

Mathias C. Bellout

Joint Optimization of Well Placement and Controls for Petroleum Field Development

Thesis for the degree of Philosophiae Doctor

Trondheim, December 2014

Norwegian University of Science and Technology
Faculty of Engineering Science and Technology
Department of Petroleum Engineering and Applied
Geophysics



NTNU – Trondheim
Norwegian University of
Science and Technology

NTNU

Norwegian University of Science and Technology

Thesis for the degree of Philosophiae Doctor

Faculty of Engineering Science and Technology
Department of Petroleum Engineering and Applied Geophysics

© Mathias C. Bellout

ISBN 978-82-326-0652-8 (printed ver.)
ISBN 978-82-326-0653-5 (electronic ver.)
ISSN 1503-8181

Doctoral theses at NTNU, 2014:369

Printed by NTNU-trykk

*Para mi Querida Mamá y Familia.
Siri, Hannah Azul y Sofia Celeste, mis Reinas.
Abuelito, eres mi Corazon, y mi Roble.*

*Hearkening to these voices, East and West, by early sun-rise, and by fall of eve, the
blacksmith's soul responded, Aye, I come! And so Perth went a-whaling.*

Moby Dick, Herman Melville, 1851

Abstract

In this thesis we optimize the drilling location and operational controls of wells in a joint manner to improve the overall development strategy for a petroleum field. In particular, in this thesis we treat the integrated problem of searching for an improved well placement configuration while also taking into account the control settings of the production and/or injector wells planned for the development of the hydrocarbon asset. In oil field development, the well placement and well control problems are commonly performed in a sequential manner. However, this type of sequential approach cannot be expected to yield optimal solutions because it relies on handling well production controls using heuristic techniques during the well placement part of the procedure. In this work, we develop a nested (joint) optimization approach that seeks to capture the interdependency between the well configuration and the associated controls during the optimization search.

This thesis summarizes the development of the joint approach; from establishing the methodology while using relatively simple cases and performing thorough comparisons against sequential approaches, to further extending and finally testing the methodology using a real field case model. This progression naturally divides the work in this thesis into two parts with different research focus. The first part of this work (Chapter 2) focuses chiefly on creating proper definitions and on establishing the proposed methodology against common approaches. The second part of this thesis (Chapters 3 and 4), on the other hand, focuses mainly on applying the developed methodology within a real field case scenario involving the North Sea Martin Linge oil reservoir. The dual aim of this application work is both to further develop the methodology, and to produce and test optimization solutions that may serve as decision-support to engineering efforts within the development work process of the Martin Linge field.

Chapter 2 establishes the core of the methodology followed in this thesis. This chapter introduces the joint and sequential approaches as different ways to solve for the coupled well placement and control problem. The joint approach embeds the well control optimization within the search for optimum well placement configurations. Derivative-free methods based on pattern search are used to solve for the well-positioning part of the problem, while the well control optimization is solved by sequential quadratic programming using gradients efficiently computed through adjoints. Compared to reasonable sequential approaches, the joint optimization yields a significant increase in net present value of up to 20%. Compared to the sequential procedures, though, the joint approach requires about an order of magnitude increase in the total number of reservoir simulations performed during optimization. This increase, however, is somewhat mitigated by

the parallel implementation of some of the pattern search algorithms used in this work.

Chapter 3 focuses on extending and applying the methodology developed in the previous chapter within a real field development scenario. A work process loop is set up to guide the entire application effort; from work model validation and problem definition, to optimization effort and solution testing. Results from the optimization effort, using an approximated work model, yield a mean increase in FOPT of close to 33% for solutions developed using the joint approach. In comparison, solutions developed using a sequential approach yield a mean increase in FOPT of almost 26%. Moreover, cost function evolution data for the joint runs using the field case work model, show that the performance of less promising locations during the well placement search may be improved significantly due to the embedded control routine. This supports the notion that the nested routine may contribute to a smoothing of the outer loop optimization surface with respect to the well placement variables, and that this smoothing may add some robustness to the well placement search conducted by the joint approach. Furthermore, as seen previously in Chapter 2, the cost of the joint approach is still substantially higher compared to the computation required by the sequential alternatives. For the application work in Chapter 3, the mean total number of reservoir simulations required by the joint runs is almost 7 times higher than the mean total number of reservoir simulations needed by the sequential runs.

In Chapter 4, the well placement solutions obtained using the work model are transferred and tested on to the original field case model for the Martin Linge oil reservoir. A main result from the overall testing shows that those well placement configurations with B wells that aggressively target the eastern lobe of the Martin Linge oil reservoir yield the greatest increases in field oil production total. Furthermore, the various well placement solutions are tested for two realistic field development considerations that were not included in the optimization effort, i.e., the original larger production time frame and a multiple realizations field case scenario. Compared to the initial well configuration for the single-realization case, the best-performing well placement solution yields an increase in field oil production total of, respectively, 25.5% for the time frame used in the optimization procedure, and of 13% for the original (larger) production horizon. However, we note that we performed the optimization effort using only a single realization. Consequently, the solutions developed from this procedure are observed to lose most of their gains once these well configurations are implemented for both the larger production time frame and over the multiple realization case. This final test result underscores the importance for future work of both improving the computational performance of the overall optimization procedure (e.g., through surrogate techniques) and of including geological uncertainty within the search routine.

Preface

This thesis is submitted to the Norwegian University of Science and Technology (NTNU) for partial fulfillment of the requirements for the degree of Philosophiae Doctor. We are grateful to the Center for Integrated Operations in the Petroleum Industry at NTNU (IO Center) and to Total E&P Norge AS for partial funding of this work.

This doctoral work would not have been possible without the key trust and support from my supervisors. First and foremost I would like to thank my main supervisor Professor Jon Kleppe at the Department of Petroleum Engineering and Applied Geophysics (IPT), and my co-supervisor Professor Bjarne Foss at the Department of Engineering Cybernetics (ITK), both at NTNU, for their constant support and thoughtful mentoring throughout the twists and turns of this work.

Furthermore, I am indebted to all the colleagues and IO Center collaborators I have had the privilege to work with. In particular, I would like to thank Professor Louis J. Durlofsky at the Department of Energy Resources Engineering, Stanford University; Dr.Eng. David Echeverría Ciaurri at IBM T.J. Watson Research Center; and Dr. Stein Krogstad at SINTEF Applied Mathematics for their considerable intellectual input, and for always demonstrating the utmost patience and integrity towards me, and the work, throughout our years of working together. I consider it an immense fortune to have worked with, and learned from, such dedicated professionals.

Moreover, I also would like to thank the Stanford Smart Fields Consortium and the Department of Energy Resources Engineering for hosting me as a visiting researcher in 2010/2011, and to Dr. Andrew Conn at IBM T.J. Watson Research Center for hosting me during the summer of 2012. Not the least, I would like to express my gratitude to the Geosciences team at Total E&P Norge AS for their considerable input and assistance during the field case application part of this work. Thanks also to all my friends and colleagues at IPT and ITK for the valuable discussions throughout the work in this thesis.

Finally, I would like to thank my family for always lovingly encouraging me to get my “Book” done. Without your love and support, none of this work would ever have happened.

Table of Contents

Abstract	i
Preface	iii
Table of Contents	v
List of Tables	ix
List of Figures	xi
Abbreviations	xv
1 Introduction	1
1.1 Field development optimization	1
1.1.1 Overview of field operations	1
1.1.2 Motivation for optimization	3
1.2 Summary and contributions from each chapter	6
1.2.1 Main contributions from Chapter 2	6
1.2.2 Main contributions from Chapter 3	7
1.2.3 Main contributions from Chapter 4	9
1.3 Thesis outline	11
2 Joint Optimization of Oil Well Placement and Controls	13
2.1 Introduction	14
2.2 Problem statement	16
2.2.1 Governing equations for reservoir production	16
2.2.2 Optimization problem	17
2.3 Optimization methodology	18
2.3.1 Well control optimization	19
2.3.2 Well placement optimization	20
2.3.3 Sequential approaches for well placement and control optimization	21
2.3.4 Joint approach for well placement and control optimization	22
2.4 Example cases	24
2.4.1 Optimization of injector location and control of five wells	26

2.4.2	Optimal location and control of five wells	33
2.5	Concluding remarks	39
3	Joint Optimization Applied to a Real Field Case	41
3.1	Targets and strategy for application development	42
3.1.1	Application targets and strategy	42
3.1.2	Work process structure	47
3.2	Optimization framework	51
3.2.1	Optimization framework	51
3.2.2	IO Center resource platform and network	63
3.3	Field case and validation work	65
3.3.1	Martin Linge field case	65
3.3.2	Martin Linge oil reservoir	67
3.3.3	Field model transfer and validation	69
3.4	Optimization work	80
3.4.1	Introduction to optimization work	80
3.4.2	Problem formulation	82
3.4.3	Methodology	84
3.5	Optimization results	91
3.5.1	Optimization runs	91
3.6	Discussion and suggestions for further work	103
4	Testing of Solutions on Field Case Model	107
4.1	Test results from solution cases	107
4.1.1	Individual analysis: Final well configurations	113
4.1.2	Collective analysis: Total oil production values	142
4.1.3	Increases in recovery versus changes in well length	152
4.2	Field model tests on multiple realizations	158
4.2.1	Solution tests on multiple realizations: 1200 days	159
4.2.2	Solution tests on multiple realizations: 5174 days	163
4.2.3	Hybrid solution tests on multiple realizations	167
4.3	Final topics on field case application	175
4.3.1	Applicability of solutions and limits of application design	175
4.3.2	Other more advanced production scenarios	178
4.3.3	Final comparison to current base case configuration	178
4.3.4	Final comment on pilot study	178
5	Chapter Summaries and Further Work	181
5.1	Chapter summaries	181
5.2	Topics for further work	189
5.2.1	Cost-reducing topics for further use of optimization framework	190
5.2.2	Topics for further development of optimization framework	191
6	Summary of results	193
	Bibliography	195

Appendix A: Validation Work Results	209
Appendix B: Reduced-Order Control Optimization	218

List of Tables

1.1	Outline describing the content of chapters 2 to 4 in this thesis.	11
2.1	Model and optimization parameters for the two examples in this chapter. .	25
2.2	Injector well location and NPV for the best solution obtained for the ex- haustive explorations.	28
2.3	Average NPV (over 12 runs) for the optimal location of one injector and control of five wells.	28
2.4	Results for NPV and total number of simulations n_{sim} for the optimal lo- cation and control of five wells.	33
2.5	Results for NPV and total number of control optimizations solved (n_{ps}) for the second example.	38
3.1	Simulation cases used in transfer of model from Eclipse to AD-GPRS sim- ulator.	71
3.2	Description of simulation model parameters modified in transfer of model from Eclipse to AD-GPRS simulator.	72
3.3	Mnemonics corresponding to the production quantities used to compare the original Eclipse model and the different work model configurations. .	73
3.4	Produced versus in place field volumes (in place volumes obtained from ECLak_orig)	74
3.5	FOPT and well lengths corresponding to the well placement solutions for joint and sequential runs.	92
3.6	Mean increases in FOPT resulting from embedded control optimizations during joint runs.	94
3.7	Performance of solutions obtained using the AD-GPRS work model when transferred to the Eclipse field case model.	97
3.8	Transfer results given as percentage increases in field and well oil produc- tion total for the field case model, $\Delta FOPT$ and $\Delta WOPT$, respectively. . .	101
4.1	Main parameters for grid, simulation and fluid properties for field model used for all solution tests.	110
4.2	Main parameters for Eclipse keywords related to well description and pro- duction strategy.	111

4.3	Final well lengths for base case and solutions.	113
4.4	Final well depths for base case and solutions.	114
4.5	Percentage increases in field and well oil production total for a 1200 day production horizon, $\Delta\langle\text{FOPT}\rangle$ and $\Delta\langle\text{WOPT}\rangle$, respectively.	143
4.6	Percentage increases in field and well oil production total for a 5174 day production horizon, $\Delta\langle\text{FOPT}\rangle$ and $\Delta\langle\text{WOPT}\rangle$, respectively.	146
4.7	Table showing mean FOPT increase ($\Delta\langle\text{FOPT}\rangle$) and well mean WOPT increases ($\Delta\langle\text{WOPT}\rangle$) for all solutions tested over the multiple realization set for the 1200 day production time frame.	159
4.8	Table showing mean FOPT increase ($\Delta\langle\text{FOPT}\rangle$) and well mean WOPT increases ($\Delta\langle\text{WOPT}\rangle$) for all solutions tested over the multiple realization set for the 5174 day production time frame.	164
4.9	Well composition of the three hybrid solutions developed from applying a simple heuristic procedure on the JNT1M1 solution.	168
4.10	Table showing mean FOPT increase ($\Delta\langle\text{FOPT}\rangle$) and well mean WOPT increases ($\Delta\langle\text{WOPT}\rangle$) for all hybrid solutions tested over the multiple realization set for the 1200 day production time frame.	170
4.11	Table showing mean FOPT increase ($\Delta\langle\text{FOPT}\rangle$) and well mean WOPT increases ($\Delta\langle\text{WOPT}\rangle$) for all hybrid solutions tested over the multiple realization set for the 5174 day production time frame.	172

List of Figures

2.1	Permeability field (mD) used for the two cases in this chapter.	24
2.2	Exhaustive search results for (a) NPV (\mathbf{x}, \mathbf{u}_0) with \mathbf{u}_0 a fixed control strategy, (b) NPV (\mathbf{x}, \mathbf{u}_0) with \mathbf{u}_0 a reactive control strategy, and (c) NPV* (\mathbf{x}).	27
2.3	Cumulative production and injection profiles for the well location and controls corresponding to the highest NPV solution.	29
2.4	Oil saturation distribution at the end of production time frame for the well location and controls corresponding to the highest NPV solution.	30
2.5	Injection and production well controls (BHPs) corresponding to the highest NPV solution.	31
2.6	Evolution of the objective function (NPV) for all nine runs versus number of simulations.	35
2.7	Oil saturation distribution at the end of the production time frame for the well controls and locations corresponding to the run from Table 2.4 with maximum NPV.	36
2.8	Evolution of NPV averaged over nine runs versus the equivalent number of control optimizations for the three pattern search optimization algorithms considered.	37
3.1	Illustration of application development.	46
3.2	Work process loop.	48
3.3	Optimization framework with three example models.	54
3.4	Martin Linge field.	65
3.5	Martin Linge field development.	66
3.6	Martin Linge oil and gas reservoirs.	67
3.7	Field pressure over time.	75
3.8	Well bottom-hole pressures.	76
3.9	Field production rates: gas, oil, water, liquid.	77
3.10	Field production totals: gas, oil, water, liquid.	78
3.11	Well water cuts.	79
3.12	Initial saturations of gas, oil and water at the Martin Linge oil reservoir model.	81
3.13	Heel and toe circular constraint regions implemented for each of the four wells in this work.	86

3.14	Example of P_{wd} applied to toe of well B.	87
3.15	Example of projection of infeasible well toe coordinate onto corresponding feasible bound area.	88
3.16	Function evolution as a function of number of function calls for 3.16(a) sequential and 3.16(b) joint runs.	95
3.17	Function evolution as a function of total number of reservoir simulations for 3.17(a) sequential and 3.17(b) joint runs.	96
3.18	Oil saturation maps (HuPhiSo) at 0, 1200 and 5174 days of production. . .	102
4.1	HuPhiSo oil saturation maps at three production times for BASECASE solution.	115
4.2	HuPhiSo oil saturation maps at three production times for BASECASE (left column) and JNT2M1 solution (right column).	118
4.3	Production profiles for JNT2M1 solution and base case wells: Well oil production rates (WOPR), well gas–oil ratios (WGOR) and well water cuts (WWCT).	119
4.4	HuPhiSo oil saturation maps at three production times for BASECASE (left column) and FXD1M1 solution (right column).	121
4.5	Production profiles for FXD1M1 solution and base case wells: Well oil production rates (WOPR), well gas–oil ratios (WGOR) and well water cuts (WWCT).	122
4.6	HuPhiSo oil saturation maps at three production times for BASECASE (left column) and JNT2OPT solution (right column).	124
4.7	Production profiles for JNT2OPT solution and base case wells: Well oil production rates (WOPR), well gas–oil ratios (WGOR) and well water cuts (WWCT).	125
4.8	HuPhiSo oil saturation maps at three production times for BASECASE (left column) and FXD2OPT solution (right column).	127
4.9	Production profiles for FXD2OPT solution and base case wells: Well oil production rates (WOPR), well gas–oil ratios (WGOR) and well water cuts (WWCT).	128
4.10	HuPhiSo oil saturation maps at three production times for BASECASE (left column) and JNT2OPT2 solution (right column).	130
4.11	Production profiles for JNT2OPT2 solution and base case wells: Well oil production rates (WOPR), well gas–oil ratios (WGOR) and well water cuts (WWCT).	131
4.12	HuPhiSo oil saturation maps at three production times for BASECASE (left column) and FXD2OPT2 solution (right column).	133
4.13	Production profiles for FXD2OPT2 solution and base case wells: Well oil production rates (WOPR), well gas–oil ratios (WGOR) and well water cuts (WWCT).	134
4.14	HuPhiSo oil saturation maps at three production times for BASECASE (left column) and JNT2CUT solution (right column).	136
4.15	Production profiles for JNT2CUT solution and base case wells: Well oil production rates (WOPR), well gas–oil ratios (WGOR) and well water cuts (WWCT).	137

4.16	HuPhiSo oil saturation maps at three production times for BASECASE (left column) and FXD2CUT solution (right column).	139
4.17	Production profiles for FXD2CUT solution and base case wells: Well oil production rates (WOPR), well gas–oil ratios (WGOR) and well water cuts (WWCT).	140
4.18	FOPT for all solutions, 1200 days production time.	144
4.19	WOPT for all solutions, from 0 to 1200 days production time.	145
4.20	FOPT for all solutions, 5174 days production time.	147
4.21	WOPT for all solutions, from 1200 to 5174 days production time.	148
4.22	WBHP for all solutions, from 0 to 5174 days production time.	149
4.23	Increase in field oil recovery compared to base case values (Δ FOPT), plotted against the difference in total drainage length (wells A, B, C, and D) computed between each solution and corresponding base case wells.	153
4.24	WOPT increases versus well length increases relative to base case values and length, respectively, for wells A, B, C, and D.	154
4.25	Field oil production totals for multiple realization case using the JNT2M1 solution (red) and base case well configuration (gray) over a production time frame of 1200 days.	161
4.26	Well oil production totals for multiple realization case using the JNT2M1 solution (red) and base case wells (gray) over a production time frame of 1200 days.	162
4.27	Field oil production totals for multiple realization case using the JNT2M1 solution (red) and base case well configuration (gray) over a production time frame of 5174 days.	165
4.28	Well oil production totals for multiple realization case using the JNT2M1 solution (red) and base case wells (gray) over a production time frame of 5174 days.	166
4.29	HuPhiSo oil saturation maps after 5174 days of production for the three hybrid solutions: J1M1H1, J1M1H2 and J1M1H3. Upper left map shows corresponding base case map.	169
4.30	Field oil production totals for multiple realization case using the J1M1H3 solution (red) and base case well configuration (gray) over a production time frame of 1200 days.	170
4.31	Well oil production totals for multiple realization case using the J1M1H3 solution (red) and base case wells (gray) over a production time frame of 1200 days.	171
4.32	Field oil production totals for multiple realization case using the J1M1H1 solution (red) and base case well configuration (gray) over a production time frame of 5174 days.	173
4.33	Well oil production totals for multiple realization case using the J1M1H1 solution (red) and base case wells (gray) over a production time frame of 5174 days.	174
4.34	Original base case wells used in this work (black lines) are drawn together with updated base case wells (blue lines).	179
6.1	Well gas production rates and totals.	211

6.2	Well oil production rates and totals.	212
6.3	Well water production rates and totals.	213
6.4	Well liquid production rates and totals.	214
6.5	Discretized well bottom-hole pressures for the four production periods comprising the entire production time frame.	215
6.6	Discretized water production rates for the four production periods.	216
6.7	Well and group controls and constraints for the two Eclipse simulation cases 1 and 2.	216
6.8	Discretized well controls used in approximated cases 3 to 6.	217

Abbreviations

<i>Symbol</i>	=	<i>definition</i>
AD-GPRS	=	Automatic-Differentiation General Purpose Reservoir Simulator
FOPT	=	field oil production total
GPS	=	generalized pattern search
HJDS	=	Hooke-Jeeves direct search
HOPSPACK	=	hybrid optimization parallel search package
ICD	=	inflow control device
ICV	=	inflow control valve
MLE	=	Martin Linge East
MLO	=	Martin Linge oil reservoir
NPV	=	net present value
WBHP	=	well bottom-hole pressure
WWPR	=	well water production rate

Introduction

1.1 Field development optimization

In this work we develop methodology for optimization of field development work tasks within the petroleum industry. In particular, we develop methodology to solve the well placement and well control optimization problems in an integrated manner. Specifically, this methodology attempts to exploit the clear interdependency between the two optimization problems. An overall goal of the development of this methodology is to serve as decision-support to field development work tasks dealing with well location and production strategy. In this chapter, we present a general discussion of topics related to this goal. The topics presented serve as general background to the more specific optimization work presented in subsequent chapters. In the following, we discuss topics regarding field operations as a set of work tasks and objectives, decisions within petroleum field development, a general introduction to well placement search, and finally, two motivations for how optimization techniques can complement current work processes. The first motivation is based on the general benefit that can be expected by introducing optimization techniques into engineering work processes, e.g., the introduction of a systematic search to complement common engineering approaches often based on heuristics. The second motivation is about taking advantage, and further developing, research ideas and applying these optimization techniques to field operations. In this regard, one of the main contributions of this work is the development of an application based on our developed methodology that was tested on a real field case. Short summaries of the main contributions from this work are presented in the second half of this chapter. At the end of the chapter we provide a broad outline of the thesis.

1.1.1 Overview of field operations

Work tasks and objectives. A petroleum field development project involves a large number of considerations, and therefore subdivision of labor and prioritization of work tasks and objectives are important work aspects within field operations. A common approach within field development is to generate well-defined work scopes that may either treat large parts of the field project, or just target specific topics of interest, e.g., decisions

regarding well placement configuration and production strategy. Specific work tasks are defined based on further differentiation of the scopes, and objectives and assessment criteria are developed for the particular topics. Considering both short and long term goals, the developed objectives and criteria are set in prioritized order and used in subsequent decision-making processes (Bratvold et al., 2014). For the development of a petroleum asset, the set of technical and operational objectives, as well as the preferences and the overall strategy for the management of the field, are introduced by the field operator in conference with partners sharing the production license. Given the stated objectives, expert engineering practice is then engaged to solve the various field development work tasks such as final well placement configuration and reservoir production strategy.

Problem complexity. The development of an offshore petroleum asset involves a wide range of decisions, from the location and power source of the production platform, to the design of the subsea facilities, to the choices regarding how and where the produced hydrocarbons should be stored and exported. The large number of decisions within a field development project, and their individual complexity, makes field development optimization a very difficult problem to solve as a whole, e.g., in an integrated manner through mathematical programming. Within the scope of well placement, decisions involving the type, number and location of wells are chiefly based on the size, composition and location of the hydrocarbon accumulations, and on the geological description and main drive mechanisms of the reservoir. In this work we regard these decisions as having medium-to-long-term time horizons, i.e., we expect these decisions to span several years and up to the life-cycle of the petroleum asset. These decisions are challenging not only because they have to be made within the larger decision space of the development project, but also because the evaluation of the different production scenarios often requires a large amount of computation. These production scenarios are costly to evaluate because a numerical simulation of reservoir fluid flow is often needed to predict the total fluid volumes produced by the different well configurations. Due to the significant number of decision variables, the high cost of simulation, and the large impact on expected oil recovery, an oil company is likely to allocate, and spend, a substantial amount of resources searching for suitable well placement configurations.

Well placement search. An oil company is likely to perform a significant number of studies to find sound well locations and subsequent well designs for the development of a field. A substantial amount of the testing and analysis that is performed will rely on the fluid flow predictions obtained using reservoir simulation models of the petroleum asset. Embedded in these models are the physical and compositional states of the hydrocarbon accumulation, the measured properties of the fluids in place, and various other geological data, e.g., the structural setting, that describe the reservoir. The simulation of these models plays a fundamental role in the evaluation of possible well placement configurations. Testing for different well configurations is usually a manual process, constrained by different operational considerations such as platform location, and the restrictive cost of simulating each well placement scenario. Once field objectives and considerations are defined, an engineering team will usually start well placement search studies based on estimates of fluids in place, expert judgment regarding the geology of the field, and experience, to produce a final well configuration. An optimization procedure, on the other hand, will use a set of mathematical principles (e.g., iterative improvement) to improve on

an initial well placement configuration¹. If so specified, its search will propose new locations based on which changes to the current well placement configurations that cause the greatest increase to a pre-defined objective function, e.g., oil recovery, net present value, or any other type of measure. In our case, the search procedure constitutes a simulation-based optimization. Crucially, the inclusion of the simulation process means the expert knowledge assembled by the operator, and embedded within the reservoir model, is an implicit part of the search routine.

Complementary approach. In this work, a main research goal is to complement the common engineering approach for solving for well placement and production controls by applying optimization techniques that are capable of exploring the solution space of these problems in a systematic manner. The premise is that a well-defined cost function and solution space can be established. This means that both the objective and associated constraints can be precisely defined, and importantly, that they together yield a reasonable representation of one, or a combination of several, of the field development work tasks of interest for the petroleum field operator, e.g., well placement, well production controls or facility configuration. The idea is then that an optimization procedure can potentially (and hopefully efficiently) solve the problem and provide the operator with valuable information regarding how to best develop and manage the asset. In this manner, the application of optimization procedures to aid field development work tasks can serve as a highly-customizable tool for decision-support.

Pilot application. In this thesis we start with developing optimization procedures using relatively simple example cases. Then, based on the methodology developed in the first part of the work, in the second part of this thesis we perform a pilot application effort using a real field case. This pilot application work tries to show that a systematic search procedure can possibly yield useful information to the work process of an industry operator. In this application, the aim is to demonstrate how the procedure can serve as a decision-support tool to the engineers in charge of developing the well strategy for a real field. Potentially, the field development team can benefit by comparing, and possibly also complementing, engineering-based solutions with solutions and/or information obtained using a systematic search. As such, the application of simulation-based optimization is meant to serve as a complementary tool to current industry work processes.

As we will discuss next, decisions regarding well placement are intrinsically related to production strategy and well control settings. This proposition is fundamental for the methodology developed in this thesis, and will be thoroughly explored in subsequent chapters. Below we preface this exploration with a discussion on the interdependence of different decisions within field development, and of the potential gains from solving integrated problems through optimization.

1.1.2 Motivation for optimization

The introduction of optimization procedures as decision-support tools within field operations has the potential to yield substantial profits. As mentioned earlier, the gains are likely to be significant if we allow the otherwise manual exploration procedure commonly used to find field development solutions to be complemented by a systematic search procedure. Another, more intrinsic, source for potential gain may come from exploiting the

actual structure of the field development problem itself, which can be described as a master problem incorporating a host of challenging, and interrelated, subproblems. Below we explore how we can apply optimization procedures to capitalize from these two sources.

