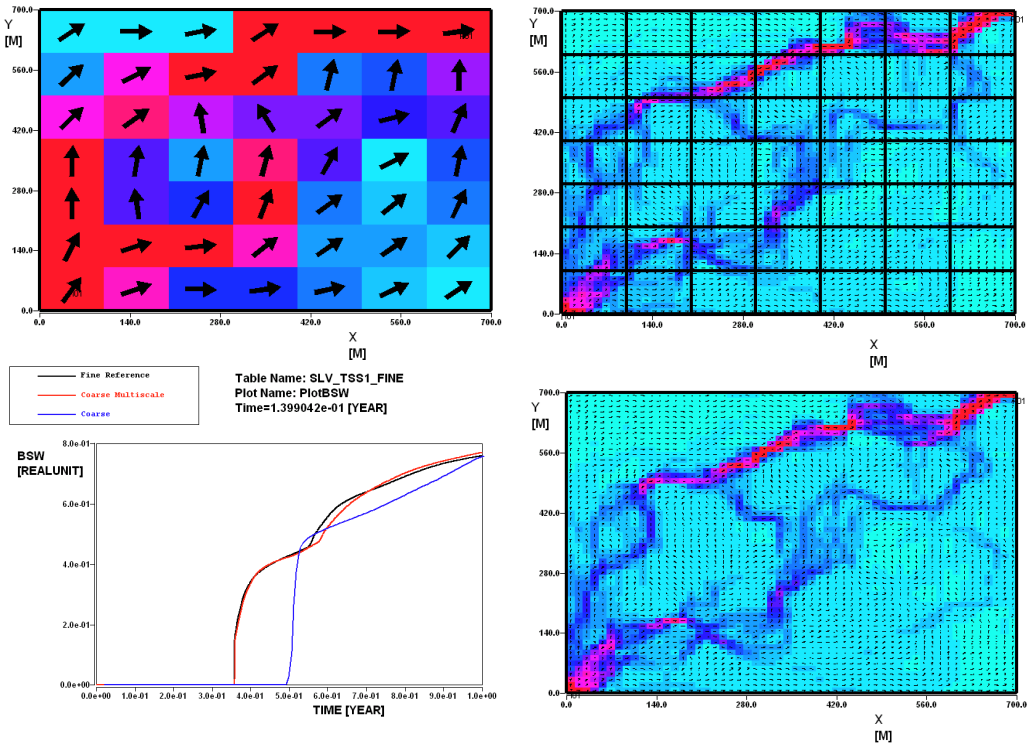

***Fig9*: Multi-scale flow simulations in MoReS. Top left shows the velocity field on the coarse grid. Top right shows the interpolated velocity field, using local mode functions. This is much easier to compute that the fine-grid velocity field shown bottom right. However, the results are quite similar. Also the water production computed with the multi-scale method (red curve in the bottom left graph) is very similar to the production computed in the fine-grid reference run (black curve). The blue cure shows the poor results computed on the coarse grid [9].**

system that follows the constant-pressure lines and the orthogonal flow lines, the transport equations decouple into a set of one-dimensional differential equations. These equations can be solved very efficiently and if the underlying (difficult, elliptic) pressure equation need only be solved occasionally, this approach is very fast. For a class of problems, which are dominated by viscous forces, such as water flooding field developments, streamlines offer a very promising alternative to the robust but slow full-physics simulations.

Also in multi-scale methods one tries to exploit the different nature of the pressure and saturation fields and the fact that the coupling between the two can be weak, or only relevant locally. With multi-scale methods one solves the pressure on a coarse grid, after which the solution is “propagated” to a much finer grid using “mode functions”. These mode functions can be computed locally and fill-in the details of the pressure field that are required to subsequently compute the saturation changes on this finer grid (see the examples in Fig. 9). Multi-scale methods are different from the more traditional (dynamic) grid refinement approaches, because pressure and saturation are solved in separate steps, whereas in grid refinement techniques, the same (coupled) equations are solved using grids with different resolutions.