Systematic search. For complex problems within field development, introducing a systematic procedure as a complementary tool in decision-making can be beneficial if the search for solutions to these problems is often too reliant on experience. Furthermore, benefits will almost always be achieved if existing configurations are frequently only improved upon by using engineering heuristics and ad hoc techniques. In these cases, the introduction of optimization procedures can complement a manual exploration of the solution space (such explorations are common for well placement problems (Zandvliet, 2008)) with a systematic search that can be both more efficient and may reduce the risk of possible bias in decision-making. Another advantage is that, in particular for field development work tasks that deal with design (e.g., well placement, pipe network and/or facilities), optimization procedures can in a straightforward manner be adapted to the different stages of the associated decision process. For example, at early decision stages, an optimization procedure with a given parametrization can search a very large solution space to provide a set of feasible configurations for further consideration, while at later stages in the process the same procedure can be adapted to fine-tune an existing configuration.

Model-based optimization in integrated systems. A particular interesting target for model-based optimization is the typical development plan for green fields. The reason is that such plans usually include a wide range of decisions, e.g., well placement, reservoir production strategy and facility layout, that may be treated simultaneously for a significant gain in profit (see Juell et al., 2010, for a discussion of integrated optimization (I-OPT), though that work uses a larger field operation perspective than the one taken here). In this work, we regard development plans of petroleum assets (and subsequent problem formulations based on these) as “integrated” in the sense that these plans and problems represent systems that are composed of interdependent subsystems. Because of the tight connection between different parts of the value chain represented by these plans, a holistic view may be required to improve the design and operation of these complex assets (Rahmawati et al., 2010). However, though the gains that could be achieved from exploiting the various decision interdependencies in such a plan are potentially large, optimization procedures that can treat multiple field development decisions at the same time are difficult to implement. With respect to problem formulation within mathematical programming, a complicating factor is that work tasks within the field development plan are often represented using different types of decision variables (e.g., integer variables for pipe-network settings, such as valves, as well as for the configuration of facility components, such as pumps and compressors, and continuous variables for well controls, such as pressures and rates). Computationally, an important reason for this difficulty is that the models used to describe the various physical processes involved in the different decisions very often require extensive calculations. For example, in the case of fluid flow from and within the reservoir, field case reservoir models commonly describe three-phase flow over highly heterogeneous porous media (e.g., permeability and porosity), involve complex geological structures and faults, and operate with advanced production settings (e.g., gas lift and well group fluid flow targets and constraints). These factors make the full-scale reser-

voir models computationally demanding, and thus difficult to implement within iterative optimization procedures that require at least one, though may require several, reservoir simulations for each cost function evaluation; and often many cost function evaluations for the entire procedure.

Note that we recognize that model uncertainty is an important topic when dealing with decision-making based on model predictions in general, and with model-based optimization in particular, but this topic is currently outside the scope of this work. This dimension is therefore not included in our general presentation. Future work, however, will also take into account the inherent error in model description during the optimization effort.

Problem interdependency. Due to the significant computational load required by these models, a common way to design an optimization effort is to target a single type of decision, such as well placement, as a separate problem (Echeverría Ciaurri et al., 2011b) and to treat interdependent problems, such as optimal production strategy, in the case of the larger well placement decision, as constant. In this regard, fixed control settings, or some type of reactive control strategy based on a heuristic treatment of production controls, e.g., shut-in of wells at some water production ratio with respect to oil, are often used during well placement optimization. However, an optimization effort that seeks to maximize the gain inherent in the complexity of the well placement decision needs to take into account improved, and preferably optimal, well control settings at the various well configurations, during the search for optimal well placement. In the type of production scenarios commonly defined for optimization purposes, production settings may refer to a series of optimized well control types operated over time, from rate, or tubing or bottom-hole pressure, settings associated with flow through single valves in wells, to the improved operation of sophisticated production systems involving well bores with multiple inflow control devices. In summary, the main premise for the work in this thesis, is that, to maximize the gain from the master-subproblem relationship between these two subsystems, the well placement and well control problems need to be treated within the same optimization effort. To this end, the work in this thesis develops a joint approach where both well placement and controls are solved for in an integrated manner.

Structure of thesis work. The first part of the work presented in this thesis develops a methodology for simulation-based well placement and control optimization. In this part we show that, in terms of cost function value, the joint approach outperforms sequential procedures that use fixed and reactive controls when optimizing the location (these procedures solve for controls only once at their final well placement configurations). That work is implemented using vertical wells and relatively simple reservoir cases. In the second part of this work, we extend and apply this methodology using a field case model provided by IO Center Industry Partner Total E&P Norge AS. In this part we compare the joint approach against a sequential procedure when optimizing the location and production of several horizontal wells using a single realization of the field model. This pilot application is our first attempt at applying research methodology to provide decision-support to well placement and production strategy work tasks within an operations environment. Substantial focus has been put on extending and adapting the previously developed methodology from fundamental examples using simple cases, into an application that can deal with a significantly more complex and challenging real field problem. An effective collaborative work process was established thanks to steadfast contributions from field operator Total

E&P Norge AS. This collaborative effort allowed our pilot application to interface, receive feedback and offer complementary decision–support to the development work of the Martin Linge North Sea field.

Next we provide a series of short summaries describing the main contributions from each chapter in this thesis. We end this chapter by providing an outline of the contents of each section in chapters 2, 3 and 4 of this thesis. (A general note describing why background work on the various topics treated in this thesis is not presented in this chapter, but rather introduced as independent literature reviews in chapters 2 and 3, is given in Section 5.1, page 182).

1.2 Summary and contributions from each chapter

Here we provide short summary points based on contributions from each of the chapters in this thesis.

1.2.1 Main contributions from Chapter 2: Joint Optimization of Oil Well Placement and Controls

Development of methodology. The main purpose of Chapter 2 is to develop methodology. The methodology presented in this chapter treats the search for optimal well placement coordinates and optimization of well control settings as separate parts of an integrated problem, allowing each part to be solved efficiently using adequate optimization routines.

Definition of core concepts. In Chapter 2 we define the joint and sequential approaches used for optimization of well placement and control in this work. The joint approach solves the well placement and control problem by nesting the control optimization within the well placement search. More specifically, the joint approach is defined as the search for optimal well placement configurations while optimizing for well controls at each well placement iterate. Importantly, because the search for optimal well locations is conducted within the space of control–optimized well configurations when using the joint approach, the well placement solution is inherently coupled with the (local) optimality found for the well control part. In the sequential approach, initial control settings are kept fixed while optimizing for well placement. Well controls are then optimized at the location found by the fixed–control well placement search.

Introduction of pattern search methods. Derivative–free optimization methodologies based on pattern search are tested as a more mathematically sound alternative for well placement optimization compared to, e.g., stochastic search procedures commonly used in the literature. In Chapter 2 we show these methods can be efficiently applied to deal with the well location part of the well placement and control problem. These methods possess the advantage of being supported by local convergence theory, as well as of having been seen to perform satisfactorily on relatively non–smooth cost functions. Finally, they have the benefit of being relatively straightforward to implement within a distributed environment. The pattern search algorithms considered in both cases in this chapter are

Hooke–Jeeves direct search (HJDS), generalized pattern search (GPS) and the hybrid optimization parallel search package (HOPSPACK).

Application to example cases. The methodologies are applied to two example cases of oil production using water–flooding, with net present value (NPV) as objective function in both cases. In both examples there are five wells (one injector and four producers in the first case, and two injectors and three producers in the second case). The controls for all five wells are optimized in the two cases. In the first example, only one well (the injector) location is optimized, while in the second case all well locations are optimized.

Exhaustive search results. Due to its low dimension, exhaustive search approximations of the optimization spaces corresponding to the sequential and joint approaches are performed in the first example case. As expected, the joint scheme used in the exhaustive search outperforms the sequential methodologies, even after the additional control optimization step. The joint approach yields an increase of 4.2% and 5.9% in net present value compared to the sequential fixed and reactive approaches.

Smoothing of optimization surface. From the exhaustive search we observe that the surface associated with the fixed control strategy is much rougher than the surfaces obtained with the other strategies. The smoothing of the optimization surface with respect to the well placement variables occurs because the performance of wells in less promising locations can be improved, sometimes significantly, by optimizing the well controls.

Test of optimization techniques. Using computed values from the exhaustive search, we test the pattern search techniques starting from 12 different initial well configurations. After the additional optimization, the average optimized NPV by the joint approach is 10.3% and 6.1% larger than the average optimized NPVs from the fixed and reactive approaches, respectively.

Optimization solutions. In the second example, both the location and control of two injectors and three producers are optimized. Optimizations are run nine times with different initial well placement configurations. In terms of NPV, the sequential fixed and reactive strategies clearly under–perform the joint approach. The average (maximum) NPV over all of the runs obtained with the joint approach is 18.2% (14.4%) and 20.6% (7.3%) higher than with the sequential fixed and reactive schemes, respectively. The average number of simulations required by the joint approach is, however, about one order of magnitude higher than that needed by the sequential methodologies.

Use of parallel implementations. The higher computational demand required by the joint approach, compared to the sequential approaches, was mitigated through the use of parallel implementations of the pattern search algorithms.

Extension of approaches. In Chapter 3 we extend the work in Chapter 2, with the aim of being able to apply the developed methodology on a real field case.

1.2.2 Main contributions from Chapter 3: Joint Optimization Applied to a Real Field Case

Main task of chapter. The main task of Chapter 3 is to extend the developed joint and sequential approaches into an optimization procedure that can treat a well design problem

provided by IO Center Industry Partner Total E&P Norge AS. The well design problem involves a reservoir that is part of the Martin Linge field located on the Norwegian Continental Shelf.

Targets and strategy for field case application work. The beginning of Chapter 3 describes the main targets for this field case application work (spanning both chapters 3 and 4). We also describe the two main strategy components for how to develop the research work into an application that can, through the use of a work model, be tested on a real field case.

Work process loop. A work process loop is created to organize the application work involving both technical and procedural issues. The work process loop is the actual execution of the strategy defined for how to extend our work into an application that can treat a real field case. Some of the challenges encountered during the application effort are briefly discussed when describing the different work process loop stages.

Optimization framework. Roughly, the implementation to treat the real field case problem, developed from the methodologies in Chapter 2, consists of a collection of algorithms, solvers, and code extensions coupled with a reservoir simulator. In this work, this implementation is presented as an optimization framework. The various parts of the framework are presented, focusing on the function of each part, and how they have been used to extend the developed methodology. Ongoing work to enhance the framework, and suggestions for further work ahead, are also described.

Collaboration effort. A description of collaboration effort with various IO Center Research Partners, as well as with IO Center Industry Partner Total E&P Norge AS, is given. The different contributions from the various partners have been crucial to increase the functionality of previous work to deal with a real field case. Furthermore, an efficient collaboration with Total E&P Norge AS has been important to both adapt the optimization procedure and test the obtained solutions according to a field development perspective.

Introduction to field case. This field case model used in this work is a particular realization of the Martin Linge oil reservoir provided by the operator Total E&P Norge AS. For application within our optimization procedure, the field case model, originally implemented in the industry-standard reservoir simulator Eclipse, has been transferred to the AD-GPRS research reservoir simulator. This transfer was followed by an extensive validation process. The field case model as well as approximations introduced during the transfer and validation process are discussed.

Application effort. The optimization framework presented in this work is a first attempt at applying the developed methodologies within a real field development scenario. The framework launches an iterative search procedure that searches for improved well trajectories while taking into account various constraints on well placement coordinates. Within the framework, both joint and sequential approaches are used to optimize for well placement and controls using the approximated work model of the field case. The well placement part of the procedure is subject to well-length, well-orientation and inter-well distance constraints. Parameters for the various constraints were specified in close collaboration with the Martin Linge field development team.

Solutions development using AD-GPRS work model. Four solutions for well placement and control optimization have been developed using the joint and sequential ap-

proach each, yielding a total of eight solutions. The solutions corresponding to each approach were developed using different formulations for well-length constraint handling, including a configuration where the maximum well-length constraint was imposed only after the optimization had ended. Optimal well trajectories and controls obtained using the joint and sequential approaches with different configurations of well-length constraint handling yield mean increases in field oil production total (FOPT) of 33 and 26%, respectively.

Cost of optimization approaches. On average, the total number of reservoir simulations required by the joint approach is approximately 7 times higher than the number of reservoir simulations required by the sequential approach.

Transfer of solutions. The obtained solutions are transferred to the Eclipse field case model. Two types of transfer are specified, involving both the well placement and well control parts of the solutions, and only the well placement part along with the original production schedule simulator settings. The solutions from these two types of transfer are tested on the Eclipse field case for the limited scope of the optimization procedure, i.e., using only a single realization and a reduced time frame. For the first type of transfer, joint and sequential solutions yield mean increases in FOPT of 24% and 19%, respectively. Joint and sequential runs corresponding to the second type of transfer yield mean FOPT increases of 25% and 20%, respectively. Further testing in Chapter 4 involves assessing the performance of the solutions for a wider set of field development considerations than those specified in the optimization scope.

1.2.3 Main contributions from Chapter 4: Testing of Solutions on Field Case Model

Individual testing of solutions. An important goal from the testing of the solutions is to make result information readily accessible to the field development work process of the operator. Individually, all well configurations from the solutions obtained using the optimization framework are plotted, and their recovery studied, using saturation maps recommend by our Industry Partner. Furthermore, the performance of each well is examined using particularly relevant production profiles (with respect to main drive mechanism) for the development of the Martin Linge oil reservoir.

Collective testing of solutions. Collectively, solutions from the optimization procedure that have been run on the Eclipse field case model, are analyzed using both the 1200 day production time frame used in the optimization procedure, and the planned production horizon for the reservoir of 5174 days. Furthermore, for both production time frames, each solution was tested over a multiple realization case scenario.

Test results using the 5174 day production horizon. When the joint and sequential solutions are run over the original base case time frame of 5174 days, their mean gains in FOPT decrease to 12 and 9%, respectively.

Test results using multiple realizations. All joint and sequential solutions (and the base case well configuration) were tested over a set of 11 model realizations using both the 1200 day and the 5174 day production time horizons. For each time horizon, the expected FOPT for each of the solutions is compared to the expected FOPT for the base case

configuration. The best solution for the 1200 day production horizon yields an increase in expected FOPT of approximately 7% with a standard deviation of 0.215 compared to a standard deviation of 0.119 for the base case configuration. (Notice that throughout this thesis, all results are expressed in terms of percentage due to confidentiality reasons. Therefore, the standard deviations associated with the FOPT results presented here also have percentage points as unit.) When using the 5174 day production horizon, the best solution yields an increase in expected FOPT of less than 3% with a standard deviation of 0.163 compared to a standard deviation of 0.110 for the base case configuration.

Test of hybrid solutions. Additionally, some simple heuristic changes were made to a joint solution to test a workflow where solutions obtained from an optimization procedure are later modified based on engineering judgment. The idea is that some solutions may be improved by applying concrete engineering experience codified as heuristic rules, e.g, a rule that was tested consisted of interchanging low-performing wells by their base case analog or by a better-performing well from another solution.

Acknowledgement

We would like to thank Martin Linge field operator Total E&P Norge AS, as well as license partners Petoro AS and Statoil AS, for allowing us to test our methodology and developed optimization framework on some of their field case data. This data has been used in chapters 3 and 4 in this thesis. The data consists of an Eclipse model of the Martin Linge oil reservoir, in addition to various Petrel projects and scripts necessary to build new well trajectories based on our solutions, and to plot the saturation maps in Chapter 4. Throughout this project, Total E&P Norge AS has kindly assisted us in the use of the various software platforms, suggested various types of analysis to treat the different results, and provided feedback on the various solutions. We are grateful for their continuous support and extensive collaboration.

Disclaimer

It should be noted that any of the content presented in this thesis, including but not limited to statements and opinions, are the sole responsibility of the author. This applies to all information presented in this thesis about the Martin Linge field and oil reservoir, as well as any presentation, discussion and/or conclusion based on data obtained from our application pilot.

Moreover, model data presented in this thesis stem from a version of the model obtained at the start of this work. Due to changes and updates, these data may no longer apply to recent versions of the model. The same is true for any information provided that is related to the Martin Linge field development. Finally, analyses and results presented in this thesis have been reviewed by Total E&P Norge AS, and partner approval has been granted. However, Total E&P Norge AS is not responsible for any of the conclusions or suggestions presented in this thesis, nor do these in any way reflect or represent the official positions or opinions of Total E&P Norge AS, nor any of the partners in the Martin Linge development.

1.3 Thesis outline

Table 1.1 gives an overview of the content of each section in chapters 2 to 4 in this thesis.

Table 1.1: Outline describing the content of chapters 2 to 4 in this thesis.

Chapter 2: Joint Optimization of Oil Well Placement and Controls	
Section 2.1	Introduction A general presentation of the well placement and control problem is given; fixed and reactive production strategies are introduced.
Section 2.2	Problem statement The optimization problem treated in this work is presented in detail.
Section 2.3	Optimization methodology The joint and sequential approaches are properly introduced.
Section 2.4	Example cases The application of the developed methodologies to two example cases is described; results from exhaustive search approximations are given; optimization runs using different initial well placement configurations are launched to compare the joint versus sequential approaches.
Section 2.5	Concluding remarks Summary and suggestions for further work are provided.
Chapter 3: Joint Optimization Applied to a Real Field Case	
Section 3.1	Targets and strategy for application development Targets and strategy components for application development are presented; a work process loop is specified.
Section 3.2	Optimization framework The implementation of methodology is presented as an optimization framework; parts and properties of the procedure are discussed.
Section 3.3	Field case and validation work The Martin Linge field case is presented; model transfer and validation issues are treated in detail.
Section 3.4	Optimization work The overall problem formulation for the field case application work is presented; various constraints for the well placement part of the problem are defined; the procedure for constraint handling is also described.

- Section 3.5 Optimization results**
Main results from the optimization work are presented; the process of transferring the obtained solutions to the field case model is discussed, and some related results are presented.
- Section 3.6 Discussion and suggestions for further work**
An overall discussion of application work is given; suggestions for both technical and procedural improvements are provided.

Chapter 4: Testing of Solutions on Field Case Model

- Section 4.1 Test results from solution cases**
Eight test cases are established; results are compared for both the 1200 day and the 5174 day production horizon.
- Section 4.1.1 Individual analysis: Final well configurations**
Individual well placement configurations are drawn; saturation maps and production profiles are given for each case.
- Section 4.1.2 Collective analysis: Total oil production values**
Total field and well oil production values for well placement solutions are treated collectively; tables and figures showing production for each well and for the entire field are presented.
- Section 4.1.3 Increases in recovery versus changes in well length**
The correlation between well length changes and increases in total oil production is studied.
- Section 4.2 Field model tests on multiple realizations**
Each solution case is tested over 11 model realizations; both the 1200 day and the 5174 day time frames are used; mean oil production total along with standard deviations are given.
- Section 4.2.1 Solution tests on multiple realizations: 1200 days**
Multiple realization results from each of well placement solutions are provided for the 1200 day production time frame.
- Section 4.2.2 Solution tests on multiple realizations: 5174 days**
Multiple realization results from each of well placement solutions are provided for the 5174 day production time frame.
- Section 4.2.3 Hybrid solution tests on multiple realizations**
A heuristic procedure is ultimately introduced to modify solutions and account for important factors not specified during optimization, e.g., geological uncertainty.
- Section 4.3 Final topics on field case application**
Final comments on field case application are given, as well as suggestions for further development of the optimization framework.
-

Joint Optimization of Oil Well Placement and Controls

This chapter develops the main methodology for this work. Specifically, it details the core definitions for the use of the joint and sequential approaches, and lays the foundation for the field case application work presented in the following chapters. The content of this chapter corresponds to an article published in Computational Geosciences (Bellout et al., 2012) in collaboration with Dr.Eng. David Echeverría Ciaurri, then at Stanford, now at IBM T.J. Watson Research Center; Prof. Louis J. Durlofsky at the Department of Energy Resources Engineering at Stanford University, and Profs. Bjarne Foss and Jon Kleppe at the Department of Engineering Cybernetics and the Department of Petroleum Engineering and Applied Geophysics, both at NTNU, respectively.

Author contributions. *The author of this thesis has performed all computations, as well as further treated and created all representations of the data (i.e., tables and figures) derived from these computations. However, the creation of the cases, the overall analysis of the data, and the development of conclusions, have all been performed in close collaboration with Prof. Durlofsky and Dr.Eng Echeverría Ciaurri, under the guidance of Profs. Foss and Kleppe. In particular, Section 2.3 (Optimization methodology), and the parts regarding result analysis and discussion, sections 2.4.1 and 2.4.2, have received significant contributions from the co-authors, in terms of theoretical review and text input. Additionally, substantial contributions from the co-authors have been the description of governing equations for subsurface flow given in Section 2.2, by Prof. Durlofsky, and a re-structuring and refinement of the mathematical explanation of the joint approach performed by Dr.Eng. Echeverría Ciaurri in Section 2.3.4.*

Abstract

Well placement and control optimization in oil field development are commonly performed in a sequential manner. In this work we propose a joint approach that embeds well control optimization within the search for optimum well placement configurations. We solve for well placement using derivative-free methods based on pattern search. Control optimization is solved by sequential quadratic programming using gradients efficiently computed through adjoints. Joint optimization yields a significant increase, of up to 20%

in net present value, when compared to reasonable sequential approaches. The joint approach does, however, require about an order of magnitude increase in the number of objective function evaluations compared to sequential procedures. This increase is somewhat mitigated by the parallel implementation of some of the pattern search algorithms used in this work. Two pattern search algorithms using eight and 20 computing cores yield speedup factors of 4.1 and 6.4, respectively. A third pattern search procedure based on a serial evaluation of the objective function is less efficient in terms of clock time, but the optimized cost function value obtained with this scheme is marginally better.

2.1 Introduction

The development of new fields for oil and gas production is increasingly complicated and expensive. Sustaining profitable production in mature fields, where water production rates are often high, also poses a challenge. For both sets of problems, it may be difficult to achieve adequate returns on investment using traditional (heuristic) production management techniques. There is, therefore, a growing interest in the development of efficient and effective simulation-based optimization procedures for well planning and operation.

This work focuses on maximizing revenue from oil production using water-flooding by optimizing medium-to-long-term (i.e., multi-year time frame) field management operations such as well placement and well control scheduling. Water-flooding, where the oil in the subsurface formation (reservoir) is driven towards production wells by a moving waterfront created by water injection wells, is a common procedure for oil production. Substantial oil volumes are often bypassed during water-flooding due to the existence of complicated geological conditions, such as high-flow regions and faults, in the reservoir. Thus, for water-flooding to be effective, the locations and control schedules of injectors and producers must be selected in an optimal manner (by control schedule we mean the well rates or bottom-hole pressures as a function of time). Here our objective function is the net present value of the asset, though other cost functions such as total oil recovered could also be used. In either case the cost function is computed by means of the numerical solution of the system of partial differential equations that describes fluid flow in the reservoir. The required simulations are very often computationally demanding, which poses challenges for optimization.

Under current procedures, the determination of well placement and well control is generally treated in a sequential manner. This means an optimal well placement configuration is first determined using a given (and thus, not optimal) strategy for handling the well controls. These controls are then optimized at the well locations found in the first step. A relatively popular (and heuristic) control strategy, which is often referred to as “reactive control”, entails closing (“shutting in”) production wells according to an economic threshold that depends on the oil price and the water production cost. This economic threshold is translated into a water-cut limit which, once reached, triggers the closure of the corresponding well for the rest of the production time frame. A reactive control strategy can be reasonably effective but is clearly suboptimal as it does not impact injection well settings and handles production wells as either fully open or closed. Any approach that does not consider well location and control jointly cannot be expected to yield optimal solutions, since it does not capture the interdependency between the well

configuration and the associated controls.

In this paper we propose a joint approach for optimization of well position and control settings. In our approach the two different optimizations are considered in a nested fashion. The outer loop involves a well location optimization, while the inner loop is based on optimizing well controls for fixed well positioning. The objective function at the outer loop (for given well locations) is an optimized value of the cost function considered in the inner optimization of the well controls. This scheme results in the solution of the outer optimization satisfying optimality conditions not only for the well placement problem but also for the well control part, because the optimal nature of the solution with respect to the controls is intrinsically inherited in the algorithm. Hence, this joint approach can be used to compute solutions that improve upon those achieved using sequential methodologies. The computational cost associated with the joint approach is, however, much higher since every upper-level function evaluation requires the optimization (not necessarily to full accuracy) of the lower-level problem.

The nested approach has been devised as a combination of methodologies that separately solve the two different types of problems, well placement and control, that constitute the joint optimization. Well control optimization is commonly stated in terms of continuous variables (well flow rates or pressures), and in some cases (bound-constrained optimization problems) has been observed to present smooth optimization landscapes with multiple optima but similar cost function values (see e.g., Jansen et al., 2005 and Echeverría Ciaurri et al., 2011a). This observation is for problems with linear constraints (which have been studied the most), and its validity is unclear for more general cases. On the other hand, well placement optimization is often formulated as an integer optimization problem (where integers correspond to specific grid blocks) with non-smooth objective functions (see e.g., Onwunalu and Durlofsky, 2010) containing multiple optima with significantly different cost function values. This non-smooth character is generally related to the strong variability (heterogeneity) in subsurface flow properties. Therefore, many of the existing well placement optimization procedures attempt a more global search. Consistent with these observations, well control optimization is often addressed using gradient-based techniques (where gradients are computed rapidly via adjoint procedures; see e.g., Jansen et al., 2005, or Sarma et al., 2006), while well placement optimizations usually use derivative-free algorithms or stochastic search procedures (see e.g., Yeten et al., 2003; Onwunalu and Durlofsky, 2010). Derivative-free and stochastic optimization approaches ordinarily require parallel computing implementations for efficiency. We note, however, that gradient-based techniques have been applied for well placement (e.g., Sarma and Chen, 2008; Zandvliet et al., 2008), and stochastic search has been used for well control (e.g., Echeverría Ciaurri et al., 2011a), so our observations here should not be viewed as absolute.

To our knowledge, no research has been published addressing in detail the joint optimization of oil well placement and control. There have, however, been approaches that use the reactive control strategy described above in well placement optimizations (e.g., Zandvliet et al., 2008). The work introduced in Wang et al. (2007), and later enhanced in Zhang et al. (2010) and Forouzanfar et al. (2010), aims primarily at well placement, and integrates indirect mechanisms for optimizing well controls. The method described in that work provides a comprehensive optimization framework, but it involves a number of

heuristics and does not treat explicitly location and control as optimization variables. The approach presented in this paper attempts to address the complicated joint well placement and control optimization problem from a mathematically sound perspective.

Other variables besides the location and controls for each well, such as the number of wells and the length of the water–flooding process, could also be included in the optimization. These variables are much more difficult to treat, however, since the number of variables in the corresponding optimization problem depends on these parameters. For example, optimizing the number of wells could be performed by adding a new set of categorical optimization variables that allow for the activation/deactivation of each well (Echeverría Ciaurri et al., 2012). The inclusion of variables of this type would significantly increase the complexity of the optimization problem. Another important effect not included in this work is uncertainty in the reservoir model; i.e., the optimization approaches studied here do not involve stochastic programming considerations. A general method for optimizing well location under uncertainty has been developed by Wang et al. (2012), and this approach could be extended in a straightforward manner to also include well controls.

This paper is structured as follows. The governing equations for the flow of oil and water in subsurface reservoirs are given in Section 2.2. This section then introduces the general problem statement and the specific formulations for the well control and the well placement parts of the optimization procedure. Next, the joint and sequential approaches used to solve the coupled system are described in Section 2.3. These approaches are applied to two example cases in Section 2.4. The first case addresses the control optimization of one injector and four producers, and the optimal positioning of the injector. For this case we are able to perform exhaustive computations, which enable clear assessments of the various optimization procedures. In the second case the well position and controls for three producers and two injectors are optimized. Section 2.5 provides a summary and some suggestions for future research.

2.2 Problem statement

In this section we briefly describe the flow simulations used to evaluate well location and control scenarios. The general optimization problem treated in this work is then presented in detail.

2.2.1 Governing equations for reservoir production

Hydrocarbons such as oil and gas are found within porous rock in subsurface formations. The equations that describe fluid flow in the reservoir are derived by combining expressions of mass conservation with constitutive and thermodynamic relationships. For clarity, our brief description here entails several simplifications (such as the assumption of incompressible flow), though in the problems considered later compressibility and other effects are included. See, e.g., Aziz and Settari (1979) or Ertekin et al. (2001) for details on the flow equations and numerical discretizations.

We consider two–phase immiscible systems containing oil (o) and water (w). Mass

conservation for each fluid i (where $i = o, w$) is given by:

$$\nabla \cdot \mathbf{u}_i + q_i = -\frac{\partial}{\partial t} (\phi S_i), \quad (2.1)$$

where \mathbf{u}_i is the Darcy velocity of phase i , q_i is the source/sink term, ϕ is porosity (volume fraction of the rock that can be occupied by fluids), S_i is the saturation of phase i , and t is time. Darcy velocity is expressed as:

$$\mathbf{u}_i = -\mathbf{k} \frac{k_{ri}}{\mu_i} \nabla p, \quad i = o, w, \quad (2.2)$$

where p is pressure (here assumed the same for both phases), μ_i is the viscosity of phase i , \mathbf{k} is the absolute (rock) permeability tensor, and $k_{ri}(S_i)$ is the relative permeability of phase i .

Combining (2.1) and (2.2) yields:

$$\nabla \cdot \left[\mathbf{k} \frac{k_{ri}}{\mu_i} \nabla p \right] - q_i = \frac{\partial}{\partial t} (\phi S_i), \quad i = o, w. \quad (2.3)$$

Eq. (2.3), written for $i = o, w$, along with the saturation constraint ($S_o + S_w = 1$), define the flow problem. Model sizes for the numerical solution of (2.3) usually range from tens of thousands of grid blocks for small models, to several hundred thousand or millions of grid blocks for large models. A typical model might require several hundred time steps.

Reservoir models are coupled to well models (via the source term q_i) to enable the computation of the volumes of fluids produced and injected at each time step. See Peaceman (1978) or Ertekin et al. (2001) for details on this coupling. Current well designs may involve vertical, horizontal, deviated and multilateral wells. These wells can be controlled by specifying either rates or bottom-hole pressures (BHPs). In this work, we will consider only vertical wells and use BHPs at various time intervals for the well control optimization parameters. The simulator used in this work is Stanford's General Purpose Research Simulator (GPRS; Cao, 2002).

2.2.2 Optimization problem

In our examples we consider oil–water systems with production driven by water injection. We seek to determine the optimal locations and BHP controls for a specified number of production and water injection wells using an optimization procedure based on a joint, rather than a sequential, approach.

The optimization problem studied here is defined as follows:

$$\min_{\mathbf{x} \in \mathbb{Z}^{n_1}, \mathbf{u} \in \mathbb{R}^{n_2}} -\text{NPV}(\mathbf{x}, \mathbf{u}) \quad \text{subject to} \quad \begin{cases} \mathbf{x}_d \leq \mathbf{x} \leq \mathbf{x}_u \\ \mathbf{u}_d \leq \mathbf{u} \leq \mathbf{u}_u \end{cases}, \quad (2.4)$$

where \mathbf{x} denotes the discrete well placement variables and \mathbf{u} are the continuous well control variables. Well placement variables are intrinsically real but are often treated as integers, since reservoir simulators require wells to be assigned to discrete grid blocks in the model. Consequently, in many cases, and in this work, \mathbf{x} is defined as discrete-valued. All wells in this work are assumed to be vertical, hence well positions can be

stated in terms of discrete areal coordinates (x, y) only. Thus $n_1 = 2(N_p + N_i)$, where N_p and N_i are the number of production and injection wells in the placement optimization, respectively (we could also optimize a subset of wells as in Example 1). In more general cases, additional variables would be needed to describe well locations. For example, the optimization variables might include the perforation interval for vertical wells (if wells are not open to flow over their entire length), or the actual trajectory for deviated wells. The controls over time for each well are represented by a piecewise constant function with N_t time intervals (i.e., well controls are held constant during an interval and then jump to their value for the next interval). Hence, $n_2 = N_t(N_p + N_i)$, assuming N_p and N_i are the number of production and injection wells in the control optimization.

In this work, we deal with bound constraints only. In order to simplify notation we introduce the well position feasible set $X = \{\mathbf{x} \in \mathbb{Z}^{n_1}; \mathbf{x}_d \leq \mathbf{x} \leq \mathbf{x}_u\}$ and the well control feasible set $U = \{\mathbf{u} \in \mathbb{R}^{n_2}; \mathbf{u}_d \leq \mathbf{u} \leq \mathbf{u}_u\}$. Non-linear constraints, which could include rate or water-cut specifications, can be handled using different techniques such as penalty functions or filter methods, as described in Echeverría Ciaurri et al. (2011a).

The objective function considered here is the undiscounted net present value (NPV) of the asset. This NPV accounts for revenue associated with the oil produced as well as for the water-handling costs incurred during production (water costs are incurred as a result of pumping and separation requirements). NPV is defined as follows:

$$\text{NPV}(\mathbf{x}, \mathbf{u}) = \sum_{k=1}^{N_s} \left(\sum_{j=1}^{N_p} p_o q_o^{j,k}(\mathbf{u}, \mathbf{x}) \Delta t_k - \sum_{j=1}^{N_p} c_{wp} q_{wp}^{j,k}(\mathbf{u}, \mathbf{x}) \Delta t_k - \sum_{j=1}^{N_i} c_{wi} q_{wi}^{j,k}(\mathbf{u}, \mathbf{x}) \Delta t_k \right), \quad (2.5)$$

where $q_o^{j,k}$, $q_{wp}^{j,k}$ and $q_{wi}^{j,k}$ are the flow rates of the oil, water produced and water injected for well j at the output interval k , respectively (expressed in stock tank barrels or STB per day, where $1 \text{ STB} = 0.1590 \text{ m}^3$), and Δt_k represents the length (in days) of each of the N_s time steps in the simulation. (Note that N_s does not in general coincide with the number of controls per well, N_t .) The oil price and the cost of water produced and injected are denoted by p_o , c_{wp} and c_{wi} , respectively. Though the problem in (2.4) is stated jointly for \mathbf{x} and \mathbf{u} , it has traditionally been addressed in practice in a decoupled manner (i.e., the well placement part is solved prior to, and independently of, the control optimization). In the next section we discuss some decoupled approaches and propose a methodology for addressing the problem jointly.

2.3 Optimization methodology

This section describes two sequential approaches and introduces a joint approach for solving the well placement and control problem given in (2.4). Both sequential approaches first seek optimal well placements using a predetermined control strategy, and then they optimize the controls for the wells determined in the first stage. Since the control and the well placement optimization problems possess clearly distinct characteristics, it is reasonable to address these two problems using different methodologies. Sections 2.3.1 and 2.3.2

describe the separate optimization problems and approaches corresponding to the continuous (controls) and discrete (well placement) parts of (2.4). Some of the methods presented in these sections will be combined in Section 2.3.4, where we define our approach for the joint problem.

2.3.1 Well control optimization

The production optimization part of the general problem in (2.4) is obtained by fixing the well placement variable to $\mathbf{x}_0 \in \mathbb{Z}^{n_1}$:

$$\min_{\mathbf{u} \in U} -\text{NPV}(\mathbf{x}_0, \mathbf{u}) , \quad (2.6)$$

and corresponds to a problem with continuous variables. The well controls $\mathbf{u} \in U \subset \mathbb{R}^{n_2}$ in this work represent BHPs. For each well, the controls are defined by piecewise constant functions over N_t intervals. The optimization bounds define upper and lower BHP limits for both injectors and producers. Other operational constraints (e.g., minimum oil and/or maximum water production over all wells) can be addressed in an efficient manner by the filter method (Nocedal and Wright, 2006; Echeverría Ciaurri et al., 2011a). The filter method is really an add-on that can be combined with most (derivative-based and gradient-free) optimization algorithms. This technique borrows concepts from multi-objective optimization, and has been observed as a very efficient means for dealing with non-linear constraints (Nocedal and Wright, 2006; Fletcher et al., 2006).

Production optimization problems can be readily solved by gradient-based techniques (Nocedal and Wright, 2006). For example, the gradient-based optimization approach used in this work to solve (2.6) is sequential quadratic programming (SQP; Nocedal and Wright, 2006). The SQP solver used in this work is SNOPT (Gill et al., 2005). Approximating gradients by, e.g., finite differences, typically requires a number of function evaluations on the order of the number of optimization variables. In addition, the quality of the approximation may depend strongly on the simulator settings. Adjoint formulations allow for efficient (though simulator-invasive) computations of gradients (Pironneau, 1974). By means of an adjoint-based procedure, gradients can be computed with a total cost of roughly one solution of a linearized system of ordinary differential equations. Adjoint-based gradient estimations have recently been implemented for optimization problems in the petroleum industry (Brouwer and Jansen, 2004; Sarma et al., 2006). In this work, we use the adjoint formulation in Stanford's General Purpose Research Simulator (GPRS).

Derivative-free methods (Kolda et al., 2003; Conn et al., 2009) have also been shown to perform satisfactorily for the control optimization problem (Echeverría Ciaurri et al., 2011a,b). These methods are applicable for problems with less than a few hundred optimization variables, and they perform fairly efficiently if implemented in a distributed computing environment. We will consider derivative-free methods for the well-positioning part of the general problem introduced in (2.4) in the next section.

It has been observed in previous work (Jansen et al., 2005; Echeverría Ciaurri et al., 2011a) that well control problems similar to (2.6) commonly display multiple local optima having similar cost function values (i.e., the cost function appears to be close to convex in \mathbf{u}). This suggests (though it does not prove) that local optimization approaches,

such as gradient-based techniques, for (2.6) may yield solutions that are acceptable from a global optimality point of view.

Finally, as discussed in the Introduction, reactive control can be applied, as a heuristic alternative to optimization, to address the issue of excessive water production. Under this approach, a production well is kept open (at its lower BHP limit in our implementation) until the revenue from the oil it produces no longer exceeds the cost associated with the water produced; i.e., the well is closed when

$$p_o q_o^{j,k} < q_{wp}^{j,k} c_{wp}, \quad (2.7)$$

where all variables are as defined previously. This relationship is used to determine a corresponding water–cut limit. A producer is permanently shut in once this limit is reached. In practice this treatment often provides satisfactory results (and this approach is inexpensive since no optimization is required), though it is clearly suboptimal since it is based on a simple rule involving only producers. It should also be noted that even though the production strategies obtained by means of reactive control can in some cases be represented by piecewise constant functions, the lengths of the control intervals are not known a priori. Thus, reactive control strategies cannot in general be identified with elements in \mathbb{R}^{n_2} .

2.3.2 Well placement optimization

The well placement optimization part of the general problem originally given in (2.4) is obtained by fixing the well control variable to $\mathbf{u}_0 \in \mathbb{R}^{n_2}$:

$$\min_{\mathbf{x} \in X} -\text{NPV}(\mathbf{x}, \mathbf{u}_0), \quad (2.8)$$

and corresponds to a problem with discrete variables. In general $\mathbf{u}_0 \in \mathbb{R}^{n_2}$, but as noted above, if \mathbf{u}_0 corresponds to a reactive control strategy, it will not necessarily have n_2 components.

Well placement problems are in a sense more challenging than well control optimization problems because reservoir heterogeneity leads to highly non–smooth objective functions containing multiple optima (see e.g., Onwunalu and Durlofsky, 2010). Therefore, the well placement optimization problem does not appear to be as amenable to solution using gradient-based methods because these approaches can get trapped in local minima. There have, however, been procedures presented for (2.8) that use gradients (see e.g., Sarma and Chen, 2008, and Zandvliet et al., 2008). These methods replace the problem with a related (though not necessarily equivalent) problem that has continuous variables.

Most of the derivative–free methods that have been used to date for the solution of (2.8) are based on stochastic search procedures. Examples include genetic algorithms (Goldberg, 1989; Güyagüler et al., 2000; Yeten et al., 2003), stochastic perturbation methods (Bangerth et al., 2006), and particle swarm optimization (Clerc, 2006; Onwunalu and Durlofsky, 2010; Echeverría Ciaurri et al., 2011b). Due to their random component, these search procedures can avoid being trapped in some unsatisfactory local optima. Most of these methods, however, are not supported by solid convergence theory, and consequently they contain tuning parameters that are often difficult to determine.

In this work we propose derivative-free optimization methodologies based on pattern search (Torczon, 1997, Kolda et al., 2003; Conn et al., 2009) as a more mathematically sound alternative for well placement optimization. These methods rely on (local) convergence theory applicable to sufficiently smooth functions of continuous variables. These local convergence results can furthermore be extended to problems with discrete variables (Audet and Dennis, 2000). Examples of these techniques are Hooke–Jeeves direct search (HJDS; Hooke and Jeeves, 1961), generalized pattern search (GPS; Torczon, 1997, Audet and Dennis, 2002), mesh adaptive direct search (MADS; Audet and Dennis, 2006), and bound optimization by quadratic approximation (BOBYQA; Powell, 2009).

Pattern search methods operate primarily through a polling procedure. Polling is accomplished by computing cost function values at points in the search space determined by a stencil which is centered at the current solution. The stencil is normally arranged along the coordinate axes, which results in a coordinate or compass search. In MADS, the stencil orientation is randomly modified after each polling. Pattern search techniques are supported by local convergence theory, but if the initial stencil size is comparable to the size of the search space (which means that, during the first iterations of the optimization, the search involves points that are distant from the initial guess), they can incorporate some global exploration features. We emphasize that global convergence is not achieved using these or other practical procedures. However, in many well placement problems, finding a reasonable local optimum following some amount of global exploration is often sufficient.

In this work, the well placement problem is solved using HJDS, GPS, and a hybrid optimization parallel search package (HOPSPACK; Plantenga, 2009). HOPSPACK is a distributed computing implementation of GPS which can be run in a so-called asynchronous mode to balance the computational load of each node in the cluster (Plantenga, 2009). In asynchronous mode, HOPSPACK avoids “idle” cores by continuously sending new polling points for evaluation. HOPSPACK dedicates a single core to handle the asynchronous assignment of polling points to each computing core. Hence, if, for example, 21 cores are available for HOPSPACK, one of these cores will not be used to evaluate polling points. HJDS is a serial computing procedure that was identified in Echeverría Ciaurri et al. (2011a) as a fairly efficient optimization procedure for oil field problems when distributed computing resources are limited or unavailable.

2.3.3 Sequential approaches for well placement and control optimization

As noted earlier, sequential procedures are commonly used for joint well location and well control optimization. Well placement is optimized first using some “reasonable” control scheme. In this work we will consider two such strategies – fixed and reactive controls. The controls are then optimized for the wells positioned in the first stage. It should be noted, however, that well placement optimization results have been observed to depend to a large degree on the control scheme used (Zandvliet et al., 2008).

Fixed control strategies belong to $U \subset \mathbb{R}^{n_2}$, i.e., the same space explored in the control optimization stage. In our approach these fixed controls correspond to the upper pressure bound \mathbf{u}_u for injectors, and the lower pressure bound \mathbf{u}_d for producers. This strategy

provides maximum injection and fluid production rates at all times. It is important to emphasize that, although fluid production is maximized, oil production is in general not maximized by this strategy if water is also being produced (as it typically is). This fixed control strategy is in general suboptimal because the water front is allowed to proceed without any “steering” (which is achieved when BHPs are varied in time). In addition, it is possible that some wells may be producing essentially all water at full capacity. The reactive control strategy operates with the same pressure settings as the fixed strategy, but it includes the capacity to shut a producer in once it is no longer profitable. In this way, the reactive control approach considers production economics, though it still does not address the efficient injection of water. The reactive control approach is often preferable to the use of fixed controls, but as we will see does not perform as well as the joint optimization procedure.

Algorithm 1 below shows the two basic steps in the sequential approaches. Here we use \mathbf{x}_s^* and \mathbf{u}_s^* to designate the optima obtained from the sequential approach. We reiterate that \mathbf{x}_s^* and \mathbf{u}_s^* do not in general coincide with the optimum of (2.4).

Algorithm 1 Sequential approach for well placement and control optimization

Require: initial locations \mathbf{x}_0 and specified control strategy (fixed or reactive) \mathbf{u}_0
for $N_p + N_i$ wells

Ensure: improved locations \mathbf{x}_s^* and control strategy \mathbf{u}_s^*

- 1: Solve $\mathbf{x}_s^* = \operatorname{argmin}_{\mathbf{x} \in X} -\operatorname{NPV}(\mathbf{x}, \mathbf{u}_0)$ using pattern search optimizer
 - 2: Solve $\mathbf{u}_s^* = \operatorname{argmin}_{\mathbf{u} \in U} -\operatorname{NPV}(\mathbf{x}_s^*, \mathbf{u})$ using gradient-based optimizer
-

2.3.4 Joint approach for well placement and control optimization

We address the joint well placement and control problem using the following nested optimization

$$\min_{\mathbf{x} \in X} \min_{\mathbf{u} \in U} -\operatorname{NPV}(\mathbf{x}, \mathbf{u}) . \quad (2.9)$$

In this bound-constrained optimization problem, it is relatively simple to see that the formulations in (2.4) and (2.9) are equivalent regarding the first-order optimality conditions. In accordance with the methods presented in the previous section, the outer well placement optimization in (2.9) is solved here by means of pattern search optimization algorithms, while the inner control optimization is addressed through a sequential quadratic programming implementation with gradients computed efficiently using an adjoint-based scheme.

The approach in (2.9) may seem impractical since it requires solving a complete optimization for every cost function evaluation of the outer (upper-level) optimization problem. However, in our application a nested procedure is reasonable because of the following two observations. First, as discussed in detail earlier, the two optimizations are of different character and it is reasonable to address them using different procedures. And second, in the control optimization we will make use of a very efficient adjoint-based gradient computation within the GPRS simulator. As noted above, the bound-constrained

well control optimization problem displays multiple local solutions, but very frequently with similar cost function values. Thus, there is little if any benefit from running this optimization from multiple starting points.

We note that the nested optimization in (2.9) could also be analyzed from a bilevel programming perspective (Dempe, 2002). However, bilevel optimization problems are often more complicated to study than the problem considered here since the two optimization levels are in general associated with different objective functions.

Using the formulation in (2.9) and the specific choice of methods for the two components of the optimization, our intent is to perform some amount of global exploration (via the use of large initial stencil size in the pattern search) in a space of dimension n_1 , and not for a search in a space of dimension $n_1 + n_2$. This is an important aspect of our procedure since the computational cost associated with the global exploration of a space of dimension $n = n_1 + n_2$ grows exponentially with n (curse of dimensionality). Moreover, as we will see in Section 2.4.1, the function optimized in the outer optimization in (2.9)

$$- \text{NPV}^*(\mathbf{x}) = \min_{\mathbf{u} \in U} -\text{NPV}(\mathbf{x}, \mathbf{u}) , \quad (2.10)$$

is much smoother in \mathbf{x} , and as a consequence, easier to explore globally, than $\text{NPV}(\mathbf{x}, \mathbf{u}_0)$, with \mathbf{u}_0 being a fixed control strategy. The smoothing of the optimization surface with respect to the well placement variable \mathbf{x} occurs because the performance of wells in less promising locations can be improved, sometimes significantly, by optimizing the well controls. The function $\text{NPV}^*(\mathbf{x})$ in (2.10) is well defined since there exists an $\text{NPV}^*(\mathbf{x})$ for every feasible \mathbf{x} . We do not, however, expect there to be a unique \mathbf{u} associated with $\text{NPV}^*(\mathbf{x})$.

Hence the joint optimization approach proposed in this work can be interpreted as a well placement problem where the cost function is an optimized NPV

$$\min_{\mathbf{x} \in X} -\text{NPV}^*(\mathbf{x}) , \quad (2.11)$$

with NPV^* as defined in (2.10).

The well control optimization required for each computation of $\text{NPV}^*(\mathbf{x})$ is not solved completely in our implementation. This is motivated by the difficulty of obtaining robust stopping criteria in practical optimization problems, and by the fact that an unnecessarily tight stopping criterion may result in an excessive number of cost function evaluations. In a preliminary study involving a problem of similar complexity to those studied in this work (in terms of the well control optimization), we determined that a moderate number of iterations for the gradient-based optimizer yields an acceptable approximation of the optimal control strategy. Thus, during the course of the joint optimization, we typically use eight iterations for the well control problems (we also consider the use of four major iterations). Then, once the optimal well locations are determined, we again run the control solution but this time with a tighter stopping criterion, which leads to a slightly improved NPV^* .

It should be stressed that the optimization in (2.11) is fully parallelizable, and indeed in this work we take full advantage of this. However, the parallel runs involve control optimizations and not simply single simulations. Therefore, the computational load in each of the nodes can be very different, because in general, two calls to NPV^* (with

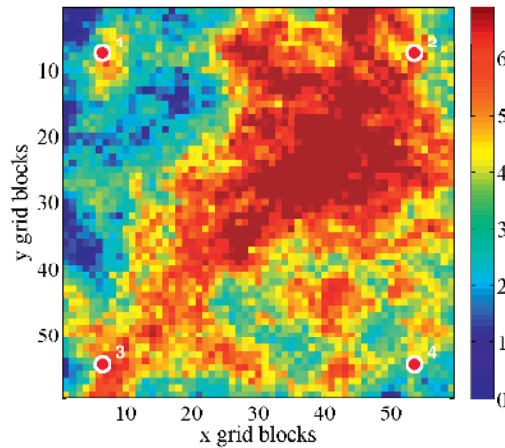


Figure 2.1: Permeability field (mD) used for the two cases in Section 2.4 (logarithm of permeability is displayed). Geological heterogeneity is clearly evident. Production wells corresponding to the first example are represented as red circles.

different well placements) will not require the same number of simulations (even using the same number of iterations in the gradient-based optimizer). This issue can be alleviated to some extent by means of asynchronous distributed computing approaches (see Griffin and Kolda, 2007, or Griffin et al., 2008, for an example within the context of pattern search).

In the remainder of the paper, the sequential optimization methodologies with fixed and reactive control strategies, and the joint technique, are denoted as *sequential fixed*, *sequential reactive* and *joint* approaches, respectively.

2.4 Example cases

In this section we apply the methodologies described in Section 2.3 to two examples. As indicated above, each control optimization problem is solved by means of a gradient-based optimizer, and the well placement part of the optimizations is handled using three different pattern search algorithms, namely, Hooke–Jeeves direct search (HJDS), generalized pattern search (GPS), and the hybrid optimization parallel search package (HOPSPACK). GPS and HOPSPACK were implemented within a distributed computing framework consisting of eight and $20 + 1$ computing cores, respectively (in HOPSPACK one of the cores is used for coordination tasks). We reiterate that, in the parallel implementations, each processor handles the full well control optimization, not just a single simulation run.

The two cases considered are based on a reservoir discretized on a 60×60 two-dimensional grid. The permeability and porosity fields are portions of layer 21 of the SPE 10 model (Christie and Blunt, 2001). These fields display strong variability in prop-

Table 2.1: Model and optimization parameters for the two examples.

Parameter	Example 1 (Section 2.4.1)	Example 2 (Section 2.4.2)
Cell size	130 ft \times 130 ft \times 20 ft	50 ft \times 50 ft \times 50 ft
Production time frame	2190 days	2920 days
Oil price (p_o)	80 \$/bbl	80 \$/bbl
Water production cost (c_{wp})	10 \$/bbl	20 \$/bbl
Water injection cost (c_{wi})	10 \$/bbl	20 \$/bbl
Injector BHP upper and lower bounds	5200 and 4100 psia	6000 and 4100 psia
Producer BHP upper and lower bounds	3500 and 1000 psia	3500 and 1000 psia

erties, as can be seen for permeability in Figure 2.1. In both examples there are five wells (one injector and four producers in the first case, and two injectors and three producers in the second case). The controls for all five wells are optimized in the two cases. In the first example, only one well (the injector) location is optimized, while in the second case all well locations are optimized. The key model and optimization parameters for both cases are shown in Table 2.1.

The gradient-based optimizer used for well control optimization is SNOPT (Gill et al., 2005), which is based on sequential quadratic programming. The initial guess in all situations is obtained by setting the injector and producer BHPs at their upper and lower bounds, respectively. This configuration provides maximum flow rates. Since the bound-constrained control optimization problem displays multiple optima, but quite often with similar cost function values, the selection of the starting point is not expected to impact the quality of the optimized solution. The stopping criteria selected for the control optimization are based on the major optimality tolerance (a value of 10^{-6} in all situations) and on the maximum number of major iterations allowed. We note that during a major iteration in SNOPT, a quadratic programming subproblem is solved to find a search direction that is used to compute the next sequential quadratic programming iterate. The solution of the quadratic programming subproblem usually requires several cost function evaluations (reservoir simulations). For more details on these stopping criteria, see Gill et al. (2007). In most cases it is the maximum number of major iterations that terminates the optimization.

As explained in Section 2.3.4, the control optimization required for finding NPV*, when called from the outer well placement loop, is not solved to full accuracy. The maximum number of major iterations is equal to eight (a relatively small number) in most cases, though in Section 2.4.2 results are also presented using a value of four. The control optimizations performed at the last iteration of both the sequential and joint approaches aim at a more precise solution. For these optimizations the maximum number of major

iterations is increased to 16 and 32 for the first and second example, respectively.

The optimizations for the well placement problem are expected to depend on the initial guess, since in general these problems are markedly nonconvex (unlike the control optimization problem when only bound constraints are present). For this reason, we perform optimization runs starting from different points. The initial stencil size in all cases is 16, which is a significant fraction of the feasible search space (41×41 and 50×50 grid blocks for the first and second examples, respectively). A stencil of this size thus leads to some amount of global exploration since regions far from the stencil center are evaluated. All pattern search algorithms terminate the optimization when the stencil size is equal to one, and the cost function value corresponding to the stencil center is lower than the cost function value associated with any other stencil point. This termination condition defines the notion of local optimality that will be considered for the discrete variables.

2.4.1 Optimization of injector location and control of five wells

Case description

In this case we consider four producers, fixed at the corners of a square, along with one injector. The four producers (designated by red circles) are located as shown in Figure 2.1. The production wells are placed somewhat away from the reservoir boundaries, which are prescribed to honor no-flow conditions. The injector can be positioned anywhere inside of the square (41×41 grid blocks) defined by the producers. The control strategies for all five wells and the location of the injector will be optimized. These strategies refer to a production time frame of six years, and except for reactive control, the strategies are divided into ten intervals of 219 days each (during each time interval the BHPs are held constant). Hence, for this problem, $n_1 = 2$ and $n_2 = 50$.