Obviously there are many variations possible by combining such multi-scale ideas, with streamline techniques and (dynamic) grid refinement methods. This is a very active area of research and hopefully will lead to faster simulations in the (near) future.

4.5 Advanced processes and physics

The last of the five directions for innovation in reservoir simulation R&D focuses on the physics and displacement processes. More than in the other topics, this is an area in which a lot of research and experiments is also done in other groups, outside flow simulation R&D. In fact, a lot of the driving force behind implementing new features in a proprietary simulator like MoReS comes from the need to test or validate research ideas that are developed in other groups in Shell; It is invaluable to be able to compare experiments or pilot studies with the prediction of the theory as it is implemented in the simulation tool. Currently there is a lot of activity around EOR and simulation is an essential tool to help develop and optimize new EOR methods. In the remainder of this section I will briefly discuss a number areas where flow simulation is going to be used more intensively.

4.5.1 Thermal simulations

There are still huge hydrocarbon reserves left if also the (ultra) heavy oils and oil sands are included in the accounting. Heavy oil can be produced by heating the fluid, usually be injecting steam (though electrical subsurface heaters could be used too). The intricate interplay between heat, rock and fluids require a flow simulator that can compute the combined effect of heat conduction, transport by flowing fluids, changing fluid properties, along with the normal viscous, capillary and gravitational effects [2].

4.5.2 Coupled flow and geomechanics

Usually the effect of rock elasticity is only accounted for by allowing the pore space to change with changing pressure. In reality, however, the effect of changing pressures cannot (always) be accounted for with such a local response and stresses and strains in the rock must be computed by solving more rigorous stress-strain equations. Rock stress extends outside the actual (oil baring) reservoir rock and usually the computational domain for such a geomechanical stress computation is much larger. Hence, performing coupled simulations of fluid flow and geomechanics requires combining effects of different scales with different gridding requirements and this makes it a difficult problem. Ignoring the effects of coupling flow and geomechanics was justifiable in many of the “easy” reservoirs. However, in many of the heavy oil reservoirs, where thermal effects add to the changing stresses in the rock and in deep-sea environments with unconsolidated sands, geomechanical back reaction effects can be very substantial. Also in fractured rocks, where the fractures can be naturally present of may have been induced by forcefully injecting water, the properties of the fracture network (opening and closing of fractures) should be computed using coupled flow and geomechanics.

4.5.3 Fractured reservoirs

Often one thinks of (subsurface) rock as solid, and it may be a surprise that there is in fact so much empty (i.e. non-rock, fluid-filled) pore space in between the solid grains. A

second important feature of certain types of rock, mainly carbonates, is that it can be fractured. Such rock contains a network of cracks and fissures, with apertures ranging from tenths of millimeters to meters. Such large, meter-scaled fractures are not fully open space of course, but are like crushed rock zones. They occur usually near faults, where the crushed rock material is much more permeable than the embedding material. From a fluid-flow point of view, the characteristic property of such fractured materials is that there are in fact two interconnected types of porous material: The original (matrix) rock, which accounts for the bulk of the volume but has low porosity and has usually a very low permeability; and the fracture network that has a low volume but is extremely permeable. Because of the permeability contrast, the processes in matrix and fractures have a quite different time scale and simulating fluid flow in such a "dual-permeability" system is difficult. This complexity increases further, if thermal and geomechanical effects have to be added (which is required when the fractures rock contains heavy oil).

4.5.4 Coupled flow and (chemical or microbial) reactions

With increasingly high oil prizes it becomes economically viable to use relatively expensive methods to improve the ultimate recovery of a field. Examples of such methods are polymer and surfactant injection. Polymers make the water more viscous and therefore more effective in sweeping the oil towards the producer. Surfactants dissolve more or the residual oil that would otherwise stay stuck in the smaller pores in the reservoir rock. Such surfactants can be chemically active and also polymers absorb to the rock. Hence it is important to be able to simulate the chemical reactions of these materials along with the fluid flow. Besides injecting chemicals, there are also attempts to use bacteria to convert heavy oil into lighter fluids. Another example where simulating chemical reactions is essential is in-situ combustion, where air is injected in oil. The oxygen in the air reacts with the oil, producing heat and lighter hydrocarbons (and coal that remains in the reservoir). Accurate simulation of this process is a wonderful computational challenge, since it combines fluid flow, chemical reactions, thermal effects; all happening in a very narrow moving front at the sub-meter scale.