In the next section we will perform an approximation of the exhaustive search of the optimization spaces corresponding to the sequential and joint approaches described in Sections 2.3.3 and 2.3.4, respectively. Thereafter, we will use this example to compare some of the optimization techniques discussed above.

Exhaustive search results

Due to the low value of n_1 in this example, it is feasible to exhaustively explore the discrete space X both for $\text{NPV}(\mathbf{x}, \mathbf{u}_0)$, with \mathbf{u}_0 being a fixed or a reactive control strategy, and for $\text{NPV}^*(\mathbf{x})$ (where the controls are determined from optimization). This exhaustive search requires $41 \times 41 = 1681$ simulations for the sequential cases, and 1681 control optimizations for the joint approach. We note that, since the cost function appears to be close to convex in \mathbf{u} , we expect the exhaustive exploration of $\text{NPV}^*(\mathbf{x})$ to be a reasonable approximation of a global exhaustive search for the complete optimization space in (2.4). This type of exhaustive search is already impractical for the example in Section 2.4.2, where $n_1 = 10$.

In Figure 2.2 we present results for the three exhaustive explorations corresponding to a fixed control strategy (injectors at maximum BHP, producers at minimum BHP), the reactive control strategy, and the optimized control strategy. It is clear that the surface associated with the fixed control strategy is much rougher than the surfaces obtained with

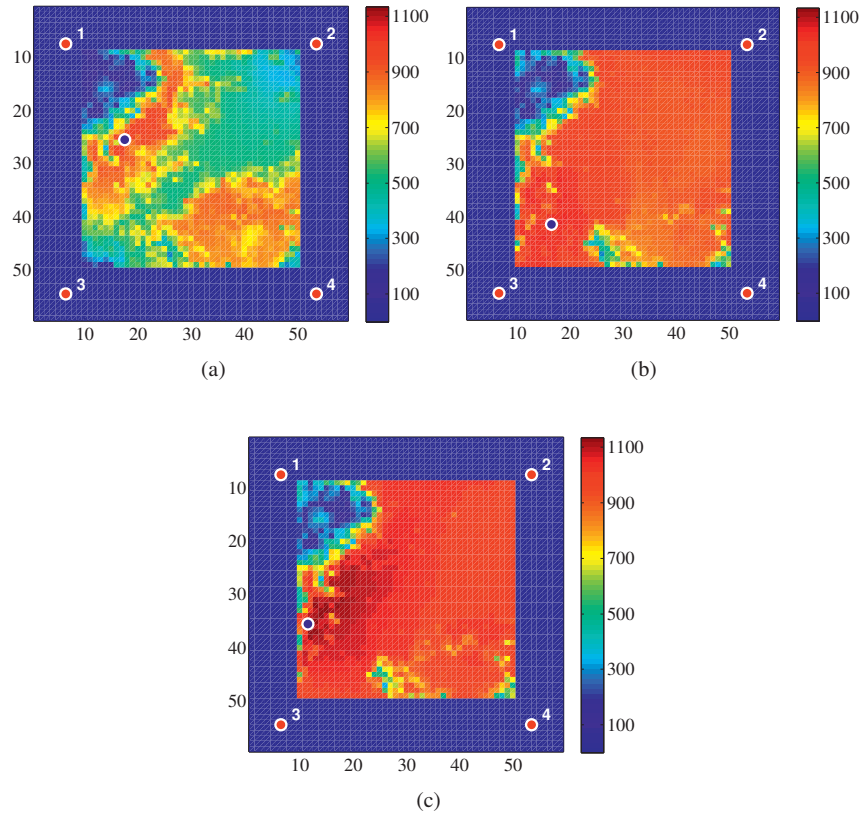


Figure 2.2: Exhaustive search results for (a) $\text{NPV}(\mathbf{x}, \mathbf{u}_0)$ with \mathbf{u}_0 a fixed control strategy (with BHPs set to provide maximum flow rates), (b) $\text{NPV}(\mathbf{x}, \mathbf{u}_0)$ with \mathbf{u}_0 a reactive control strategy, and (c) $\text{NPV}^*(\mathbf{x})$. Production and injection wells are represented as red and blue circles, respectively. The dark blue region near the boundaries is infeasible. The scale indicates 10^6 \$.

the other strategies. This demonstrates that it is possible to somewhat compensate for less promising well locations with a proper control strategy (in terms of net present value). As a consequence, the associated optimization landscape $\text{NPV}^*(\mathbf{x})$ can be expected to be smoother than the landscape corresponding to $\text{NPV}(\mathbf{x}, \mathbf{u}_0)$, for \mathbf{u}_0 a fixed strategy. This suggests, consistent with our earlier discussion, that the joint optimization landscape may be somewhat easier to explore globally.

The well locations with the highest net present value resulting from the three exhaustive explorations are given in Table 2.2. The “fixed” and “reactive” results are for the best well locations in Figure 2.2(a) and 2.2(b). The “sequential fixed*” and “sequential reactive*” results additionally apply gradient-based optimization for the well controls using the positions found in the exhaustive search. This optimization is performed with a tight tolerance (a maximum number of major iterations of 16), which is why we include the * designation. For the joint optimization, in the exhaustive search we use

Table 2.2: Injector well location and NPV for the best solution obtained for the exhaustive explorations. The * indicates that an additional gradient-based control optimization is performed.

Approach	Location [x,y]	NPV [10^6 \$]
fixed		976
sequential fixed*	[18,26]	1091
reactive		1061
sequential reactive*	[17,42]	1074
joint		1135
joint*	[12,36]	1137

Table 2.3: Average NPV (over 12 runs, expressed in 10^6 \$) for the optimal location of one injector and control of five wells. For the sequential approaches, n_{ps} is equivalent to the average number of reservoir simulations needed in the entire optimization process. For the joint approach, n_{ps} indicates the average number of control optimizations required in the complete search.

Approach	HJDS			GPS			HOPSPACK			Exhaustive NPV
	NPV	σ	n_{ps}	NPV	σ	n_{ps}	NPV	σ	n_{ps}	
fixed	901		42	883		29	891		23	976
sequential fixed*	1015	60		994	50		1002	45		1091
reactive	1003		33	1015		25	1004		21	1061
sequential reactive*	1034	44		1053	44		1044	39		1074
joint	1117		47	1109		32	1093		25	1135
joint*	1118	32		1110	33		1094	42		1137

a maximum number of major iterations of 8 (these results are designated “joint” in the table). Using the best well location found during the exhaustive search, we again run the control optimization, this time using 16 major iterations. These results are designated “joint*”.

As expected, the joint scheme out-performs the sequential methodologies, even after the additional control optimization step. The joint approach yields an increase of 4.2% and 5.9% in NPV with respect to the sequential fixed and reactive approaches. In this simple example, these improvements correspond to \$46 million and \$63 million. These amounts, as will be seen in the next example, can be even greater in larger and more realistic problems. It is interesting to note that while the reactive approach obtains a better solution than the fixed scheme before the final control optimization, the situation changes

after the control optimization. This reiterates that the control optimization can somewhat compensate for well locations that are suboptimal in terms of NPV. Furthermore, the sequential reactive approach may in some cases serve as a good approximation of the joint approach.




Figure 2.3: Cumulative production and injection profiles for the well location and controls (after the additional control optimization) corresponding to the highest NPV solution: (a) sequential fixed, (b) sequential reactive, and (c) joint approach.

In Figure 2.3 we show, for the three exhaustive explorations performed (plus the additional well control optimization), the cumulative injection and production profiles for the configurations with the highest NPV. From these plots, it is evident that the joint optimization provides more cumulative oil than the other two procedures. The joint optimization scenario also involves more water injection than the other scenarios, but this is more than compensated for by the increase (of about 5.1%) in cumulative oil. Figure 2.4 displays the oil saturation distributions at the end of the production time frame for the three optimizations. These plots illustrate how the different approaches perform in terms of reservoir




Figure 2.4: Oil saturation distribution (blue indicates water and red indicates oil) at the end of the production time frame for the well location and controls (after the additional control optimization) corresponding to the highest NPV solution: (a) sequential fixed, (b) sequential reactive, and (c) joint approach. Injection and production wells are represented as blue and red circles, respectively.

“sweep” efficiency. It is evident that there is less bypassed oil in the joint approach than in the sequential approaches.

The well controls (BHPs) corresponding to the highest NPV solutions are shown in Figure 2.5. The BHPs for the injectors (blue lines) for the various optimizations are in the top row and the next four rows (red lines) represent the producers. Upper and lower BHP bounds are indicated by dashed lines. The time axes span the entire production period (2190 days). Note that the BHPs for Producers 1 and 4 stay at the minimum BHP limit in all cases, presumably because these wells are outside of the large (diagonally-oriented) high-permeability region evident in Figure 2.1. The BHPs for Producers 2 and 3 are, by contrast, away from the lower BHP limit, for at least some part of the simulation, for all three optimization schemes. This is likely due to the fact that these wells, along with the injector, fall within the high-permeability region. If these two wells produced at their




Figure 2.5: Injection and production well controls (BHPs) corresponding to the highest NPV solution: sequential fixed (left), sequential reactive (center), and joint approach (right). Top graph corresponds to injector and next four graphs to the producers.

lower limits for the full simulation, significant water production would result. In order to avoid this, the optimizations reduce the flow rates (which leads to water breakthrough at later times) for these two wells.

Optimization solutions

The results in Table 2.2 required an exhaustive search, which is not feasible in practical situations. In this section, rather than search exhaustively, we apply pattern search optimization for the well location part of the problem. The control optimization is again handled via gradient-based optimization, with all derivatives computed efficiently using adjoint-based procedures.

We reiterate that most derivative-free optimization techniques (such as pattern search algorithms) can be readily applied to problems with discrete optimization variables, and that these methodologies have been observed to perform satisfactorily on relatively non-smooth cost functions such as that in Figure 2.2(a). Although the sequential reactive and joint strategies displayed relatively smooth cost functions (see Figures 2.2(b) and (c)), the degree of smoothness observed for high-dimensional searches may differ from that for these low-dimensional ($n_1 = 2$) cases. In any event, as we will see below, all of the derivative-free algorithms considered yield solutions that are on average relatively close, in terms of NPV, to the results from the exhaustive explorations.

As mentioned earlier, the pattern search algorithms considered here are Hooke–Jeeves direct search (HJDS), generalized pattern search (GPS) and the hybrid optimization parallel search package (HOPSPACK). These algorithms (all of which are supported by local convergence theory; see e.g., Torczon, 1997) rely on the same principles, and this facilitates meaningful comparisons. Pattern search optimization is based on evaluating a stencil whose size decreases along iterations (the reduction in the stencil size is performed when all the stencil points have a higher cost function than the stencil center). The stencil used in all cases here has $2n_1$ points distributed along the coordinate axes from the stencil cen-

ter (as in a compass). The initial stencil size is always equal to 16, and this value allows a rough exploration of the search space (for any initial guess) since the lower and upper bounds for \mathbf{x} are 10 and 50 for this example, and 6 and 55 for the second example. The sequence of stencil sizes $\{16, 8, 4, 2, 1\}$ is consistent with the optimization variables being discrete. Hence, all the algorithms stop when the stencil size is equal to 1, and the stencil center cost function value improves on every stencil point. Upon termination, the solution obtained is a (discrete) local optimizer for the $2n_1$ -point (compass) neighborhood.

Hooke–Jeeves direct search does not compute the cost function for all $2n_1$ stencil points. As soon as a point in the current stencil improves on the cost function value for the stencil center point, the stencil is moved to a new center (this strategy is known as opportunistic polling). This makes HJDS a serial strategy that can be attractive when distributed computational resources are scarce, or when commercial software licensing issues limit massive parallelization. Since both GPS and HOPSPACK evaluate the $2n_1$ points for every stencil, the use of distributing computing is very beneficial for these algorithms.

In this example the three cost functions, $\text{NPV}^*(\mathbf{x})$ and $\text{NPV}(\mathbf{x}, \mathbf{u}_0)$, with \mathbf{u}_0 corresponding to all wells at their BHP limits, and to a reactive strategy, are based on a lookup table constructed with the results from the exhaustive explorations. Therefore, for this case, GPS and HOPSPACK do not take real advantage of being implemented in parallel. In the example in Section 2.4.2, this feature will be effectively exploited.

Separately, and this is applicable to pattern search methods in general where the stencil only changes its size along iterations, some points in the optimization are revisited at different times. The cost function computation in these cases can be avoided if all (or just a number of) evaluations are stored in a cache. In this work caches are implemented for the three pattern search algorithms considered.

The results from the three approaches, together with the NPVs obtained in the exhaustive explorations, are summarized in Table 2.3. The NPVs for the exhaustive explorations are taken from Table 2.2. Because different initial guesses result in different locally optimal solutions, we run each optimization 12 times, starting at different initial points. Each pattern search run is followed by a gradient-based control optimization with tight tolerances (as above, * denotes the use of a maximum of 16 major iterations). The NPVs, expressed in 10^6 \$, are averaged over the 12 runs. Standard deviation of the NPV (σ) over the 12 runs is also reported for each case. The average number of iterations n_{ps} for each pattern search procedure is also reported. It is important to note that for the sequential fixed and reactive approaches this number is equivalent to the average number of reservoir simulations needed for the entire optimization process. However, for the joint approach it indicates the (average) number of control optimizations required in the complete search. In this example, each control optimization requires on average 14 reservoir simulations.

The differences (in terms of NPV) between the results obtained by the sequential and joint approaches before performing the additional control optimization are somewhat larger than the corresponding results for the exhaustive explorations (shown in Table 2.2). This may be because the cost function for the joint approach is globally smoother, which makes it easier to optimize. The additional control optimization to some extent reduces the discrepancies in the results. Before the control optimization step, the average opti-

Table 2.4: Results for NPV (expressed in 10^6 \$) and total number of simulations n_{sim} for the optimal location and control of five wells. GPS is used for the well location optimization. The highest NPV for the nine runs for each approach is underlined.

Run #	seq. fixed* _{GPS}		seq. reactive* _{GPS}		joint* _{4,GPS}		joint* _{8,GPS}	
	NPV	n_{sim}	NPV	n_{sim}	NPV	n_{sim}	NPV	n_{sim}
1	<u>336.6</u>	321	334.5	295	347.4	1075	<u>385.1</u>	2683
2	300.5	505	354.8	422	353.0	2021	355.2	3276
3	328.9	426	314.7	310	329.1	1770	346.7	3327
4	328.0	511	192.7	240	325.8	1922	372.3	2481
5	326.7	477	240.9	377	<u>355.2</u>	1936	354.8	5003
6	294.9	468	253.3	361	336.6	2741	336.0	3278
7	263.4	423	345.4	329	344.7	2031	360.0	3941
8	256.8	644	279.6	420	339.5	2602	357.2	4187
9	293.7	587	<u>358.8</u>	447	330.6	1938	358.0	4855
Mean	303.3	485	297.2	356	340.2	2004	358.4	3670
σ	29.2	94	58.6	68	10.6	479	14.0	890

mized NPV by the joint approach is 24.1% and 9.8% larger than the average optimized NPVs from the fixed and reactive approaches, respectively. After the additional optimization, these percentages decrease to 10.3% and 6.1%.

It is not clear from the results in Table 2.3 if one pattern search algorithm is preferable over the other two. GPS and HOPSPACK are slightly faster than HJDS, but they yield lower average cost function values. As noted earlier, however, GPS and HOPSPACK can be accelerated, in terms of clock time, if a cluster is available (and in this situation, they will outperform HJDS). In this relatively simple case ($n_1 = 2$, and cost function computed via a lookup table) the performance of GPS and HOPSPACK seems to be comparable. Differences between the various approaches and algorithms will be more evident in the next example, which is more realistic and more complex.

2.4.2 Optimal location and control of five wells

Case description

In this example we optimize both the location and control of two injectors and three producers. Some of the reservoir parameters are different than those used in Section 2.4.1. Specifically, the reservoir area is reduced, the production time frame is longer, and the costs for injected and produced water are doubled (the corresponding model and optimization parameters are given in Table 2.1). A water-flooding configuration with two

injection wells is richer in terms of variety of sweeping strategies than an arrangement with only one injector (as considered previously). Thus we aim at increasing the diversity of production scenarios, which renders the search more challenging. In addition, our reactive control strategy (which addresses only producers) is now less appealing because water injection is costly.

The number of grid blocks in the reservoir model, and the permeability and porosity values for each grid block, are the same as before (see Figure 2.1). Because all five well locations are optimized we now have $n_1 = 10$. The production time frame is again divided into ten intervals (hence, $n_2 = 50$). As in the previous example, the gradient-based optimization algorithm embedded in the joint approach is SNOPT, and the pattern search methods considered for the well placement search are HJDS, GPS, and HOPSPACK. Both GPS and HOPSPACK are implemented within a distributed computing framework. While HOPSPACK uses 21 cores (one core is dedicated to the coordination of the concurrent jobs, leaving effectively 20 computing cores), GPS, due to a limited number of licenses, is applied only on eight cores. The markedly nonconvex character of the well placement optimization is dealt with by running the optimizations nine times with different initial guesses. These initial guesses were not randomly selected – rather, they correspond to well placements that are reasonable from a reservoir engineering perspective.

Optimization solutions

In these optimization runs, the parameters for GPS are the same as in Section 2.4.1, i.e., the sequence of stencil sizes is $\{16, 8, 4, 2, 1\}$. The control optimization in the joint approach is solved with two different values (four and eight) for the maximum number of major iterations. In all cases, one supplementary control optimization is performed with a maximum number of major iterations of 32 (one optimization for the sequential fixed approach needed 64 iterations because convergence was not obtained after 32 iterations).

The results using GPS in the well placement optimization part for all of the approaches and each of the nine different well location initial guesses are presented in Table 2.4 ($\text{joint}_{4,\text{GPS}}^*$ and $\text{joint}_{8,\text{GPS}}^*$ refer to the joint approach with the maximum number of major iterations in the gradient-based control optimization equal to four and eight, respectively). The total number of simulations performed in each of the runs is denoted by n_{sim} . The final control optimization is performed in all cases, and the simulations required for this step are included in n_{sim} .

We observe that the average (maximum) NPV for $\text{joint}_{8,\text{GPS}}^*$ over the nine runs is 5.3% (8.4%) higher than for $\text{joint}_{4,\text{GPS}}^*$. This observation is consistent with the much larger number of simulations performed in $\text{joint}_{8,\text{GPS}}^*$, and indicates that a maximum number of major iterations of four tends to terminate the optimization prematurely. An insufficient maximum number of major iterations may yield a clearly suboptimal solution, and an excessively large value could lead to prohibitive computational requirements. Therefore, a tuning process for this parameter might be beneficial when applying the joint approach. For the remainder of this section, our remarks on the joint approach will refer to the case with maximum number of major iterations equal to eight.

In terms of NPV, the sequential fixed and reactive strategies clearly under-perform the joint approach. The average (maximum) NPV over all of the runs obtained with the joint

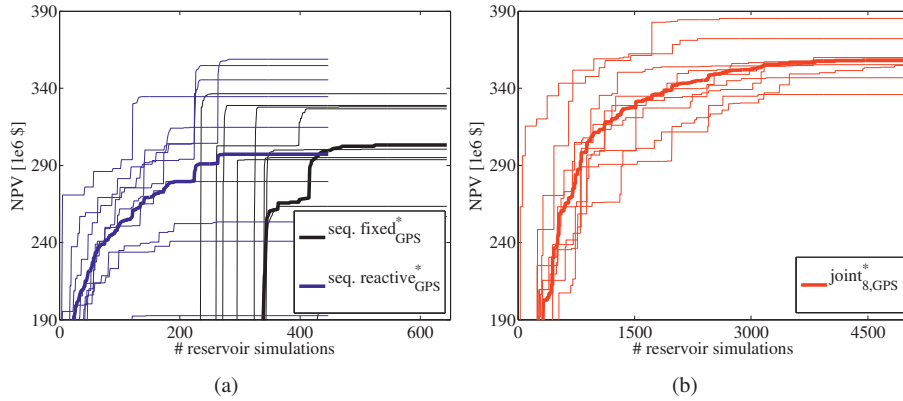


Figure 2.6: Evolution of the objective function (NPV) for all nine runs versus number of simulations: (a) sequential fixed and reactive approaches, (b) joint approach. Corresponding averages over the nine runs for each approach are represented by thick solid lines. All runs include the supplementary control optimization.

approach is 18.2% (14.4%) and 20.6% (7.3%) higher than with the sequential fixed and reactive schemes, respectively. The average number of simulations required by $\text{joint}_{8,\text{GPS}}^*$ is, however, about one order of magnitude higher than that needed by the sequential methodologies. Along these lines, it is important to realize that the maximum NPVs reported in Table 2.4 for the sequential strategies are based on a fraction of the computational effort dedicated to the joint approach. Thus, in order to complement the results in the table, we tested the sequential reactive scheme (with supplementary control optimization) for 100 new random initial well locations (in that manner, the associated total computational cost is comparable to that for $\text{joint}_{8,\text{GPS}}^*$). The average and maximum NPV over these 100 runs are \$288.0 million and \$353.8 million, respectively. These values are lower than the corresponding values in Table 2.4 (\$297.2 million and \$358.8 million), which is in accordance with the fact that the nine initial well placements were not selected randomly but rather based on engineering judgement. The key observation, however, is that, even when we compare based on the same number of total simulation runs, $\text{joint}_{8,\text{GPS}}^*$ still outperforms the sequential reactive scheme.

As can be seen in Table 2.4, an advantage of the joint approach is that it results in smaller standard deviation σ of the NPV than the sequential methodologies. This fact is consistent with the smoothing of the well placement optimization landscape observed for the joint strategy (which was illustrated earlier in Figure 2.2). Note further that the results for the sequential reactive approach are not in this case as close to those for the joint strategy as in the previous example (indeed, here they are more comparable to those for the sequential fixed approach). This may be explained by the increased complexity of this problem and by the elevated cost of injected water.

The optimization results for the sequential and joint approaches are further illustrated




Figure 2.7: Oil saturation distribution at the end of the production time frame for the well controls and locations corresponding to the run from Table 2.4 with maximum NPV: (a) sequential fixed, (b) sequential reactive, and (c) joint approach. Injection and production wells are represented as blue and red circles, respectively.

in Figures 2.6(a) and (b), where the evolution of the objective function (NPV) versus the number of forward simulations is represented for each of the runs. The corresponding averages over the nine runs are plotted as thick solid lines. In order to enable clear comparisons, all figures use the same vertical scale. We note that, prior to the supplementary (final) control optimization, all solutions for the sequential fixed scheme have NPVs lower than \$200 million. From Figure 2.6(a) it is clear that the additional control optimization is crucial in the sequential approaches. We reiterate that in both the fixed and reactive strategies the water injectors operate at maximum BHP, and this may negatively impact the objective function. Hence, the supplementary control optimization can again be seen as a means to compensate for suboptimal well locations. The lower standard deviation in the joint approach compared to the sequential strategies is also evident in Figure 2.6.

The oil saturation distributions corresponding to the solutions with maximum NPV, at

the end of the simulation time frame, are presented in Figure 2.7 (injection and production wells appear as blue and red circles, respectively). The amount of bypassed oil is noticeably less for the joint approach than for the sequential strategies. The well locations obtained generally tend to be toward the boundaries of the domain. In a few cases, some wells are placed very close to each other, as can be seen e.g., in Figure 2.7(b) for the sequential reactive approach. This type of solution might not be acceptable in practice, and can be prevented in the optimization by including (non-linear) constraints that ensure a minimum distance between wells. The computation of these constraints does not involve time-consuming function evaluations, and for that reason, they are not as complicated to handle as other simulation-based constraints that may be present.

The results obtained for HJDS, GPS and HOPSPACK are shown in Table 2.5 for the same nine initial well locations considered in Table 2.4. The settings and stopping criteria for these derivative-free optimizers are the same as were used for GPS. In all cases the maximum number of major iterations in the gradient-based control optimization is equal to eight (and again all runs include an additional control optimization with a maximum number of major iterations of 32). The total number of control optimizations solved (n_{ps}) coincides with the number of times the function NPV* is called within each pattern search algorithm, and can also be assumed to be roughly proportional to the total computing cost. Each call to NPV* involves approximately 12-15 forward simulations, so the total number of simulations is around 4000 for the different approaches (consistent with the 3670 value given in Table 2.4). It is useful to express the results in terms of n_{ps} (rather than in terms of the total number of simulations n_{sim}) when the optimizations are parallelized, and because all the algorithms compared are pattern search algorithms of the same type.




Figure 2.8: Evolution of NPV averaged over nine runs versus the equivalent number of control optimizations for the three pattern search optimization algorithms considered. The number of equivalent control optimizations solved is the total number of optimizations divided by an estimate of the speedup. The speedup factors estimated for GPS and HOPSPACK are 4.1 and 6.4, respectively.

Table 2.5: Results for NPV (expressed in 10^6 \$) and total number of control optimizations solved (n_{ps}) for the second example using Hooke–Jeeves direct search (HJDS), generalized pattern search (GPS) and a hybrid optimization parallel search package (HOPSPACK). The highest NPV for the nine runs for each approach is underlined.

Run #	$\text{joint}_{\text{HJDS}}^*$		$\text{joint}_{\text{GPS}}^*$		$\text{joint}_{\text{HOPS}}^*$	
	NPV	n_{ps}	NPV	n_{ps}	NPV	n_{ps}
1	<u>386.6</u>	331	<u>385.1</u>	188	386.6	222
2	380.8	391	355.2	225	<u>388.0</u>	333
3	327.6	216	346.7	225	358.5	375
4	386.4	316	372.3	175	380.0	358
5	377.5	321	354.8	331	343.1	361
6	344.9	456	336.0	244	333.9	286
7	377.9	556	360.0	285	353.9	331
8	313.8	441	357.2	311	350.6	344
9	371.9	306	358.0	332	358.0	468
Mean	363.0	370	358.4	257	361.4	342
σ	27.3	101	14.0	60	19.3	66

In the absence of distributed computing resources, HJDS performs marginally better than GPS and HOPSPACK. However, these two algorithms are preferable to HJDS once they are implemented in parallel. The effect of distributed computing on GPS and HOPSPACK is shown in Figure 2.8. In that figure the vertical axis represents the evolution of NPV averaged over all nine runs, and the horizontal axis corresponds to the equivalent number of control optimizations solved for each pattern search optimization algorithm. The number of equivalent control optimizations solved is defined as the total number of optimizations divided by an estimate of the speedup obtained through parallelization. We note that HJDS is inherently serial, and for that reason the number of equivalent optimizations coincides with the total number of optimizations solved. For all algorithms, the horizontal axis in Figure 2.8 is roughly proportional to total clock time. Though GPS and HOPSPACK are parallelized on eight and 20 computing cores, respectively, the speedup factors estimated for these procedures are 4.1 and 6.4, respectively. Consequently, as can be seen in Figure 2.8, HOPSPACK outperforms GPS in terms of total elapsed time.

It is worth noting that the ratio of the two speedup factors is different than the ratio of the numbers of computing cores used for the two algorithms. This discrepancy is related to an observed increase in the reservoir flow simulation clock time with the number of nodes used. This may be explained by the concurrent sharing of common libraries by the parallelized simulations or by excessive input/output data traffic within the cluster. Thus,

in some practical applications there might be an optimal number of nodes to use in a distributed computing framework (in other words, a larger number of nodes does not always provide a higher speedup factor). In this example, for HOPSPACK, we do not observe clear effects associated with the asynchronous parallelization mode, but it is reasonable to expect an increase in performance for more complicated cases. From Figure 2.8 it can also be concluded that HJDS could be an alternative to the other two derivative-free methods if distributed computing resources are limited or unavailable, particularly if the optimization algorithm need not be run to full convergence.

2.5 Concluding remarks

In this work we considered the joint optimization of oil well placement and well controls. These two problems, though clearly coupled, have been treated as separate optimizations in most previous studies. We devised a nested optimization approach where the outer (high-level) optimization addresses the well placement problem. For each well configuration, the optimization cost function is defined as the optimal objective function value after performing a well control optimization for the particular well arrangement. Since well control optimization often displays a more convex character than well placement optimization, the former optimization can be approached from a more local, and thus efficient, perspective than the latter optimization. Therefore, in the well control optimization we apply a gradient-based procedure, with gradients provided by an adjoint solution. For the well placement part of the optimization, several pattern search algorithms were considered. Although these are local optimizers, by using a large initial stencil size we achieve some amount of global search.

We considered two optimization problems involving different numbers of optimization variables. Three basic optimization strategies were considered – two of these were sequential schemes that involved particular assumptions regarding the well controls (specifically, fixed and reactive controls) used during the well location optimizations, and the third was the joint optimization procedure. In all cases, after the basic optimization had converged, we performed an additional well control optimization (for the optimized well locations) using the gradient-based procedure with a large number of iterations (i.e., a tighter stopping criterion tolerance). In the first example, the location of only one well was optimized, so we were able to perform an exhaustive search. This enabled a clear assessment of the performance of the different optimization methods. The exhaustive search results showed that the optimization landscape corresponding to the well location in the joint approach was smoother (suggesting that global exploration can be more readily accomplished in this case) than the optimization landscape for the sequential fixed strategy.

The joint procedure was shown to consistently outperform the sequential schemes in terms of the optimized cost function (net present value in our examples). For the second (more challenging and more realistic) example, the increase in net present value obtained by the joint approach exceeded that achieved by the sequential methodologies by around 20% on average. The joint approach does, however, require around an order of magnitude more reservoir simulations than are required for the sequential approaches. This high computational demand can be mitigated through use of parallel implementa-

tions of the pattern search algorithms. Two of the algorithms considered, generalized pattern search (GPS) and hybrid optimization parallel search package (HOPSPACK), parallelize naturally and such implementations were in fact applied.

The control optimization in the joint approach can be interpreted as an effective means to compensate for well placements that are suboptimal from the objective function perspective. This reasoning can be extended to optimizations that include variables in addition to those considered here. For example, the negative effects resulting from using too short of a production time frame, or from an insufficient number of wells, could be alleviated to some extent by the optimization of well locations and controls. We reiterate, however, that the “smoothing” of the optimization landscape typically entails a significant increase in computational cost.

The joint optimization procedure presented here can be extended in several interesting directions. For problems involving more general (non-linear) production constraints, it is not clear if multiple optima with very similar cost function values will continue to be observed in the well control optimizations, as they are for bound-constrained problems. If this is not the case, then this issue must be addressed in some way; e.g., by performing multiple control optimizations using different initial guesses. Another useful direction for future research is to consider the use of surrogate models to accelerate the optimizations. Specifically, in some of the computations, the optimized net present value could be estimated using a sequential reactive strategy. This approach would be most effective if the particular reactive strategy is “tuned” (including some treatment for injection wells) based on the joint optimization results. Further effort should also be expended toward including inter-well distance constraints and non-linear simulation-based production constraints (such as maximum water cut in production wells), possibly through use of a filter method (see e.g., Echeverría Ciaurri et al., 2011a). It will also be useful to consider global exploration techniques such as particle swarm optimization (Eberhart et al., 2001) or genetic algorithms (Goldberg, 1989) for the well placement part of the optimization. Uncertainty in the reservoir model should also be included in the optimization using, for example, the stochastic procedure recently presented by Wang et al. (2012).

In Chapter 3 we focus on extending the functionality of the joint procedure so it can be applied to a real field case. A strategy is chosen where we emphasize the full use of the joint approach, even though this is particularly costly when dealing with a field reservoir model. An important component of this strategy is therefore to target for further development specific capabilities within the procedure, e.g., the parallelization of pattern search methods. Also, some of the developments mentioned above, e.g., inter-well distance constraints, have been introduced in the new implementation of the joint procedure, referred to in Chapter 3 as an optimization framework. Development along some of the other topics, e.g., the use of surrogate models, is currently underway, and is also briefly discussed in the next chapter. Several of the other possible developments mentioned above still remain important topics for further research, and will be the subject of future work.

Joint Optimization of Well Placement and Controls Applied to a Real Field Case

Introduction

Field case application. In Chapter 2 we introduced and tested the joint and sequential approaches for the optimization of well placement and controls. In this chapter we develop an optimization framework based on these approaches, and test it using a real field case. The field case model tested in this work is significantly more challenging than the cases used in Chapter 2 to develop the methodology. The purpose of the framework is to extend and adapt the developed methodology so that we are able to efficiently apply it to a real field case. A pilot application of the framework is performed using a real field model provided by field operator Total E&P Norge AS. This model includes four horizontal well trajectories planned for the development of a North Sea field. This application project constitutes a first attempt at testing our approaches for well placement and control optimization within a field development operational context (see Chapter 1).

The field case model tested in this work is associated with the development of the Martin Linge oil reservoir. The Martin Linge oil reservoir is part of the Martin Linge field located on the Norwegian Continental Shelf. The field production license is owned by Total E&P Norge AS together with partners Petoro AS and Statoil AS. This chapter presents the work concerning the development and implementation of the optimization framework. During optimization, the framework developed in this chapter uses a work model approximation of the field model. The next chapter (Chapter 4) deals with testing the solutions obtained in this chapter on the original field case model. This pilot application is the result of a collaborative effort between NTNU's Center for Integrated Operations in the Petroleum Industry (IO Center), and IO Center Industry Partner Total E&P Norge AS. Some of the contents in this chapter have been adapted for submission to an *SPE* publication.

Additional note: All information presented in this chapter regarding the Martin Linge field and oil reservoir has been obtained from publicly available sources, the reservoir simulation model and from meetings with the field development team. Presentation of the material has been approved by Total E&P Norge AS and partners Petoro AS and

Statoil AS. The author is solely responsible for all the presentation and analyses of the material. We are grateful to Total E&P Norge AS, and partners Petoro AS and Statoil AS, for allowing us to use the Martin Linge oil reservoir simulation model in our work.

Introduction to chapter. This chapter regards the development and implementation of the proposed optimization framework using a real field model. It contains five sections. The first two sections cover work process aspects, while the last three sections conduct the technical application. The testing of the solutions obtained from the optimization effort in this chapter is covered in Chapter 4. Together, this chapter and the next constitute the pilot application effort that is the use of our methodology on a real field case. The first two sections in this chapter serve as introduction and treat the entire application work as a whole. The first section presents the main goals and strategy for the development of the joint procedure into an optimization framework that can treat a real field case. The pilot application of this procedure is then described in terms of parts in a work process loop. The second section presents the general design of the optimization framework, and discusses some of the main developments introduced to treat the real field case. The various developments are characterized as either enhancements, extensions, additions or replacements, to the methodology presented in Chapter 2. The third section gives a general presentation of the Martin Linge field, followed by a description of the Martin Linge oil reservoir model, and extensive notes on the transfer and validation of the field model to the research reservoir simulator used in this work. The fourth section provides detailed definitions of the optimization problem, non-linear constraints for the well placement part of the optimization, as well as a presentation of the algorithmic procedure. Finally, the fifth section in this chapter, shows the results from the optimization effort. The various data obtained from the application of the optimization framework are discussed, and suggestions for further work, specially to improve the application of the procedure within a field development operations environment, are presented.

3.1 Targets and strategy for application development

This section presents two targets for application development. Together, these targets underlie all the work in this chapter. These targets are concerned with the progression of the research developed in Chapter 2, and with the application of the core methodology on a real field case. A general work strategy is outlined to achieve the stated targets. The components defined for this strategy serve as guiding principles for the development of our work, i.e., the re-implementation of our methodology and its extension to a real field case. A work process loop has been developed to structure the work performed in this application effort. This loop, and the work process definitions included in it, represent the actual execution of our strategy. Some of the challenges encountered during the course of this work are commented at corresponding stages in the work process loop.

3.1.1 Application targets and strategy

Preface for strategy terminology. We have previously developed a research methodology to treat relatively simple example cases dealing with the well placement and control

optimization problem. However, the application work described in this chapter involves a field case model and a more advanced well placement configuration. Among other things, this application work has required that we gain understanding about field development work tasks related to well placement strategy, and that we use this knowledge to improve the functionality of our application, in particular with respect to what kind of results we should produce, and how these might contribute to the work process of the operator. (By “functionality” we broadly mean what kind of problems the application is able to solve.) An important challenge has been to identify relevant information which could then be used to adapt our optimization effort, so as to best align the results from the application with the observed needs of the operator.

Moreover, applying basic methodology on a complex real field case has demanded that we amplify our research work, both technologically and in terms of work processes. With respect to work processes, some of our work effort has been spent (in a non-expert modality) on dealing with topics related to the area of technology management (e.g., technology development and integration of new technology within established workflows). Technology management is not our field of study, and we will not use specific terminology to treat the various aspects of this work where this theory might be relevant (and indeed helpful). However, we have found that some of the main principles within this area of study are useful to describe some of the broad aspects of this work. These principles can therefore be used to structure our thinking, which will in turn help us organize and guide our research and application effort. For example, a fundamental principle within the area of technology management is the need to create and effectively execute a coherent strategy (Læg Reid, 2001). In this chapter, we will use this and related concepts regarding the formulation and execution of strategy to treat (only descriptively) the developmental aspects of our work, i.e., the general extension of our research methodology for real field case application.

The main reason behind this particular setup is that we believe it is important for us to understand our work within a broader industry and research context (in this case from a very broad technology management perspective). We furthermore believe obtaining better distinctions regarding our general area of work (research development and application) will help us be clearer about decisions that have been taken during the course of this project. Moreover, we think this understanding will help us improve the planning and guidance of future work developments. We therefore adopt the terminology of strategy targets and execution to offer an overarching description of this work, and how it has been conducted¹. (Most of the theoretical foundation supporting our use of the concept of strategy and related terminology is based on our reading of M.Sc. thesis "Technology Strategy and Innovation Management in the Petroleum Industry" by Læg Reid (2001); see also note².)

Targets for application work

We identify two specific targets for the application work presented in this and the next chapter. (Notice that in this section we use the terms “methodology” and “technology” interchangeably).

First target. The first target for the work in this and the next chapter is to further test the core methodology developed in Chapter 2, i.e., the joint approach versus the sequential

strategies. At the end of Chapter 2, we mentioned several important work topics for further development. Among these, we believe testing the approaches on more sophisticated example cases pose an interesting next target for further development. A first attempt along these lines would be to target cases that include horizontal well-bores, and that involve more advanced control optimization problems, e.g., production scenarios that deal with inflow control valves and/or the enforcement of non-linear production constrains. However, as we will discuss later, this first target co-exists with a second target (stated below). This co-dependency prompts a realignment of the first target with respect to the second one. In summary, due to the complexity of the field case involved in the second target, we are required to scale down on applying a more complicated production optimization problem. Still, the original intention of advancing the developed methodology by treating more sophisticated cases is, for this application effort, amply met by the need to introduce a suite of non-linear constraints to solve for the well placement part of the field case problem.

Second target. The second target for this application work is to extend and adapt the basic methodology developed in Chapter 2 into an application that can deal with a real field case. With this we mean the developed application should be able to handle the new technological challenges that emerge when treating a field case, e.g., the increased computational demand that entails using a field case model, or close approximations of one, within an iterative optimization procedure. It also means the functionality of the research application should, to a reasonable degree, be aligned to the business needs of the end-user, e.g., field operator. The connection of this target to the first one is then clear, in that making the research application useful to the field development work process of the operator becomes the end-point for the methodology development outlined for the first target. To reach this end, an important target in itself is to establish an effective collaboration environment. Through an effective collaboration process, our aim is to gain case knowledge that will help us build problem and constraint definitions so that results from the optimization procedure are able to address work task topics that are important to the operator. The translation process is bi-directional, and we therefore emphasize using tools and software platforms commonly employed by the operator to test the solutions obtained by our optimization procedure. This also includes analyzing the results from the general perspective of the field development work process (this analysis is of course limited by our capacity to meaningfully apply this perspective given our knowledge background and experience regarding the specific field development).

Clearly, these targets are tightly coupled, in that the first target draws a line for development from a current state of functionality, while the second target represents a final state of functionality for the application. Crucially, a conflict arises due to the computational demand from each of the targets. Roughly, the former target requires optimizations involving both well placement and control variables, while the latter targets demands a high-fidelity representation of the field case reservoir model to be used during optimization. Below we develop a strategy to deal with these targets.

Strategy

A single strategy is developed to solve the coupled target system. Though we treat the targets as a coupled system, our understanding is that a single work path to reach them will confront diverging forces because the targets pull differently for the same resources. We therefore develop a strategy composed of two main components meant to work in tandem to satisfy both objectives. The first component deals mostly with design issues and expansion of research topics. The second component focuses on solving the field case problem by making use of and adapting the designed capabilities set up by the first component, and on establishing collaboration work. This latter point aims at gaining understanding of work tasks and problem features, and on translating these into practical problem and constraint definitions.

In Figure 3.1, a single horizontal line is drawn to represent the development process for the application work in this thesis. Importantly, the figure illustrates how the two strategy components ultimately help guide and define the application and testing of our methodology. The process moves from left to right, starting with basic research (Chapter 2), then going through our application effort (Chapter 3) and finally reaching an area of field case testing (Chapter 4). The two strategy components are shown as red ellipses with black arrows illustrating their basic mode of operation within the application work. The two strategy components are further described below.

First strategy component. As illustrated in Figure 3.1, the first strategy component points towards achieving a real field case application (which is the main part of our second target). However, this component is grounded on the current state of the methodology, as developed in Chapter 2, and its primary focus is on expanding the original research along the lines discussed at the end of that chapter. Basically, the first strategy component enables further development of the methodology by re-implementing the work in Chapter 2. Crucially, this re-implementation consists of re-building the joint approach as a framework to solve embedded, or integrated problems. A useful benefit of the framework design we propose here is that we can extend, add or replace components within it. In our work, we use these operations to reach the second target described above. At the same time, the chosen extensions, additions and replacements are selected based on how much they contribute to advance the original research, thus serving to satisfy the first target. In particular, the current research is moved forward by which developments that are implemented, and how they are treated, e.g., dealing with horizontal trajectories and non-linear well placement constraints still enables us to further explore the gain of the joint versus sequential approach. As mentioned, to accommodate for further developments, this strategy component puts a strong focus on building infrastructure. Concepts of design and modularity are specially important in this process. Finally, developing the joint approach into an optimization framework enables us to exploit the relationship between the parts, e.g., an interesting line of work is to attempt to accelerate the solution procedure by introducing surrogate techniques to decrease the cost of the embedded optimizations (indeed, this possible enhancement will be discussed at a later point).

Second strategy component. As illustrated in Figure 3.1, the second strategy component is about shaping the end-point problem definition. This definition is an interface between the development driven by the first strategy component, and the testing of our application

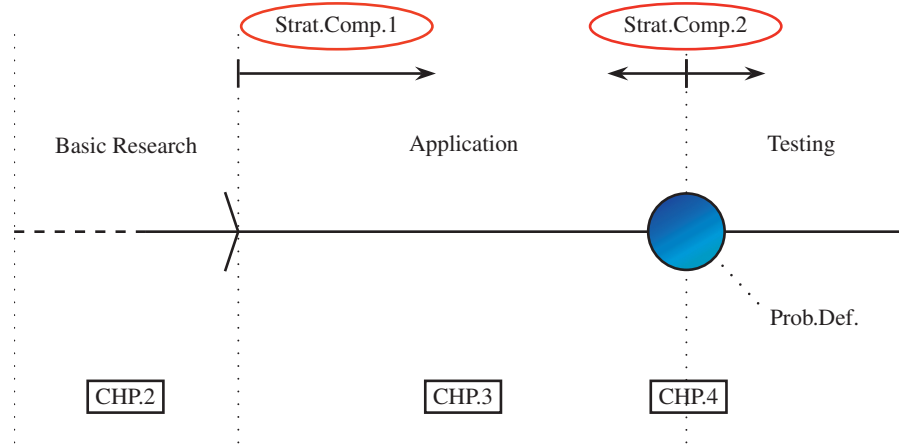


Figure 3.1: Illustration of application development. Development process is represented as moving left to right on horizontal black line. Strategy components are drawn as red ellipses with arrows showing their main mode of operation. Thesis chapters are positioned according to where they treat the different parts of the work in this thesis. “Prob.Def.” label refers to the overall problem definition for the field case application.

on a real field case. The fundamental task of the second strategy component is then to balance the two types of activity (i.e., petroleum engineering and research) that contribute to shape the definition. This mode of operation is further illustrated in Figure 3.1 by the two black arrows exerting influence towards opposite sides of the development line. The balancing task is to establish a problem definition that may possibly yield interesting results for the engineering side (preferably involving the field case model or a close approximation of it), but that can also be solved, in a relatively efficient manner, by the research application effort. (Admittedly, in this task, the main mode of operation of the first strategy component may on some occasions be a constraining element.) In total, we describe the second strategy component as both being focused on using and adapting the designed capabilities set up by the first component to solve the field case problem, and also, at the same time, as being engaged in a collaborative process aimed at shaping a reasonable and interesting problem description by gaining deep understanding about the field problem and by facilitating feedback and contributions from expert sources.

Clearly, then, an important part of the second strategy component is an efficient collaboration effort. An efficient collaboration effort, in our perspective, encourages the processing of information about the field development and oil reservoir, and facilitates the translation of this case information into problem and constraint specifications. Furthermore, one of the main tasks of the collaboration effort, as we see it within the context of application work, is to use expert understanding about the field problem to limit the scope of the problem definition. A definition with a tightly defined scope can serve as a better target for application development, and solutions obtained from solving such problem may possess a high degree of operational relevance. To this end, the operator team may provide expert knowledge and feedback to perform these limitations, while general

knowledge to further shape the problem definition is gained from our interaction with the team and through study of case documents.

As a final note, the first and second strategy components are meant to operate iteratively because, based on our experience, the definition of the end–point problem is an ongoing process that is also dependent on the development of the application functionalities, and vice versa.

Next we will discuss execution of the strategy. In the following we will describe how the stated strategy has been executed, and some of the main decisions taken during the course of the execution. In particular, we will briefly mention the various challenges contained in each decision, and how these were resolved to satisfy the two stated application targets. As might be expected, several of the more difficult challenges encountered in this application work stem from the bi–objective nature of the project.

3.1.2 Work process structure

In this section we describe the execution of the outlined strategy. We describe how we have structured our work such as to deliver the results presented at the end of this chapter and in Chapter 4. The structure of our work represents the actual execution of our strategy.

Work process structure. The background for this work is a collaborative effort between NTNU/IO Center and IO Center Industry Partner Total E&P Norge AS. Total E&P Norge AS is the Martin Linge field operator, and has provided a model of one of the Martin Linge field reservoirs for use in this project. For this work, we have developed a structure of work processes that extends our research methodology and yields an application to optimize the set of initial well locations in the test case provided by Total E&P Norge AS. The creation of this structure is the main effort to solve the second of the work targets (the one focused on end–point application), while the effort to solve the first target (dealing with research development) is mostly defined by the choices taken within this work process structure.

Motivation. The main motivation for this work has been to produce well placement results that offer direct value to the current field development work process of our Industry Partner. The overall strategy has been to expand our core methodology for well placement and control optimization, i.e., the joint and sequential approaches, into a practical real field application. The application work includes not only expansions of the methodology itself, but also work process issues such as collaboration efforts and the integration of results back into the information stream of the operator. Clearly, all these issues are interrelated. We have created a work process loop to help organize and allocate our efforts along these various types of application work.

Work process loop. To structure the work in this project, we developed a work process loop with four different parts, or stages. This work process loop is shown in Figure 3.2. Each part consists of a different type of application work and a different mode of collaboration between the research team and Industry Partner. It was important to establish a clear work process structure since both the technical and collaborative type of efforts needed at one part could be quite different from the work efforts required at another part. The work process loop ranges from model transfer and validation work to problem design, optimization effort and solution testing. Mainly, the four parts represent

clear changes in work focus at different stages of the project.

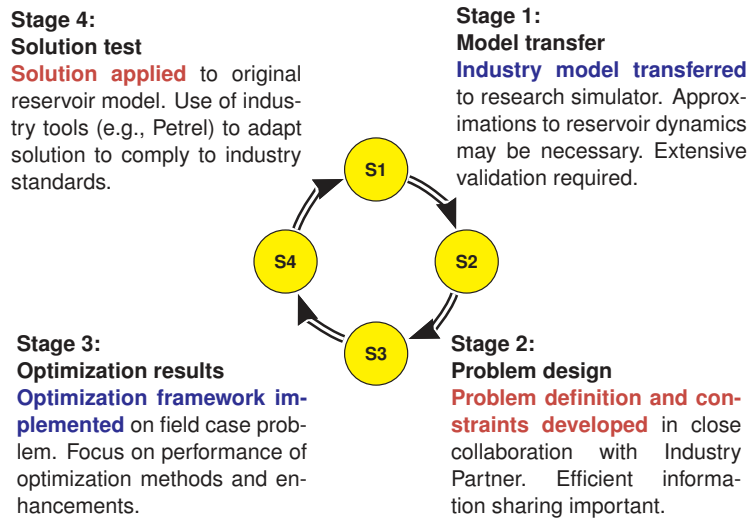


Figure 3.2: Work process loop.

Work process loop: first part. Every part of the work process loop consists of a set of technical and collaboration tasks, each generating a particular set of challenges. The first part of the work process loop deals with model transfer and validation work, and focuses on developing a work model that provides reasonably accurate fluid flow predictions to be regarded as trustworthy by the operator. Also, to efficiently serve as the computational engine underlying the optimization routine, the work model had to show a sufficiently fast and robust performance. The work model was implemented in a research reservoir simulator, which required several approximations from the original Eclipse model implementation (the simulator, AD-GPRS, will be presented in further detail later). A substantial amount of additional work can be spent on this part of the work process loop since any change to the original model, e.g., in initial well configuration or grid geometry, or update of information, e.g., new relative permeability curves or gas lift tables, may require a redo of model transfer and time-consuming validation work. In fact, because the industry reservoir model provided to us was regularly being updated and reworked during a period, this part did on several occasions during this project require costly supplementary work. However, the frequency of such rework is likely to diminish as coordination with the industry work process of the operator is improved. We believe, e.g., that once the optimization effort is well-understood, an improved communication with the reservoir team will help determine which changes to the original model that warrant an update of the work model.

Work process loop: second part. The second part of the work process loop deals with problem definition, optimization scope and application design in general. This part can pose various difficulties, given the significant amount of information analysis required to understand and reformulate operator knowledge into an optimization problem and scope. Changes to reservoir management strategy or base case configuration, for instance, are

likely to trigger updates in either problem definition, scope, or both, and may require substantial realignments of application development, and possibly loss of work. Moreover, the actual progression of the problem definition can be challenging. Throughout the development of the problem definition we considered several simplifications to the work model to test whether we could use more abstract problem definitions. For example, models with simplified grid structures, or models where free gas was removed from the simulation to decrease runtime and accelerate the control optimization part of the procedure, were tested. Eventually, through several consultations with the reservoir team, most of these alternatives were reconsidered (though gridding issues were addressed), since several of the implementations disregarded fundamental aspects for the production of the reservoir. Ultimately, this process was an important and very instructive part of the problem definition, though it demanded a costly back-and-forth between this part and the first part of the loop (since some of the simplified models were extensively compared to the original model).

Work process loop: third part. The third part of the work process loop is concerned with implementing and running the designed application. Primary focus has been on accurate computation of original code with subsequent extensions, and on developing a robust optimization framework. Building a robust optimization framework entails successfully integrating the original code with new code additions and/or extensions (e.g., a new way of handling parallel objective function evaluations on server), and finally also with replacements, such as the introduction of the new model and reservoir simulator, all into one reasonably efficient and coherent whole. Apart from implementing the designed features, much of the framework development work was spent on handling challenges that emerged during optimization. This has resulted in a programming structure with a large amount of patches and ad hoc solutions, which may complicate possible future developments towards a more general-purpose implementation. Moreover, the rudimentary construction itself is likely to reduce overall performance. However, these type of problems can be readily solved if we use the expertise acquired during this first build to redesign and rebuild the complete optimization framework into a more efficient programming structure for further research.

Secondary focus for this part of the work process loop has been algorithmic performance of each of the structure elements (as opposed to overall framework performance discussed above). In order to efficiently produce solutions we need to tune the various elements that make up the structure. For elements mainly dealing with optimization, central tuning parameters are those that control well placement search, e.g., the range of coordinate perturbation sizes (for the pattern search algorithm), and those that determine the extent of the embedded control optimization routine, e.g., the maximum number of method iterations, and the cap on total number of simulations for the routine. Other structure elements deal mainly with computation, e.g., the reservoir simulator. For these elements, tuning efforts are aimed at highest performance, but need to maintain accepted levels of accuracy, e.g., to achieve accurate produced volume calculations. Also, these efforts need to be configured to run at reasonable computational loads (given the limited amount of computer cores available). This involves tuning the extent of parallelization both at the lower reservoir simulation level, and at the upper level of the well placement algorithm. Ultimately, these type of tuning efforts need to be balanced against robustness in

cost function call execution. This means that an overall framework configuration needs to be found such that the work model simulations during optimization are able to efficiently handle (in the large majority of cases) any trial solution within a relative wide range of both well placement coordinates and controls. Principal configuration parameters in this respect are simulator solver tolerances, ranges of parallelization in each simulator run and for server job batches, and the establishing of a monitoring system to manage jobs and enforce kill-criteria for poor performance jobs, if necessary.