4.5.5 CO2 sequestration

Finally it is worth to mention an area of new physics and application of flow simulations that is growing rapidly. Injecting CO2 in depleted gas reservoir or water-baring layers deep in the subsurface may be a practical way to curb or at least slow-down the growing concentration of this greenhouse gas in the atmosphere. Injecting CO2 in depleted gas reservoir is relatively easy: these formation have proven to be able to contain gases for (very) long times. However, there are not too many such reservoirs and to make large scale disposal of CO2 near the source (energy plants) possible, also injection in water baring layers, saline aquifers, is being investigated. Here one must assess the various processes that keep CO2 contained in the water: dissolution, trapping in pore space through capillary forces, binding into solids through chemical reactions etc. Reservoir flow simulators can be used to study these processes and help finding reliable injection locations and disposal strategies.

5. Conclusion

Numerical simulations of oil and gas fields have been valuable in the past to help planning for efficient development of these hydrocarbon resources. It is not hard to predict that the role of simulation and (reservoir engineering) technology in the future will only become more important: the easy oil and gas fields are mostly under development, new fields almost invariable pose technological challenges. For instance because they are in ultra deep water, in arctic regions, or the hydrocarbons contain acid or other pollutions, or the oil is so heavy it does not flow. The huge hydrocarbon reserves in the form of oil sand and very have oil pose an interesting dilemma: On the one hand they can offer a source of much needed energy for many decades to come, on the other hand it takes a lot of energy to get it out of the ground. Furthermore, the CO2 content in these hydrocarbons is very high, and – like burning coal – using such fuels may be unacceptable in a world where CO2 emission has to be reduced. Technology can help on both counts: new methods may allow producing heavy oil such that much of the CO2 remains subsurface, as is the case with in-situ combustion; The same technology that is used to inject CO2 in reservoirs to enhance oil recover, can be used to dispose CO2 in depleted gas reservoir or saline aquifers.

Shell has made the strategic choice to tackle the challenges of developing difficult hydrocarbons in an efficient and responsible way, through technology innovations where needed. Hence, reservoir-engineering research, also in the area of flow simulation is reinforced. Technology will not provide the means to indefinitely delay the time when oil and gas cease to be the easy energy providers they are now. But technology, and reservoir simulation as part of it, will make it possible to better, more responsibly use the remaining hydrocarbon reserves. In this way a gradual conversion to different, sustainable energy sources can be made.

References

[1] K. Aziz, A. Settari, *Petroleum Reservoir Simulation*, BlitzPrint LTd, Calgary 1979

[2] R. M. Butler, *Thermal Recovery of Oil and Bitumen*, GravDrain's Blackbook, Calgary 1997

[3] J.R. Fanchi, *Principles of Applied Reservoir Simulation*, Elsevier, 2005

[4] L.P. Dake, *Fundamentals of Reservoir Engineering*, Elsevier, 1978

[5] J. Bear, *Dynamics of Fluids in Porous Media*, Dover, New York, 1988

[6] See for example the proceedings of the bi-yearly Reservoir Simulation Symposium, [http://www.spe.org/spe-app/spe/meetings/RSS/2007/index.htm]

[7] J.D. Jansen, Model-based Control of Subsuface Flow (and references therein), 8th international symposium on Dynamics and Control of Process Systems (DYCOPS) Cancun, 2007.

[8] H. Kraaijevanger et al, *Optimal Waterflood Design Using the Adjoint Method,* SPE 105764 (2007)

[9] H. Dijk, *Multiscale simulations in MoReS* [unpublished]