Work process loop: fourth part. Finally, in the fourth part of the work process loop, solutions from the optimization procedure are tested on the original reservoir model, and on a selected set of model realizations provided to us by Total E&P Norge AS. Essentially, solution testing involves adapting the solutions found by the procedure to the original model using standard industry tools, i.e., Petrel, and then running these new configurations using the original simulator, i.e., Eclipse. Solution testing is important because it enables us to communicate the results within the industry perspective of the operator, and it provides us with information about the effectiveness of the various work model approximations. Significant emphasis has been put on implementing all solutions using the original model. For this reason, Chapter 4 in this thesis performs somewhat comprehensive presentations and analyses of each of the re-implemented solutions. It also tests the solutions for case configurations that were either approximated (e.g., production time frame) or out of scope (e.g., multiple realizations) during optimization. The overall intention of this part of the work process loop is to provide potentially useful information back to the work process of the operator. Beyond this purpose, an additional function of the set of analyses is to serve as a practical result-interface for communication with the reservoir team. By having solutions recreated using standard industry tools and results presented and analyzed in common formats, the purpose is to facilitate commentary and feedback on the results. This information can then be used to further align the problem and constraints definitions to the business needs of the operator (thus starting a new iteration of the work process loop). Finally, a way of improving the communication task in this part of the work process loop could be to develop graphical user interfaces of core concepts of the problem definition. The function of these interfaces would be to serve as graphical representation of main features of the problem description, e.g., the different well placement constraints. Preferably, the re-modeling and adjustment of these interfaces would be a process performed directly on the graphical representation by the reservoir team, in interaction with the research team. Possibly, then, the use of these interfaces would facilitate the translation of expert problem understanding into specialized input for, in this case, constraint parametrization.

Conclusion. The work process structure just described represents the execution of the strategy components described in the beginning of this section. At the engineering level, this work process structure joins model validation, problem definition, extensions of developed methodology, and final testing of solutions on field case model, to solve for the second application target. The first target specified for this application work is addressed by the decisions and trade-offs considered during the design of the structure. IO Center Industry Partner Total E&P Norge AS has contributed to the development by providing expert knowledge and the field case model for us to test our application, in addition to offering substantial feedback on the results obtained. Below we will see how the specific

research developments are realized in the procedure that deals with the optimization part of the loop. The optimization procedure, or framework, is the actual application of our methodology, and the engine producing the solutions within the entire work effort. This framework is described in the next section.

3.2 Optimization framework

In the previous section we introduced the overall targets and strategy for this project. The plan for the execution of the strategy was presented in the form of a work process loop. In this section, we introduce the actual implementation of our methodology. The final implementation is made up of a collection of algorithms, solvers, and code extensions coupled with the reservoir simulator. We therefore present this implementation as an optimization framework³. An advantage of this framework is that the coupled parts may be developed independently, or new parts added, to deal with different planned and unplanned challenges during the implementation of the methodology. We will describe this modularity, and how it has been used to extend the functionality of our application to deal with a real field case. We will also present the various parts of the framework, focusing on their function, and how they have been assembled. Finally, we discuss some of the challenges in putting this application together, and ongoing work to enhance the framework.

3.2.1 Optimization framework

Framework introduction. Several challenges are associated with the application of our methodology when using a real field case. Throughout this section we will discuss some of these challenges and present the optimization procedure designed to deal with them. (Notice that in this section, we use the concepts “framework” and “procedure” interchangeably.) How these challenges have been resolved can be seen as the execution of our first strategy component, namely the emphasis on developing research based on our application work.

In the following, we offer a general description of the modular property of the framework. The idea is not only to show how the procedure is organized, but also to highlight how it can be improved by taking advantage of its modular structure. Using this property, the procedure may be enhanced by replacing components or adding extensions to solve the well placement and control problem in more efficient ways. The procedure could possibly also be developed to solve for other problems with similar integrated structure, e.g., problems involving the combination of well placement and/or controls with the design of facilities and pipe network (Litvak et al., 2002), and/or the routing of well streams (Foss et al., 2009). We then describe the various parts of the procedure, and how they function within the framework. We start by describing the algorithms for the well placement and control optimization. A procedure to enhance the joint optimization is also briefly discussed. We then explain how various capabilities from our previous application have been extended to deal with the real field case. Additional software introduced to handle some complex aspects of the case (e.g., the MRST software for gridding) are discussed thereafter. Finally, we describe the coupling of the reservoir model and simulator to the

framework.

In the discussion of each part we will refer to other sections in this work that deal with those topics in detail, and/or will provide references to relevant literature. At the end, we will also discuss the role of the IO Center as a collaborative network of Research and Industry Partners, and how the various partners have contributed to different parts of this work.

Overall purpose and design. The overall purpose of the optimization procedure is to provide decision-support to field development work tasks involving well placement. The optimization procedure is shown in Figure 3.3. (The details of the figure will be described later.) The design of the optimization procedure has mostly been guided by the first strategy component. Consequently, the optimization procedure shown in the Figure 3.3 makes possible the application of the main aspects of the developed methodology on the Martin Linge case provided to us. In accordance with the stated goals, the application design enables us to treat the real field case without compromising the use of the embedded optimization routine that is crucial for the joint approach. In fact, it allows for this approach to be readily implemented, and possibly also further enhanced. However, the unmitigated implementation of the joint approach forces the procedure to rely heavily on distributed computing and on an effective computation of gradients for the optimization of controls (among others issues, discussed further below). In total, the framework represents a first attempt at fulfilling the stated targets. We will proceed with describing the overall framework and its modular property, and then introduce each of its constituent parts.

Overall framework

Framework design. In this application work, we deal with an integrated problem that includes an optimization of controls at those well configurations that are explored during the search for optimal well placement. Moreover, the overall problem deals with a real field case which is more complex and has significantly higher computational demands, compared to our previous test cases. With this as our starting point, we think it is useful to consider the application of our methodology as a framework, or a collection of algorithms and solvers that, coupled with tools for reservoir grid handling and simulation, is specifically designed to deal with the embedded, or integrated, nature of a problem containing a challenging case. Given one of the main design features in our methodology is that the well placement and control problems are treated separately, it is important that the master search, the embedded procedure, and additional tools, are all well integrated. This means that once the various problem tasks to be solved are clearly differentiated, the procedures for solving these tasks need to be coupled together, and to additional software, to ultimately produce an application that functions reasonably well as a whole, and that is capable of handling a real field case. Conceptually, the framework perspective is useful because it promotes the effective integration of the optimization procedures, while enabling each procedure to specialize on its particular problem part. A key benefit of the framework design is that the relatively loose coupling between the procedures allows the development of the application to be more flexible. In our experience, this flexibility has been crucial when dealing with the planned and unplanned challenges posed by the optimization of a real field case.

Framework properties

Differentiation of problems, and integration. The optimization framework enables a procedure to solve a specific part of the well placement and control optimization problem, thus taking full advantage of the strength of each routine. By solving these problems separately, different search routines can be applied based on the particular characteristics of the well placement and control problem. For example, derivative-free methods are used to solve the well placement optimization problem since gradients with respect to well coordinates are not readily available. (Several approaches for approximating gradients for this problem have been proposed. Some of these approaches will be discussed later in this section.) Moreover, a significant advantage of the derivative-free pattern search routine chosen for the well placement part of the procedure, is that its function evaluations can be solved in a distributed manner. Separately, for the well control part of the procedure, the well controls at each well location are optimized efficiently using gradients computed by an adjoint-formulation in the AD-GPRS simulator. These gradients are routed to a derivative-based solver to solve the control problem as an independent part within the framework. Finally, by solving the well placement and control problem in a nested optimization, the joint approach integrates the two problems within the framework. At the end, the joint procedure is coupled to the computation of the set of well blocks with corresponding well indices that represents each of the well trajectories, and to a routine that manages the launching and performance of reservoir simulations.

Modularity. The modular nature of the optimization procedure is apparent from the differentiation and integration choices described above. (Modularity of system elements as a concept is further described in Barton (1992) within the context of system modeling and dynamic simulation of industrial processes.) One advantage of having a modular structure is that individual algorithms treating each of the problem parts can be readily changed. Also, needing to extend, add or replace only some of the parts eases implementation and may in some cases be sufficient to increase the capability of the application as a whole significantly. In fact, in the following, we will describe the various features (besides the algorithms themselves) and developments of the optimization framework, as either extensions, additions or replacements applied to elements of, or to the whole structure itself, to meet the different application challenges of the real field case.

Of importance for future work, the modular design also lends us the opportunity to develop enhancements aimed at exploiting the interaction between the different parts. Briefly discussed below, work is currently underway to improve the overall performance of the optimization procedure by using reduced-order models (Doren et al., 2006; He et al., 2011) built from sets of control solutions as surrogates for the control optimization routine during parts of the well placement search. Independently, we can, in a straightforward manner, increase or decrease the complexity to the individual parts, or insert distinct approximations, e.g., by replacing a full physics model with a surrogate model (Onwunali et al., 2008). Finally, by drawing on these various capabilities, we can further develop and adapt the framework to meet new challenges, cover more complex configurations, and even make customized applications for other types of integrated problems (see, e.g., Rahmawati et al., 2010, for an example of an integrated problem).

In the following we will describe the function of each of the parts in the optimization

procedure and explore some of the opportunities for further improvement. We start by giving brief introductions to the algorithms and solvers used in the optimization engine for the well placement and the control part, respectively. We then describe the various parts of the framework in terms of Extensions, Additions or Replacements. (These distinctions are here seen as useful to describe the current functionality of the solution system. Conceptually, we believe these distinctions will be helpful for further development of the optimization framework.)

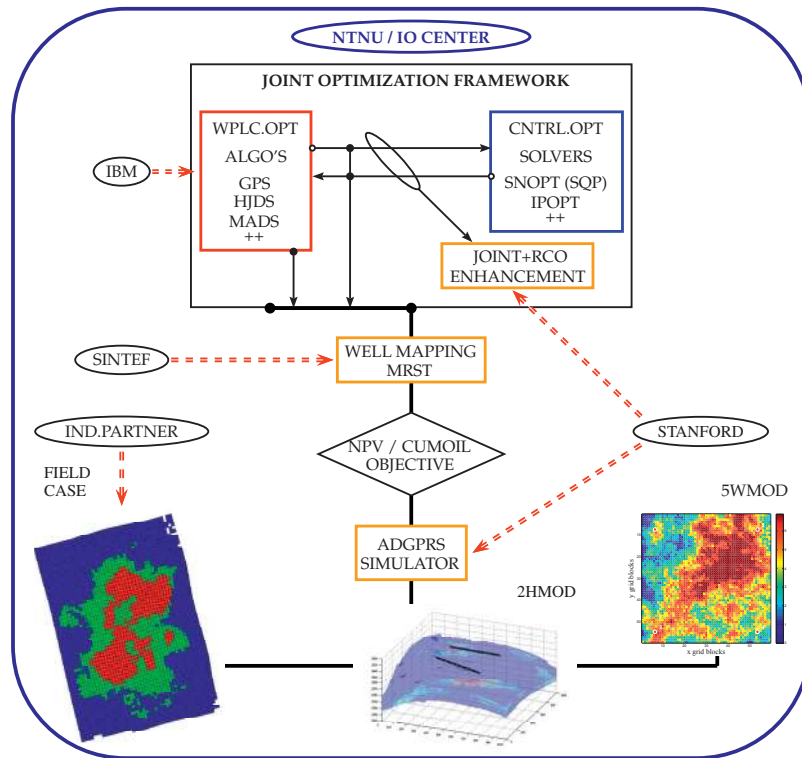


Figure 3.3: Optimization framework with three example models. Contributions from IO Center Research and Industry Partners to different parts of the framework are shown as red arrows.

Algorithms

Figure 3.3 shows the optimization framework. In this figure the core of the optimization procedure is represented by the black rectangle shown at the top. The rectangle consists of two main parts. The first part represents the well placement optimization procedure, and is exemplified by the inner red rectangle. The second part corresponds to the well control optimization, and is represented by the blue rectangle. An enhancement approach

currently in development is represented by the orange rectangle. In the following we give a general description of the well placement part of the optimization framework. A technical description of the complete optimization problem is given in Section 3.4.

Well placement part

The well placement part of the procedure is solved by a local search algorithm based on pattern search procedures (these have been introduced earlier). A very general description is that the algorithm searches the neighborhood of feasible well placement coordinates, and alters the location of the wells depending on which changes in trajectories that increase the pre-defined objective function (e.g., net present value (NPV), or in our case, field oil production total (FOPT)). The horizontal well trajectories in this application are defined by the real-space spatial coordinates corresponding to their heels and toes. Though the heel and toe coordinates of the horizontal wells are real variables, their actual representation within the reservoir simulator is as a set of discrete well blocks. As mentioned in Chapter 2, the discrete location representation of wells within the system of equations solved by the reservoir simulator causes an immediate lack of objective function derivatives with respect to well placement variables. For this reason, several key approaches within the literature focus on obtaining approximate information regarding well placement gradients.

Gradient-based approaches. In general, these approaches associate small displacements of the well location variables to continuous outputs, e.g., well rates, for which derivatives may be readily available, e.g., through an adjoint formulation. In this manner, for example, a chief approach has been to calculate improvement directions from efficiently computed production gradients corresponding to pseudo wells (i.e. wells operated at a negligible rate) that have been placed around the current well location (Zandvliet, 2008). The main advantage of this approach is that only one (forward) reservoir simulation, and then a (backward) adjoint simulation, of comparable computational cost, are required to compute improving directions for all wells (Zandvliet et al., 2008; Handels et al., 2007). Vlemmix et al. (2009) later applied this approach to horizontal well trajectories. Horizontal well trajectories, in addition to well type, have also been treated in Yeten et al. (2003).

Also other approaches that focus on converting the discrete well placement problem to a problem dealing with continuous variables have been presented. For example, in Wang et al. (2007) the location of a vertical injector is optimized by putting injectors in every grid block that does not contain a producer. The optimization procedure optimizes the rates of the wells, eliminating a well at an iteration if the rate of that well is close to zero, and using a maximum total injection rate allocated among the remaining injector wells. The implementation of the procedure in that work was seen as inefficient, though, since only one well could be eliminated at each iteration of the optimization process. In Forouzanfar et al. (2010) the well placement problem is also converted into a continuous optimization problem, while an initialization step is also introduced to determine total injection and production rates for the problem. Other approaches have mimicked a more traditional approach of computing finite differences. Bangerth et al. (2006), e.g., uses stochastic perturbation of well location variables to obtain derivative information

that may then be used within a standard solver. Complementary references can be found in Chapter 2.

Derivative-free approaches. A different approach within the literature is to target the well placement part using derivative-free methods for optimization. The general advantages of derivative-free methods are that they are non-invasive, that they can use the subsurface flow simulator as a black box during optimization, and that the global search characteristics of some of the methods may be well suited for problems with multiple optima, or rough optimization surfaces, such as the well placement problem. Different types of derivative-free methods have been used to treat the well placement problem, e.g., simulated annealing by Beckner and Song (1995) and neural networks by Centilmen et al. (1999). Recently, though, we see a greater use of derivative-free methods based on evolutionary algorithms, e.g., genetic algorithms (GA) and particle swarm optimization (PSO), being presented in the literature. See Echeverría Ciaurri et al. (2011b) for an extended comparison of gradient-based and derivative-free methods for common optimization problems in petroleum operations. In that work, the authors show that derivative-free methods, in particular genetic algorithms, particle swarm optimization and pattern search methods, are viable methodologies for a range of oil field applications.

Pattern search methods. In this work we have chosen to implement derivative-free pattern search methods to solve for the well placement part of the procedure. We use pattern search procedures because these methods are deterministic and rely on solid convergence theory developed in the last decades (Torczon, 1997). In particular, we have used general pattern search (GPS) in our implementation (Kolda et al., 2003). This particular algorithm is relatively straightforward to implement, and fairly robust, e.g., against possible crashes of the reservoir simulator for some cost function calls, which is an important characteristic when dealing with a field case application. Moreover, we draw heavily on the advantage of these algorithms on computing many of the trial solutions in a distributed manner. A thorough description of the optimization procedure for the GPS algorithms is given in Echeverría Ciaurri et al. (2011b). Because of the high cost of each objective function evaluation, in our implementation we have taken particular care that function calls are only made for unique trial solutions. The reason is that the implementation of the various non-linear constraints sometimes resulted in identical or very similar trial solutions during optimization, i.e., a varying number of the vectors in the set of trial solutions during one iteration of the algorithm were sometimes equal, or very similar by some tolerance. A check was therefore implemented to only compute unique instances of trial vectors. Also, the implementation relied on the use of a cache of previously computed function evaluations. Further discussion of derivative-free methodologies can be found in Conn et al. (2009).

Advanced applications. Significantly more advanced implementations of pattern search algorithms exist that can be introduced to solve for the well placement part of the framework. One of the main goals of applying more sophisticated pattern search implementations is to reduce the total number of cost function evaluations, which is particularly important when dealing with expensive reservoir simulations. For example, interesting implementations of pattern search algorithms exist that approximate some gradient information using cost function values evaluated earlier in the procedure (Custodio and Vicente, 2007). Also, there is further interesting research activity directed at the devel-

oment and performance analysis of derivative-free algorithms, both model-based and direct search type of procedures, that use expensive simulation-based objective functions (Wild, 2009). Finally, further work should explore insights from bilevel optimization theory. The study of the well placement and control problem from that perspective could possibly add interesting improvements to how the optimization procedure is solved.

Well control part

Optimization of production strategy has been discussed in Chapter 2, and for this particular case is further defined in Section 3.4. Here, we will only describe the role of this optimization part within the optimization framework. The well control optimization part is represented by the blue box in Figure 3.3. Different from previous work, in this case we performed the well control optimization using the optimizer module implemented in the AD-GPRS simulator. As before, though, we used the SNOPT solver, which is the SQP implementation developed by Gill et al. (2005), this time pre-installed within the module. An important consideration of launching the control optimization in an embedded manner is to tune the algorithm to keep the number of reservoir simulations required by the control optimizer as small as possible, while still allowing the SQP solver to yield a significant increase in cost function during well placement search. Importantly, for our implementation of the joint approach, this consideration needs to hold for most part of the procedure, i.e., we need to allocate enough resources to the control optimizer so that an increase in cost function for very different well configurations is likely. In this implementation, the main tuning parameters to achieve a sufficient number of SQP steps at various well locations were the maximum number of major iterations, and the total number of reservoir simulations. The embedded control optimization was stopped once either one of these two limits was reached. The final setting of these limits was achieved after extensive testing. The application of other solvers, e.g., IPOPT (Wachter and Biegler, 2006) or Method of Moving Asymptotes (MMA; Svanberg, 2002), which may yield similar cost function improvements while using fewer calls to the reservoir simulator, is currently underway.

Enhancements

A guiding principle behind the design of the framework is that each of the constituent problems is solved using an adequate optimization methodology. Because the subproblems are clearly structured, we might take advantage of the coupling between the different solution procedures. In this context, we refer to such a development of the framework as an enhancement. An enhancement that is currently being worked on is represented by the orange rectangle in Figure 3.3. In the term Joint+RCO, “RCO” is an abbreviation for Reduced-Order Control Optimization. Joint+RCO is an enhancement where the control optimization routine is accelerated through the use of a surrogate. The surrogate is built based on reduced-order techniques, and is meant to decrease the total number of reservoir simulations required by the control routine. We briefly outline the procedure for this enhancement below. Further details are given in Appendix B (page 218).

Reduced-order control optimization. The procedure aims at maintaining the gain obtained from the embedded optimization of controls, but to reduce the total number of

reservoir simulation calls that this approach entails. The GPS algorithm proceeds by solving for sets of well placement trial solutions (each complete set is hereby referred to as a poll set, and the computation of the poll set is referred to as the polling procedure). For the joint approach, the polling procedure entails solving a (full-order) control optimization at each well placement trial solution using a relatively high limit of major iterations in SNOPT. The idea behind the Joint+RCO enhancement is to replace a certain number of poll sets that use the regular, full-order, optimizations of controls, with poll sets that instead perform reduced-order control optimizations. Because these optimizations are performed on a much lower number of control variables, the optimization procedure is likely to require fewer function evaluations (simulations) than the full-order control optimizations to achieve an adequate solution for our purposes. The reduced-order vector of well controls is projected back to full-order space to run the simulations.

Within the optimization framework we are presenting, the enhancement is implemented as follows. Initially, a well placement poll set using full-order control optimizations is run to yield training data for the surrogate procedure. Once the required data for the surrogate procedure is assembled, and a projection matrix computed, a number of subsequent poll sets are then run using reduced-order control optimizations at their well placement trial solutions. After a number of poll searches, and possibly other criteria, e.g., change in poll step size, a new poll set using full-order control optimizations is run to re-train the procedure. The purpose of the enhancement is solving embedded (reduced-order) control problems between training runs that yield gains in objective function value that are similar to those from solving the full-order problem. Furthermore, because of the smaller dimension of the reduced-order problem, the idea is that these problems can be solved satisfactorily using a lower limit of major iterations in SNOPT (or, some other finishing criterion, if using another optimizer). Thus, the joint approach is approximated by performing surrogate control optimizations at intermediate poll sets. These surrogates provide sufficient improvement over initial controls but are less costly than standard procedures, and are therefore likely to reduce the total number of reservoir simulation calls needed for the approach. Below, we provide a general explanation about how the surrogates are built.

Overall, the enhancement consists of two main steps. As mentioned, the first step consists of performing a regular polling procedure where a (full-order) control optimization is launched at each well placement trial solution. Once this polling procedure finishes, all the control solutions are collected and stored in a “snapshot” matrix \mathbf{U} (the process is somewhat similar to the one in Cardoso and Durlofsky (2010) where saturations and pressures are collected into snapshot matrices at different times during simulation, though for a very different implementation). Using this snapshot matrix, we calculate a projection matrix Φ by performing a singular value decomposition of \mathbf{U} .

At a subsequent number of GPS polling procedures, Φ is used to reduce the embedded control problem of each trial solution, i.e., we perform $\mathbf{z} = \Phi^T \mathbf{u}$, where \mathbf{z} and \mathbf{u} are the reduced-order and full-order control vectors, respectively (corresponding gradient vectors are reduced in identical fashion). The control problem within these GPS polling sets is solved using the same gradient-based solver as before, but the optimization is now solved in the \mathbf{z} -space. As mentioned, the dimension of \mathbf{z} is much lower than the dimension of \mathbf{u} , so we are able to run SNOPT on the reduced-order control problem using a

lower limit of major iterations, and thereby, in the majority of cases, reduce the number of reservoir simulations required for the embedded control optimization. This enhancement has been implemented on a case similar to the one in Chapter 2 (two injectors and three producers, though with double the number of control variables) with promising results. Because the full-order dimension of the control problem is significantly higher in the new case compared to the original one, there is a greater potential for reducing the number of cost function calls by making use of the surrogate procedure. In general terms, the implementation of this enhancement can potentially yield substantial reductions in total computational cost once more complex control optimization problems are introduced in the framework. Pseudo-code given in Appendix B (page 218) further describes the implementation of the two steps of the enhancement procedure. We now proceed to describe how existing functions from the implementation in Chapter 2 have been extended to deal with the real field application.

Extensions

In the current context of our optimization framework, an extension refers to a development that adds to one or several existing capabilities. In this section, we discuss extensions to capabilities that were already developed in previous work, though in a much simpler form. In particular, we refer to the calculation of well indices for arbitrary well trajectories, and to the more extensive application of distributed computing. These two extensions are discussed in the following.

Horizontal well index calculation. This extension entailed developing auxiliary code to be able to optimize on the trajectories of horizontal wells. Previous work had only dealt with vertical wells and used the standard Peaceman relationship that applies to vertical well bores traversing grid blocks perpendicularly (Peaceman, 1978) for the calculation of the well index (also referred to as *connection transmissibility factor*, see Schlumberger, 2012b). The well index calculation serves as a proportionality constant relating the fluid flow entering or leaving the well block, and the pressure differential existing between the well block and the reservoir. This connection factor depends chiefly on the geometry of the grid block, the well-bore radius, and the rock permeability (Schlumberger, 2012b). In the current implementation, Peaceman's formula is still at the core of our well index calculation. However, the formula is now applied in piecewise manner to segments of wells. During optimization, the well heel and toe coordinates given by the well placement algorithm are converted into well trajectories, where each trajectory traverses a number of grid blocks. (These traversed grid blocks are hereon referred to as well blocks.) From here it follows that each well block contains a length segment of the overall well trajectory. In our extension, Peaceman's formula is applied to each individual well segment to calculate the connection factor of each well block. Crucially, the formula is modified by a proportionality constant that relates the length of the well segment to the well block geometry. Roughly, a short penetrating well segment will produce a low well connection factor, and vice versa.

Making the necessary geometry calculations, specially when dealing with a corner-point grid and deviated well trajectories, is not trivial. To implement the extension, the main functions that needed to be solved were which reservoir model grid blocks that were

traversed by the deviated well bore, and the length of the well bore actually penetrating each of these grid blocks. The current implementation is represented by the orange box labeled “WELL MAPPING MRST” in Figure 3.3. This part of the framework was developed in close collaboration with SINTEF Applied Mathematics who are the main developers of the MATLAB Reservoir Simulation Toolbox (MRST; Lie et al., 2012) used for this implementation. This collaboration work has been important for this application effort, and is further discussed at the end of this section. The field case Eclipse model was imported onto MATLAB using the MRST toolbox. The open-source MRST toolbox allowed us to efficiently treat all necessary grid data, and to develop the different well mapping functions used by the optimization procedure to make the conversion from optimization variables to well perforations with corresponding well connection factors.

Distributed computing. The computational load of the optimization procedure has increased substantially due to the effort to achieve the targets set forth for this application work. Broadly, the two targets have been a reasonable field case application, and the possibility to fully test the developed research approaches even though using a field case model. To meet these targets, it has been important that we improve the ability of the framework to efficiently distribute computational load. In fact, meeting the increased demand for calculations by extending the capability for parallelization possessed by several of the core functions in the procedure is a critical element in both the strategy components stated before. This means we rely heavily on being able to extend the parallel capabilities of our previous work to meet the challenges posed by the new application. In the following we focus on two main developments aimed at increasing the parallel capability of the optimization framework.

The first development was to step away from a parallel computation of jobs using the MATLAB Distributed Computing Toolbox. The main reason was that, due to licensing constraints, this implementation was too restrictive on the number of function evaluations that could be performed concurrently. To bypass this constraint, an execution process was engineered that interfaced directly to the Torque job scheduler on the server. In this manner, the pattern search poll set can be sent directly to the server as a batch of independent jobs, each job representing a different polling point in well coordinate space. A monitoring function was developed to read and combine state data from the scheduler regarding each job with simulation log data, and to manage overall job execution according to different performance criteria.

The second development involves the implementation of the AD-GPRS simulator itself. The advantages of avoiding any licensing limitations by implementing the reservoir model on a research simulator are discussed elsewhere in this chapter. In this argumentation we would like to characterize the choice of selecting a research simulator as our main computational engine as a choice pertaining the further development of the parallel capabilities of the optimization framework. The reason for this description is to emphasize that the substantial development in execution and management of cost function evaluations as independent server jobs described above, is only useful if proprietary licenses are not serving as bottlenecks. That is, we can take full advantage of the first development described above only if the number of simulation calls that we can make concurrently is unconstrained. Our overall point is that the strategy element of putting substantial focus on the development of further parallel capabilities for the optimization framework,

creates a powerful argument for selecting to perform optimizations using a work model implemented in a research simulator. Next, we will discuss the contribution of additional functionalities to the further development of the optimization framework.

Additions

In the current context of our optimization framework, an addition refers to a development that introduces a new capability into the system. An important addition to this application effort has been the introduction of non-linear constraints into the well placement part of the procedure. A detailed description of the constraints, and the constraint handling procedure in general, is given in Section 3.4. Therefore, only a brief description of the non-linear constraint handling capability, and its development within the optimization framework, is given here.

Non-linear constraints for well placement part. A clear new addition to the optimization procedure has been the introduction of non-linear constraints to the well placement part of the framework. These constraints, including bounds, are important because they define the scope of the well placement search. The implementation of such constraints may be challenging for reasons like complex grid geometry, irregular reservoir bounds and a highly intermittent pattern of inactive grid blocks within the grid itself. (Although, computationally, the introduction of these constraints is relatively cheap since their implementation does not require additional reservoir simulations.)

Constraint development has progressed following feedback from Total E&P Norge AS. For example, an important consideration within the well placement work task has been the degree of uncertainty associated with drilling relatively long well bores. The uncertainty may originate from the actual production from long drains, from concerns about drilling beyond good reservoir sands, or other field development factors. Once this consideration was specified by the field development team we proceeded with adding what eventually became different implementations of the well-length constraint (see Section 3.4).

Based on our experience with this field case application effort, we notice that translating field case understanding into workable constraint definitions can be a challenging process. For this argumentation, we view this process as an effort in information analysis and retrieval. We furthermore classify three types of information met during our work as (1) artifacts, (2) values and rules, and (3) tacit, underlying assumptions. (These distinctions are loosely based on the three types of levels that Schein (1992) defines to analyze organizational culture. We do not claim any type of analysis power from our distinctions; rather, we use these here only to help us describe the process of constraint definition from our point of view.) Artifacts in this context are the documents and data provided and/or otherwise available to us, e.g., published articles, publicly available documents (e.g., *Konsekvensutredning* from PUD; see Total E&P Norge AS, 2011), and, most importantly, the reservoir simulation model. These artifacts are the actual representation of underlying values and rules about how to develop the reservoir, which rest on engineering knowledge and assumptions gathered over time about the prospect. Our point is that developing a set of constraints that answer reasonably well to the information needed by the field development team can be challenging for various reasons. First, because we may have to filter

through a large number of artifacts to both gain a general understanding of the problem, and then to find specific information regarding the work task that we wish to complement with our optimization routine. Secondly, we may have to both infer from this data and, preferably, through a series of constructive dialogues (Isaacs, 1999), an idea about what kind of solutions that are reasonable with respect to overall values and rules for the development. Thirdly, through extensive collaboration we might be able to understand underlying assumptions that will help us create and assess the various constraints needed for optimization. In summary, in our experience, constraint and problem specification can be demanding, and may often require an efficient collaborative effort. This type of effort is further described in our discussion of the IO Center as a platform for collaboration between Research and Industry Partners, in the last part of this section. Next, we end the description of the framework by discussing how elements of the procedure can be replaced to attain different functionalities.

Replacements

Lastly, the modularity property of the optimization framework allows us to replace any of the constituent parts, if necessary. New parts that serve the same function as older ones may be introduced if they perform the same function more efficiently, or if they possess additional properties that may make their inclusion advantageous to the overall structure, or if they constitute a necessary fit for a new case. A development that substitutes a part of the framework with another part having at least equivalent function, is referred to as a replacement. Future developments to the framework are likely to focus on replacing the well mapping module, as well as introducing alternative reservoir simulators, and/or substituting reservoir models to deal with other cases. We briefly discuss each of these developments below.

Replacement of well mapping module. As described earlier, the well mapping module contains important functions that convert well placement variables into well completions that can be used in the reservoir simulator. The current implementation works appropriately for the current application, but requires substantial redevelopment if it is to operate robustly across different case models. (The planned rework is significant enough to warrant this development to be labeled a replacement rather than an extension.) The redevelopment would also yield the opportunity to perform comprehensive tests of the well connection factors produced by the module, e.g., across an extensive suite of different well trajectories, against well indices computed using standard industry tools, such as Petrel. An additional benefit is that this would help validate the module for the future coupling with an industry-standard simulator like Eclipse.

Alternative reservoir simulators. Conceptually, any type of fluid flow predictor (from full-physics simulators to surrogates) can be introduced into the optimization framework. Current efforts are focused on introducing the MRST simulator as the main simulator engine for the framework, and on developing future cases within this platform (more on this below). However, we also recognize the advantage of achieving optimization results using the field case model, often implemented in Eclipse, directly in the procedure. We are therefore also exploring workflows that would permit the use the Eclipse simulator within the optimization framework. In this respect we believe, given current capability,

that a reasonable workflow is to use the Eclipse simulator only to solve a problem with very limited scope. Preferably, this narrow scope would be the end–result of an extensive optimization procedure performed using a work model implemented in a research simulator.

New models for new cases. We end this part with the straightforward replacement that entails the introduction of a new model into the framework. A new model may signify a new case, or it may mean optimizing for the same case but using an alternative, e.g., simpler model, or surrogate. At the bottom of Figure 3.3 we show three models representing three different cases. The rightmost model is the same model as implemented in Chapter 2, while the leftmost model is the Martin Linge oil reservoir model treated in this chapter. The center model at the bottom is a horizontal well model currently under development. This model is mainly meant for research purposes and will ultimately be implemented on all three simulator platforms (MRST, AD–GPRS and Eclipse).

3.2.2 IO Center resource platform and network

We end this section by describing NTNU’s Center for Integrated Operations in the Petroleum Industry (NTNU/IO-Center, 2014), as a resource platform, and the activity of its Research and Industry Partners as an open–source network for collaboration. This platform has been a crucial source of knowledge with respect to our work, and the collaboration with the various partners has been an important drive for the various solutions developed in this thesis.

Since its start in 2007, the IO Center has delivered noticeable contributions to the market of ideas and applied research in the Petroleum Industry, with particular relevance to the Norwegian sector. In the following we will briefly describe the IO Center as a resource platform for integrated research. This platform (represented in Figure 3.3 as the large blue square encircling the entire optimization framework) supports the network of Research and Industry Partners that in turn make collaborative activities possible. This network has facilitated the different contributions to the work in this thesis. (Notice that the following description of activities within the IO Center is based on our perspective and work within, and on topics related to, the center’s IO4 Program: “Production Optimization and Subsurface IO”.)

Resource platform. One of the main functions of the IO Center is to serve as a platform for Industry Partners (comprising of both petroleum field operators and suppliers) to expose current industry topics to ongoing research activities. At the same, the platform allows Research Partners to fully engage current operational problems, and challenges them to further extend and develop the scope and applicability of their methods. Phase I of the IO Center emphasized the development of methodologies within several research topics important to the industry. The focus of Phase II of the IO Center is to apply the methodologies developed in Phase I on current industry cases. The aim is to provide benefits by improving operations and industry work tasks through research methodology (i.e., industrial applications and innovation enabling; NTNU/IO Center, 2013). To this end, one general strategy has been to extend and adapt existing methodologies into applications that may contribute directly to industry work processes. In our case, the developed methodology for well placement optimization was coupled with ongoing operational work aimed

at the development of a North Sea field operated by Industry Partner Total E&P Norge AS. Furthermore, throughout this work, the IO Center resource platform has allowed us to draw from and combine technologies between the different Research Partners to solve the challenges posed by the field case application. We describe the various contributions as a network of collaboration.

Collaborative network. In Figure 3.3, red arrows represent the various contributions from Research and Industry Partners to the different parts of the optimization framework presented in this section. As stated, our main collaboration partner from the industry side has been Total E&P Norge AS. Total E&P Norge AS has contributed substantially to the development of this work. One primary contribution has been the allocation of resources that enabled ample collaboration between us and the engineering team dealing with the development of the Martin Linge field. Aspects of this collaboration have already been described as part of the work process loop in Section 3.1.2. Through this collaboration we have been able to closely treat and discuss project stages such as problem and constraint definition, model validation, effective implementation of optimization framework, and to obtain expert feedback from our solution tests on the field case model.

Furthermore, we need to emphasize the collaborative, open-source, network composed of IO Center Research Partners that has been a main enabler in this application effort. Our primary research partners have been Stanford University, IBM T.J. Watson Research Center, and SINTEF Applied Mathematics. These partners have each contributed directly, not only with applications developed within their respective area of expertise, but also with specific knowledge about how to apply and customize those applications to devise the extensions, additions and replacements necessary to make possible the application of our methodology.

Overall, we have had a case provided by Total E&P Norge AS, which we have solved using algorithms suggested by IBM research, using simulator and control optimization software from Stanford, and using SINTEF code to resolve any issue regarding wells within the reservoir model grid. An example of collaboration has been, e.g, while using the AD-GPRS reservoir simulator developed by the SUPRI-B group at Stanford. In this case, we collaborated on how to formulate the control problem within the optimizer module given the particular model features introduced by the validation work. We also received substantial support from the AD-GPRS development team to modify some aspects of the optimizer module to perform the control optimizations more efficiently within our application framework. Moreover, all functions within the optimization procedure dealing with reservoir grid coordinates and well definitions have required substantial customization of the MRST software provided by SINTEF.

Framework summary

Finally, one of our main tasks has been to integrate these different contributions into a framework for well placement optimization that both embodies the original research developed in Chapter 2, and manages to perform a reasonable optimization using a field model. We have organized our complete application in terms of a framework that emphasizes the coupling between developed procedures and support software that exists with the IO Center platform. We believe this organization provides a useful flexibility to the

application effort, allowing it to expand and adapt, and thus to efficiently handle the different challenges posed by a field case application. The next section describes the Martin Linge field and oil reservoir in detail, and deals with the model validation effort. Subsequent sections, Sections 3.4 and 3.5, provide the problem formulation and optimization results.

3.3 Field case and validation work

3.3.1 Martin Linge field case

Martin Linge field introduction. This section starts with a brief introduction of the Martin Linge field development project. It then focuses on the Martin Linge oil reservoir, where we describe, in general terms, those aspects of the reservoir development that have the greatest impact on our work. In Section 3.3.3 we go through the validation process following the transfer of the reservoir model to our research simulator. (Reasons for transfer and other issues have previously been discussed in Section 3.2). A comprehensive description of the base case setup for the Martin Linge oil reservoir model, as well as production strategy and specific model parameters, is given in the beginning of Section 4.1.

Martin Linge field reserves. Martin Linge is an offshore oil and gas field on the Norwegian Continental Shelf. The field is located at North Sea coordinates ($E 2^{\circ} 0' 53.403''$, $N 60^{\circ} 30' 22.302''$), near the border of the British sector (Total E&P Norge AS, 2011); see Figure 3.4. Though considered a single asset, the Martin Linge field consists of two independent hydrocarbons reserves. The largest of these reserves, Martin Linge East (MLE), contains mainly gas and is found in the Brent formation at depths between 3700 to 4400 meters (NPD, 2013b). In this work, we consider only the main hydrocarbon accumulation at the smallest of the Martin Linge reserves, i.e., the Martin Linge oil reservoir (MLO). This reservoir is found in the Frigg formation at a depth of 1750 meters (NPD, 2013b). Current plan for field development is to produce both the Brent and Frigg reserves



Figure 3.4: Martin Linge field (Total.com, 2013).

concurrently. Four horizontal production wells are planned in the base case for the development of the Martin Linge oil reservoir. This field case was implemented within the optimization framework presented in Section 3.2.

Economic basis for development. The Frigg and Brent hydrocarbon reserves were discovered already in the late 1970s. The major reason why these reserves were not considered for development until now was the high uncertainty associated with the structural settings in several of the reservoirs (Boutaud de la Combe et al., 2012). Since then, several studies of the underground geology, e.g., seismic and extended well tests (Douillard et al., 2009), have contributed to a better understanding of the structural complexity of the reserves. The increase in knowledge about the subsurface geology over the last 40 years, and the current availability of efficient production techniques (e.g., horizontal well drilling) have been important factors in reassessing the likelihood of a successful development of the reserves. The field development plan is being assembled by field operator Total E&P Norge AS. Total E&P Norge AS holds an ownership share of 51% of the production license. Partners are Petoro AS and Statoil AS, which hold shares of 30% and 19%, respectively (NPD, 2013a). Field production is planned to start by the end of 2016 (NPD, 2013b). General features of the overall development plan are briefly outlined below. Following this description, we focus solely on the Martin Linge oil reservoir.




Figure 3.5: Martin Linge field development (TU, 2013).

Overall field development. An illustration of the overall field development plan is shown in Figure 3.5. The Martin Linge field will be developed using a sub-sea installation and topside facilities, which will be supplied with electric power from shore (Thibaut and Leforgeais, 2012). The development comprises a platform with a jackup rig and a Floating Storage Offloading unit (FSO) for oil and condensate storage, where oil, water and condensates will be stored and processed (Total E&P Norge AS, 2011; NPD, 2013b). Water will be separated for re-injection in the FSO. Oil will be exported via shuttle tankers and processed gas will be exported to St. Fergus (United Kingdom) via a new gas pipe

link to the existing Frigg UK Pipeline (FUKA) (Thibaut and Leforgeais, 2012). The well program is to drill six producers to the Brent accumulation, and four producers for the production of the oil reservoir in the Frigg formation (i.e., MLO). One additional well will be drilled for injection of produced water (Total E&P Norge AS, 2011).

3.3.2 Martin Linge oil reservoir

Martin Linge oil reservoir introduction. The Martin Linge oil reservoir is a relatively shallow reservoir containing viscous oil overlaid by a small gas cap. The four horizontal drains specified in the production base case are set to produce with gas lift. Fluid flow rates predicted for this case are obtained using lift gas injection rates optimized by the reservoir simulator. The overall production case is characterized by early water breakthrough due to a high mobility ratio favoring water and the presence of a strong aquifer. The aquifer is thought to provide sufficient pressure support during production, so additional measures for pressure maintenance, e.g., water injection wells, are not included in the current development plan. The net-to-gross (NTG) ratio of the sands comprising the Martin Linge oil reservoir is in the order of 70 to 80%, with a moderate amount of shales thought to be distributed over the region in an intermittent pattern. Overall, reservoir sands are regarded as having good porosity and high permeability. Figure 3.6 shows the Martin Linge oil and gas reservoirs. MLO is the shallow (i.e., topmost) reservoir in the figure.

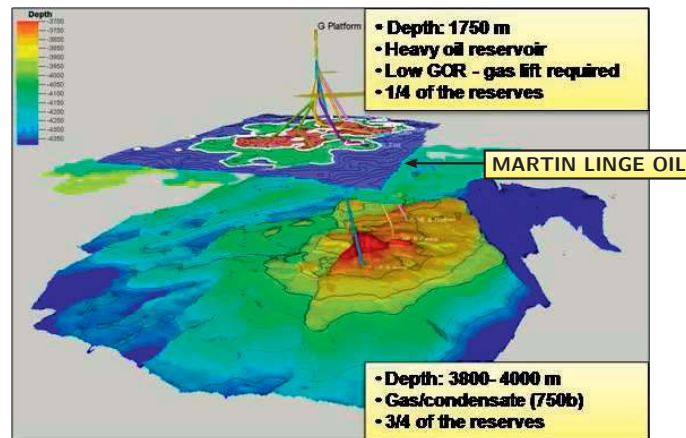


Figure 3.6: Martin Linge oil and gas reservoirs. MLO is the shallow reservoir in the figure.

Reservoir structure. MLO is a sandstone reservoir formed on a sedimentary fan. The sands in this system are considered to be highly permeable, with good vertical and horizontal pressure communication, and limited faulting. However, these sands are reckoned to be poorly consolidated, and a moderate amount of formation shales is expected to be found more or less intermittently throughout the reservoir structure⁴. Moreover, given the sedimentary structure of the reservoir, the quality of the sands is thought to degrade toward the outer fan area, i.e., north. The reservoir contains a structural height that stores the main part of the hydrocarbon accumulation. Locally, the height contains two small

elevations with a corresponding depression in the middle, creating a central saddle area in the reservoir. Relatively small amounts of free gas are accumulated at these elevations points. The oil layer is found below the saddle area, and runs more or less continuous, i.e., laterally, throughout the rest of the structure. The oil column for the MLO accumulation is somewhat around 20-plus meters, while the overlying gas column is considered to be maximum about 20 meters thick.

Flow properties. The MLO reservoir contains saturated oil with dry gas in its overlying gas cap. The oil has a relatively high viscosity of around 5 cP at reservoir conditions. The difference in viscosity results in a high mobility ratio, making water significantly more mobile than oil. The high mobility ratio, coupled with the existence of a large aquifer below the oil layer, yields an important flow dynamic in the reservoir. Measurements suggest the reservoir will receive substantial pressure support from the aquifer. If this is the case, it is likely the displacement of oil by water coming in from the aquifer becomes less effective due to water rushing past oil, and we can expect water breakthrough times and high water production rates early and throughout most of the production time frame (see also Section 3.3.3 for more discussion about fluid flow topics, e.g., hysteresis). It is worth mentioning that a possible upside of geological uncertainty, is that the inflow from the aquifer turns out to be weaker than expected. In this case, the reservoir might lose some overall pressure support, but the weaker aquifer inflow is likely to make the displacement of oil more effective, and thus lead to higher recovery.

Production schedule. The field operator has developed a base case reservoir simulation model for the production of the Martin Linge oil reservoir. The base case model specifies both the well trajectories and the management of the wells for the entire production time frame. The model is specified with a production time frame of 5174 days, starting late 2016. Within this time frame, the reservoir simulator operates all wells following a set of scheduled controls. The control schedule includes well definitions that specify the preferred phase of the wells, e.g., if wells are mainly expected to produce gas or oil, and their planned state of operation, e.g., whether they are to be opened or closed at any particular time during simulation. Three of the four wells in the base case scenario are set to start producing at the same time (the late-2016 production start date). The fourth well is set to start producing from the reservoir more than a year after general production start.

Well control. Importantly, the well schedule specifies the main parameter setting for the operation of each well at any given time during simulation. The setting determines if well flow is primarily controlled by rate (e.g., gas, oil, water or liquid rate) or pressure (well tubing head or bottom hole pressure). Wells may switch primary control setting during simulation, e.g., from rate to pressure control, and vice versa, if limiting pressures or target rates have been specified for their operation, and if these have been reached during simulation. Limiting pressures and target rates are commonly specified for both individual wells and groups of wells. Additionally, general field production constraints, often determined by overall well design and planned capacity for the field, may be included into the well control specifications. The more advanced production techniques, e.g., production using lift gas injection, and enforcement of group well controls and field constraints, have their own operational parameters, as well as independent production targets and limits. These techniques also usually have overarching control over individual well pressures and rates (Schlumberger, 2012b). For the MLO base case, it is particularly important to

honor the field liquid and gas rate constraints imposed on production, since these originate directly from the designed fluid and gas handling capacity of the facility. These topics are also discussed in the Field model transfer and validation section next. In Section 4.1 we describe in further detail the specific reservoir simulator keywords, and their parametrization, used for field and well controls in the base case.

Introduction to validation section. In the next section we discuss the validation effort and present the results from the comparison between the original base case model Eclipse implementation and our implementation on the AD-GPRS research simulator. We have previously discussed the reasons for the transfer in Section 3.2: Optimization framework.

3.3.3 Field model transfer and validation

For efficiency reasons, we develop a work model based on the field case simulation model of the Martin Linge oil reservoir. This work model is involved in all reservoir simulations launched by the optimization framework presented in Section 3.2. This section explains the transfer procedure and presents the validation data from the development of the work model. A general purpose research simulator is used for work model simulations, while the original field case model is implemented in an industry-standard reservoir simulator. This section is divided into two parts. The first part describes the field case simulation model and production strategy. This part also introduces the research simulator, and outlines the transfer procedure for the development of the work model. Finally, the first part presents the main validation results and extensive comparisons of the developed work model configuration against the original field model. Clearly, the extensive transfer and validation effort has focused on reaching sufficient overlap in production curves from the original field model and approximated work model. At the same time we have emphasized developing a work model that is reasonably fast and sufficiently robust during optimization. The second part of this section briefly describes those field case model features and simulator functions that required modification. This part provides rough explanations of how these modifications were implemented in the different cases. Challenges from this transfer and validation effort have been discussed in Section 3.2. For overview, the first part of this section may be read independently of the second.

First part: Model transfer and validation results

Here we present the field case simulation model and production strategy for the development of the Martin Linge oil reservoir. We also introduce the general purpose research simulator used to develop our work model, and briefly describe the model transfer procedure. The procedure consists of testing different model approximations by running a progression of simulation cases. At the end we choose which collection of approximations we think is the best based on simulation performance and accuracy with respect to the original field model. Validation results with corresponding production profiles are presented at the end of this section.

Field case simulation model. IO Center industry partner Total E&P Norge AS has provided the field case reservoir simulation model used for the development of the Martin Linge oil reservoir. The field case model is implemented using the industry-standard Eclipse reservoir simulator (Schlumberger, 2012a). This model consists of approximately 55000 active grid cells. The version of the base case model provided to us, runs a fluid flow simulation for approximately 14 years, spanning a planned production time frame set from late 2016 to early 2031.

The base case production strategy for the Martin Linge Oil reservoir is based on production from four horizontal wells. Three of these wells will be drilled on stream at start-up. The fourth well will be drilled almost two years after production start. Current development strategy builds on natural pressure depletion drive for the wells, gas injection for artificial lift, and sustained pressure support from a large aquifer flanking the reservoir. While managing the existing drive energy towards increased oil recovery, the development strategy for the reservoir needs to operate within individual and group well rate and pressure constraints. Operational bounds on single wells and groups of wells are mostly determined by individual well bore configurations (e.g., tubing diameter) and the fluid-handling capacity allocated to the platform. For this particular development, close control of individually and collective gas and fluid production well rates is an important element in the production strategy.

AD-GPRS simulator. The work model is implemented using Stanford's Automatic-Differentiation General Purpose Reservoir Simulator (AD-GPRS; Voskov and Zhou, 2012; Tchelepi and Aziz, 2012). The optimization framework applied in this work uses the AD-GPRS research simulator to both run simple work model simulations and to optimize well controls for increased field oil recovery. Optimization problems are handled using the AD-GPRS optimizer module (included in the simulator). The optimizer module relies on the effective computation of objective function gradients implemented in the AD-GPRS simulator. This implementation is based on a discrete adjoint formulation using automatic differentiation (Volkov and Kourounis, 2012). The efficient computation of gradients allows the optimizer module to deal with problems that may require a large number of simulator runs. Optimization problems may depend on a wide range of simulator model variables. The objective function gradient with respect to these variables is supplied to an external solver. Both the IPOPT (Wachter and Biegler, 2006) and SNOPT (Gill et al., 1997) solvers are implemented within the optimization module in AD-GPRS. The gradients are used within the solver routines, interior point algorithm for IPOPT, sequential quadratic programming for SNOPT, to find an optimum.

AD-GPRS is a research simulator that does not require a user license to run it, nor does it require an additional license to solve the system of equations in parallel. This means that the well placement algorithm within the optimization framework can launch a large number of reservoir simulations in a distributed manner without any type of license limitations. Still, the distributed implementation may be limited by other restrictions, such as server load, number of processors available, and, particularly in our case dealing with many reservoir simulations running in parallel, server hardware constraints on read-and-write access to disk (Chang and Moyer, 2010).

Table 3.1: Simulation cases used in transfer of model from Eclipse to AD-GPRS simulator. Each simulation case is described in terms of simulator, and the main approximations to simulation parameters and functions (these approximations are summarized in Table 3.2).

Case # and label	Simulator	Comp.	Aqu.	Gas lift	Controls	Rel. Perm.
1 ECLak_orig	Eclipse	Petrel	Anl.	Yes	Orig.	Hyst.
2 ECLbk_ovaq	Eclipse	Petrel	Vol.	Yes	Orig.	Hyst.
3 ECLck_rvqw	Eclipse	Petrel	Vol.	No	Disc.	Hyst.
4 ADGa_rqwdom	AD-GPRS	ExtM	Vol.	No	Disc.	Drain.
5 ADGa_rqwdfom	AD-GPRS	ExtM⁺	Vol.	No	Disc.	Drain.⁺
6 ADGc_rqwdfoc	AD-GPRS	Petrel	Vol.	No	Disc.	Drain.

Model transfer and validation procedure. The work model built from, and validated against, the original Martin Linge oil reservoir model has required several approximations to important model parameters and simulator functions. The approximations include modifications to aquifer support, saturation functions (e.g., relative permeability), well group control handling, gas lift optimization, and calculation of well completions. Roughly, the model transfer and validation process has consisted of finding suitable approximations to the effects these parameters and functions have during simulation. Several simulation cases were created to account for the effect these modifications have on final fluid flow. The configurations of these simulation cases are listed in Table 3.1. Table 3.2 summarizes the model parameters and functions modified in the transfer from Eclipse to AD-GPRS simulator.

The model transfer process starts from the original Eclipse reservoir model. Gradually, the different approximations are introduced into the field model until reaching a point where the Eclipse model can be transferred to the AD-GPRS simulator in a straightforward manner. The crucial element in the transfer process is the transplant of discretized well bottom-hole pressures and water production rates from an approximated Eclipse model to a replicate AD-GPRS model. Once production curves are reproduced in the AD-GPRS model, the model can be further tuned to produce curves that yield a better fit against the original fluid flow predictions. In particular, further tuning efforts involve modifications to relative permeability tables and calculations of well transmissibility factors. Below we discuss the final results from the validation effort without going into detail regarding the underlying transfer process. The simulation cases are therefore broadly defined. Further details regarding the different simulation cases, the approximations made to the original Eclipse reservoir model and other changes required by the transfer can be found in the next part of this section, Section 3.3.3 on page 75.

Validation results

The original Eclipse model is compared to different work model configurations for various production quantities. For further reference, mnemonics corresponding to the different quantities are listed in Table 3.3. The comparisons are spread over several fig-

Table 3.2: Description of model parameters modified in transfer of model from Eclipse to AD-GPRS simulator.

Parameter	Description
Simulator	Eclipse: Model run in Eclipse simulator. AD-GPRS: Model run in AD-GPRS simulator.
Completions (Comp.)	Petrel: Well connection factors computed by Petrel. ExtM: MRST code for well connection factor calculations developed for horizontal wellbores (see Section 3.2). ExtM⁺: ExtM well connection factors multiplied by constant to approximate Petrel well connection factors.
Aquifer (Aqu.)	Anl.: Aquifer modelled by Eclipse analytic aquifer functions. Vol.: Aquifer modelled by pore volume multiplication of boundary grid cells.
Gas lift	Yes: Lift gas injection is active for production wells. Injection rate optimized by Eclipse. No: No injection of lift gas.
Controls	Orig: Original well and well group controls including bottom-hole pressure and rate constraints (see Figure 6.7). Disc.: Modified well control setting (see Figure 6.8) using discretized well bottom-hole and water production rates from simulation case 2 (ECLbk_ovaq).
Relative permeability (Rel. Perm.)	Hyst.: Hysteresis is active in model, i.e. separate saturation function tables are used for drainage and imbibition. Drain.: Saturation function tables for drainage (water-oil and gas-oil) used for both drainage and imbibition. Drain.⁺: Irreducible gas saturation from imbibition gas-oil saturation table replaces irreducible gas saturation in gas-oil drainage saturation table. Modified gas-oil drainage table used for drainage and imbibition.

ures. Comparisons of main production quantities such as field reservoir pressure (Figure 3.7), well bottom-hole pressures (Figure 3.8), field production rates and totals (Figures 3.9 and 3.10), and well water cuts (Figure 3.11), can be found in this section. Other quantities secondary to the validation effort, such as well production rates and totals for gas (Figure 6.1), oil (Figure 6.2), water (Figure 6.3), and liquid (Figure 6.4), can be found in Appendix A (on page 209).

The primary concern of the validation effort is to obtain a reasonable agreement of to-

Table 3.3: Mnemonics corresponding to the production quantities used to compare the original Eclipse model and the different work model configurations.

Mnemonic	Description
WGPR / WGPT	Well gas production rate / total
FGPR / FGPT	Field gas production rate / total
WOPR / WOPT	Well oil production rate / total
FOPR / FOPT	Field oil production rate / total
WWPR / WWPT	Well water production rate / total
FWPR / FWPT	Field water production rate / total
WLPR / WLPT	Well liquid production rate / total
FLPR / FLPT	Field liquid production rate / total
WBHP	Well bottom-hole pressure
WWCT	Well water cut
FPRH	Reservoir pressure weighted by hydrocarbon pore volume

tal field production values, in particular oil recovery. Table 3.4 shows the normalized total field production values for gas, oil and water for the various simulation cases, including the original Eclipse field case model (ECLak_orig). Runtimes for the different simulation cases are also shown. Broadly, the simulation cases are configured as follows. The original Eclipse reservoir model is run in simulation case 1. Simulation case 2 modifies case 1 by adding the numerical aquifer approximation. Still using Eclipse, simulation case 3 tests the discretization of the bottom-hole pressure and water production rates obtained from simulation case 2. In simulation case 4, we have transferred the Eclipse model approximated in case 3 to AD-GPRS. Both approximations to saturation functions and our own calculation for completions with associated well connection factors are introduced here. Simulation case 5 tunes the relative permeability to better approximate the original hysteresis. Simulation case 6 tests our well completion calculations by running the AD-GPRS simulation using the original Eclipse well completion list. Detailed configurations for the different simulation cases and summaries of the individual approximations are given in Tables 3.1 and 3.2. Based on the simulation case results shown in Table 3.4 we select the work model configuration that offers both the best match in terms of production, and the best performance considering runtime and robustness.

Table results. The simulation case results presented in Table 3.4 show how the various model transfer approximations influence total production. From the table and corresponding production profiles, we see a clear advantage in obtaining discretized well pressure and production rates from an Eclipse simulation that already includes the aquifer approximation (i.e., from case 2). The interrelated forces from the main drive mechanisms, i.e., natural depletion, aquifer support and gas lift enhancement, are then all embedded

Table 3.4: Produced versus in place field volumes (in place volumes obtained from ECLak_orig)

Case # and label	Gas [-]	Oil [-]	Water [-]	Time [min]
1 ECLak_orig	1.000	1.000	1.000	302
2 ECLbk_ovaq	0.994	1.002	0.999	280
3 ECLck_rvqw	0.994	1.001	0.994	160
4 ADGa_rqwdom	1.164	0.941	1.001	259
5 ADGa_rqwdfom	1.003	0.989	1.000	338
6 ADGc_rqwdfoc	0.987	0.991	0.948	322

in the discretized quantities transferred to the AD-GPRS model. Due to their combined influence, it seems reasonable these forces have a dampening effect on the difference caused by replacing the analytic aquifer with a numerical aquifer approximation (see also page 76).

Aquifer approximation. In Table 3.4, we notice this dampening effect in that field production volumes only vary slightly when the aquifer approximation is added in case 2, even though the pressure situation in the reservoir has changed significantly, as can be seen in Figure 3.7 for the field pressure, and in Figure 3.8 for the well bottom-hole pressures.

Discretized pressures and rates. Results from Table 3.4 for simulation case 3 show that using the final well pressures and water rates after simulating with well and group control handling and gas lift optimization is a good approximation to these simulator functions (see page 78 for more on the gas lift function). We see that the three Eclipse cases are in good agreement for all types of production profiles.

Transfer to AD-GPRS. The transfer from Eclipse to AD-GPRS that occurs from case 3 to case 4 yields a marked increase and a noticeable drop in total produced volumes of gas and oil, respectively. We attribute these changes to only using the drainage relative permeability curves in our implementation since no hysteresis option is available in AD-GPRS (see page 76 for more on hysteresis). In Figure 3.10, we see that case 4 has a substantial gas over-production and a significant reduction in oil recovery.

Hysteresis fix. It is clear that the simpler relative permeability setup in the AD-GPRS model allows for a greater movement of free gas in the reservoir, than what was originally intended. We wanted to check whether this movement was the main factor contributing to the significant decrease in total field oil production. Following advice from the Total E&P Norge AS reservoir team, we increased the gas irreducible saturation in the drainage table to its analog value from the imbibition data. This modification was implemented in case 5. In Figures 3.9 and 3.10 we see that this modification keeps a larger amount of

previously free gas trapped in the rock, and results in a substantially improved match in production curves.

Final selection of work model configuration. Among the AD–GPRS cases, simulation case 5 yields the best match against the original field production data. Still, due to robustness considerations, we have chosen to implement the model configuration from case 4. The main reason is that the hysteresis fix introduced in simulation case 5 creates a non-smooth gas relative permeability curve which we observed increased numerical instability in the research simulator. Even for a simulation using base case wells (which we expect to be computationally less demanding than many trial solutions during well placement search) we noticed the hysteresis fix produced numerical difficulties that significant increased simulator runtime, as can be seen in Table 3.4 for this case (338 versus 259 minutes for case 5 and 4, respectively). Given our well placement and control routines will launch large numbers of model simulations configured with very different well trajectories and control schedules, these problems are likely to become worse, and may even cause simulator crashes. In this situation, we have chosen to prioritize simulation robustness over prediction accuracy. Thus, our final selection of base case work model configuration is based on simulation case 4 which has an approximately 6% lower total oil production and a roughly 16% higher total gas production compared to the original field model.




Figure 3.7: Field pressure over time. Suffix “b” on pressure mnemonics signifies values have been scaled for confidentiality reasons.

Second part: Main approximations to field model

The transfer of the Eclipse model to AD–GPRS implies that we approximate some features of the original Eclipse model, and some of the simulator functions used in the production strategy. In the following, we briefly describe each of these parameters, and their modification during model transfer.

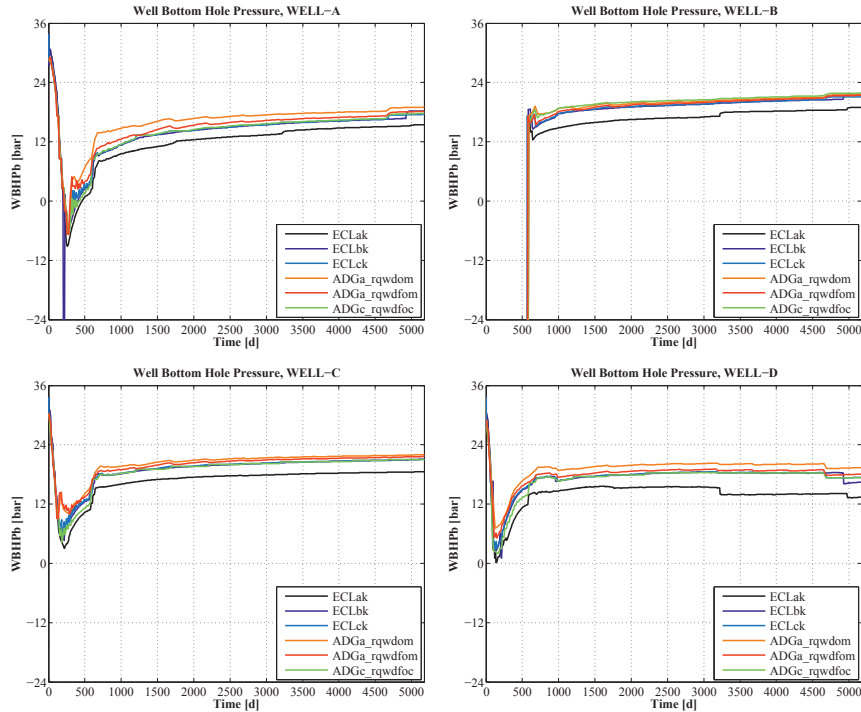


Figure 3.8: Well bottom-hole pressures. Suffix “b” on pressure mnemonics signifies values have been scaled for confidentiality reasons.

Aquifer. The original Eclipse model is implemented with an analytic aquifer formulation (Schlumberger, 2012b). This formulation calculates how aquifer pressure and influx varies over time depending on aquifer properties such as porosity, permeability and water viscosity. In our implementation, the analytic aquifer has been converted into a numerical aquifer description represented by a one-dimensional row of cells within the simulation grid (Schlumberger, 2012a). In this approximation, the pore volumes of the boundary row of cells that represent the aquifer is multiplied by a large constant to simulate constant pressure support and water influx.

Hysteresis. Fluids will usually have different adherence to the reservoir rock. The wettability of a reservoir rock is a measure that explains the preferential adherence of fluids to the rock (Dake, 1978). For example, in a water-wet oil reservoir, water will have a preferential adherence to the rock and cover the pore walls in a thin film. Water will then be termed the wetting phase for this reservoir rock. Drainage or imbibition are the processes where the saturation of the wetting phase of the rock is either decreased or increased, respectively. Because the rock usually prefers one phase over another, the displacement of a non-wetting phase, e.g., oil, by the wetting phase, e.g., water, will have different characteristics that when the displacement is in reverse order. This causes a hysteresis of the saturation functions (i.e., relative permeability and capillary pressure curves). In reser-

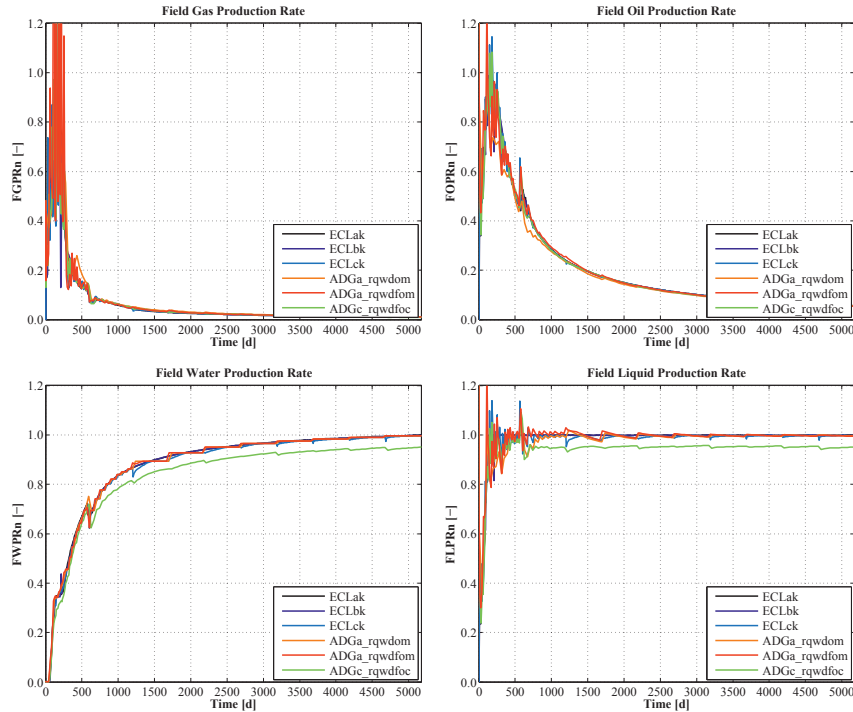


Figure 3.9: Field production rates: gas, oil, water, liquid. Suffix “n” on production mnemonics signifies values have been normalized for confidentiality reasons.

voir simulation, the hysteresis effect for, e.g., oil–gas displacements, is accounted for by having two sets of oil and gas relative permeability and capillary pressure tables (Schlumberger, 2012a), and correspondingly for water–oil displacements. In the original Eclipse model, hysteresis was only specified for oil–gas displacements. For water–oil displacements, equal sets of saturation function data were specified for drainage and imbibition processes. In our work, we approximated the oil–gas hysteresis effect by using only the drainage saturation function data to simulate for both the drainage and imbibition type of displacements.

Well group controls. Overall production management is usually guided by a set of constraints that function on groups of wells, on the entire field, or both. These field and group constraints come in addition to the constraints operating individually on each well, e.g., water and gas target production (i.e., rate) limits, and minimum bottom and tubing hole pressures. Specifically, group constraints function by imposing upper limits on field production quantities or on the total production of a selected group of wells (Schlumberger, 2012a). Any group production rate that surpasses the set limit will trigger an action rule in the simulator. The user selects the type of rules that will be activated if the group or field constraints are violated. For example, if an upper limit for group water production rate is exceeded, the simulator may scale back fluid production from wells

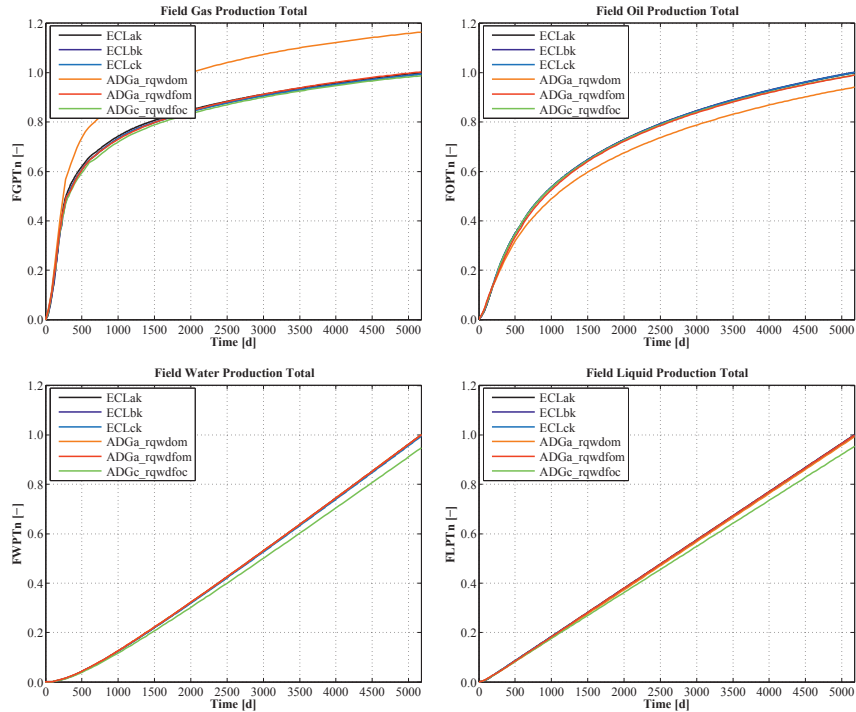


Figure 3.10: Field production totals: gas, oil, water, liquid. Suffix “n” on production mnemonics signifies values have been normalized for confidentiality reasons.

with high water cut, in order to keep group water production at or below the water rate limit. Other type of heuristics may also apply. Production control of groups of wells is currently not available in our AD-GPRS implementation. The effect of field-wide production constraints has been approximated by introducing into our implementation the resulting well bottom-hole pressures and production rates from the fully-constrained Eclipse simulation. These quantities were discretized at different intervals for the entire production time frame. In Appendix A (page 209) we present the different discretization periods (shown in figures 6.5 and 6.6), and illustrate how the control parametrization changes from the Eclipse to the AD-GPRS simulator (figures 6.7 and 6.8).

Gas lift. Production from wells for the development of the Martin Linge oil reservoir will be enhanced by artificial lift. Artificial gas lift is an advanced production technique that injects gas into production fluids at bottom-hole level. Lift gas will mix and lower the density of fluids so that a lower pressure differential between surface and reservoir is needed to bring the fluids to the surface. At each time step during simulation, a gas lift allocation routine determines an adequate lift gas injection rate for each well. Specifically, the gas lift routine determines how much lift gas to allocate to each well in order to meet well, group or field production targets (Schlumberger, 2012b). Among other constraints (such as compressor capacity), the gas lift allocation procedure is bounded by the mini-




Figure 3.11: Well water cuts.

imum tubing head pressure that needs to be present for proper operation of the platform separators. Enhanced production by lift gas injection is not an available production option in our current AD-GPRS implementation. This production technique, similar to field and group production controls, has been approximated by bottom-hole pressures and water production rates taken from an original Eclipse simulation and inserted into our work model.

Well completions. Base case wells in the original Eclipse model are built using the Petrel E&P Software Platform developed by Schlumberger (2012c). Petrel allows the user to design arbitrary well trajectories on the grid, and to specify well completions such as type and diameter of tubing, and perforations. Once well details are specified, Petrel uses grid information, e.g., block geometry and permeability, to convert the wells into a well completion list. Importantly, the well completion list is the actual representation of the wells within the simulator model. Each well is represented by a set of completions with a corresponding well transmissibility factor (also known as well index, or well connection factor; Schlumberger, 2012b). The transmissibility factor determines the fluid influx through the completion and into the well given the existing pressure differential between reservoir and well-bore. In our implementation, we use an external code developed from the MRST (Lie et al., 2012) software as an approximation to the Petrel computation of

well transmissibility factors. This external piece of code uses grid information from the original Eclipse model read in by MRST functions into the MATLAB platform. During optimization, this code is called by the well placement algorithm to find the well blocks and compute the corresponding well transmissibility factors for each of the trial well placement solutions.

3.4 Optimization work

Section introduction. This section and the next deal with methodology and results. This section starts with a brief recap of some of the work up to now, and then presents the overall problem formulation, and application of methodology. The next section, Section 3.5, presents the main results from our application, as well as discussions and suggestions for further work. The solutions developed in Section 3.5 are tested on the original field case model in the next chapter, Chapter 4. The optimization work presented in this and the next section is the result of the procedural work conducted in the first three sections of this chapter. Those sections dealt mostly with work process issues such as strategy and framework, and model validation work. Here we focus on application and solutions.

3.4.1 Introduction to optimization work

We start this section by giving brief summaries of issues that have particular relevance to the optimization work that will be presented in this and the next section. These issues have been treated in Chapter 2 and in the previous three sections in this chapter. We summarize these topics here with the intention of updating the foundation for the optimization work ahead. The issues discussed below are: joint versus sequential approach, recap of the reservoir simulation model, transfer of field case model to research simulator, embedded control optimization, development of work model for optimization, production time frame approximation, and finally, challenges and collaboration work.

Joint versus sequential approach. In this work, we wish to realize some of the gain inherent in the complexity of the well placement decision. To reach this end, our approach requires the search for optimal well placement to take into account its dependency on optimal production controls. In Chapter 2 we described a joint approach where the well placement and control problems are solved in an integrated manner. In that work we showed that, in terms of cost function value, the joint approach outperformed sequential approaches by almost 20%. For the optimization work ahead we compare the joint approach against a sequential procedure when optimizing the location and production of several horizontal wells using a real field model.

Field case reservoir simulation model. The previous section, Section 3.3, introduced the field case model and discussed related validation issues at length. Here we show the initial saturations and recap some of the main features of the simulation case. The initial gas (red), oil (green) and water (blue) saturations for the particular realization provided to us are shown in Figure 3.12. We reiterate the model for the Martin Linge oil reservoir consists of approximately 55000 active grid cells, and has a base case production strategy




Figure 3.12: Initial saturations of gas (red), oil (green) and water (blue) at the Martin Linge oil reservoir model.

based on production from four horizontal wells. Tables 4.1 and 4.2 in Chapter 4 provide specific details about the simulation case.

Transfer of field case model to research simulator. The implementation of the field case model using the AD-GPRS simulator enables us to take full advantage of the extensive capacity for parallelization inherent in the pattern search algorithm used in this work. In this manner, our application can perform a large number of cost function calls (in the order of 50) in a distributed manner without any type of license limitations usually imposed on commercial software.

Embedded control optimization. Given the integrated solution setup of the optimization framework devised for this work, an efficient optimization of controls is particularly important. The optimization module included in the AD-GPRS simulator provides an efficient optimization of well controls through an adjoint-based computation of gradients. This optimization relies on the effective computation of objective function gradients using a discrete adjoint formulation relying on automatic differentiation (Volkov and Kourounis, 2012).

Development of work model for optimization. Simulator functions for well group control handling and gas lift rate allocation procedure are not currently present in AD-GPRS. The operation of these functions was approximated by using the well bottom-hole pressures (WBHP) and water production rates (WWPR) from a simulation of the Eclipse model as initial well controls and target/limit rates in our AD-GPRS implementation. The work model implementation in AD-GPRS, including all the above-mentioned approximations, yields about 6% less total oil production and roughly 16% higher total gas pro-

duction compared to the Eclipse field model⁵.

Production time frame approximation. For this particular reservoir, the field and well oil production rates are seen as likely to peak shortly after a couple of years, and are not expected to have a substantial plateau. Because a large part of the oil production is predicted to occur in the first few years of production, the work model production horizon has been reduced to about a fourth of the original production time frame.

Application challenges. The complexity of field development operations, in addition to the significant computational demand of field models, will often require that the problem definition and scope exclude various factors otherwise considered important to the overall development of the field⁶. From an applied research perspective, to achieve a reasonable problem definition and scope, it is important to resolve the task of gaining a clear understanding of which factors that should be taken into account. Resolving this task is essential because it will ultimately determine how the research methodology is to be developed towards a field application.

Collaboration work. Through the IO Center research network we established a robust collaboration with field operator Total E&P Norge AS. Also, working with IO Center Academic and Research Partners, we have developed a practical problem definition, re-implemented the field case reservoir model, and extended the previous implementation to deal with the significantly more challenging problem of testing our methodology on a field case.

Next we provide the concrete problem formulation used for optimization, and a description of the the non-linear constraints implemented for the well placement part of the procedure.

3.4.2 Problem formulation

In this section we provide the field case well placement and control optimization problem formulation.

Optimization problem

In this work we apply our optimization procedure to find improved locations and controls for four production wells. At the time the field case model was provided to us, these wells corresponded to one base case solution created for the development of the Martin Linge oil reservoir. The optimization problem presented next treats these wells as horizontal wells subject to non-linear constraints including maximum well-length and minimum inter-well distance. The optimization problem studied here is defined as follows:

$$\min_{\mathbf{x}_p \in \mathbb{R}^n, \mathbf{x}_c \in \mathbb{R}^m} -\text{FOPT}(\mathbf{x}_p, \mathbf{x}_c) \quad \text{subject to} \quad \begin{cases} \mathbf{c}_{wl}(\mathbf{x}_p) & \leq \mathbf{l}_{max} \\ \mathbf{c}_{rb}(\mathbf{x}_p) & \leq 0 \\ \mathbf{c}_{wd}(\mathbf{x}_p) & \geq d_{min} \\ \mathbf{x}_c^d & \leq \mathbf{x}_c \leq \mathbf{x}_c^u \end{cases}, \quad (3.1)$$

where \mathbf{x}_p and \mathbf{x}_c represent well placement coordinates and well control variables, respectively. Well placement variables are denoted by $\mathbf{x}_p \in \mathbb{R}^n$, where $n = 6 \cdot N_w$ with N_w

being the number of horizontal wells in the optimization (for each well we need to determine six coordinates, three for the heel and three for the toe). Thus, $N_w = 4$ wells yields 24 well placement variables. $c_{wl}(\mathbf{x}_p)$, $c_{rb}(\mathbf{x}_p)$, and $c_{wd}(\mathbf{x}_p)$ are non-linear constraints on well-length, reservoir bounds, and inter-well distance. The implementation of these constraints is fully explained in the next section.

In this work, we optimize the controls for only the second half of the total production time frame of 1200 days designated for optimization. Since the volumes of free gas initially in place for this development are considered to be relatively small, a main target in the established strategy for this asset is to produce most of the free gas during early production. In accordance with this field development target, in our application we do not optimize for controls during the first 600 days of production, when most of gas production occurs, but rather implement the well control schedule preset by the operator (i.e., obtained from the Eclipse model, see Section 3.3) during this time period. At the same time, this configuration alleviates several computational concerns involving high gas rates in the first phase of production. For example, during early implementation tests, we observed that very high gas rates caused difficulties in the convergence of the solutions for the fluid flow equations, a situation that resulted in substantial simulation runtime increases. At these early production times, this type of behavior was observed already in the Eclipse reservoir simulations, and was seen to become markedly worse in the AD-GPRS implementation. From subsequent tests, where optimization of controls was also performed within this early time period (testing both the base case and other well placement configurations), we furthermore observed several reservoir simulations for various trial solutions either stagnating due to very short solver time steps or straightaway crashing. Thus, by not optimizing controls at early production times, we also avoided the introduction of this potential instability to the overall optimization process which involves running simulations for a whole range of different well placement and control trial solutions.

Well control variables are denoted by $\mathbf{x}_c \in \mathbb{R}^m$, where $m = 2 \cdot N_t \cdot N_w$ with N_t being the number of time intervals in the piecewise constant function over time that represents the controls for each well. The piecewise constant function for each well between 600 to 1200 days contains $N_t = 8$ time intervals. As mentioned before, our work model approximation was validated by implementing the well bottom-hole pressures (WBHPs) and well water production rates (WWPRs) from an Eclipse field case model simulation as well pressures/target rate settings in the AD-GPRS work model. This means that, during simulation, each well in the AD-GPRS work model is operated using either its WBHP or WWPR control setting. Recall that, from one simulation time step to another, the operation of a well may switch from WBHP to WWPR control if water production for that well has reached the water rate limit specified for that well at that time period. Otherwise, the well will remain under (or possibly switch back to) WBHP control if, and as long as, the well water rate is below the given limit.

In the well control optimization procedure implemented in this work, both WBHP and WWPR settings are used as control optimization variables defined over the 600 to 1200 day production horizon. Within this time frame, each well is represented by N_t control variables corresponding to WBHP, and N_t control variables corresponding to WWPR. In our implementation, initial WBHP and WWPR values for all wells are the

same as those used for the validation effort. Using both WBHP and WWPR as variables for control optimization is an approximate way of optimizing for controls in the presence of simulator-imposed production constraints. The idea with this treatment is not to lose any gradient sensitivity whenever the control shifting occurs. During the optimization of controls, \mathbf{x}_c is therefore only subject to bound constraints. See Kourounis et al. (2014) for a detailed discussion and comparison of formal and heuristic approaches for how to impose non-linear constraints on output rates and/or pressures during control optimization. The heuristic treatment the authors describe in that work could be an attractive alternative for future optimization of controls in our case. The main reason is that the heuristic treatment they propose only requires one additional reservoir simulation to enforce the non-linear constraints after an optimization using only bound constraints has been performed. Similar treatments have also been tested by Møyner et al. (2014). Though optimizing for both WBHP and WWPR doubles the number of control variables, this number is still relatively low (less than 100), and does not incur any significant cost for the highly efficient, adjoint-based, computation of gradients. Having $N_w = 4$ wells, yields a total of $m = 2 \cdot 8 \cdot 4 = 64$ well control variables.

We define our objective function as the cumulative oil produced (or field oil production total; FOPT) for the 1200 day production time frame. The FOPT is defined as follows:

$$\text{FOPT}(\mathbf{x}_p, \mathbf{x}_c) = \sum_{k=1}^{N_s} \left(\sum_{j=1}^{N_w} q_o^{j,k}(\mathbf{x}_p, \mathbf{x}_c) \Delta t_k \right), \quad (3.2)$$

where $q_o^{j,k}$ is the oil rate for well j at the output interval k , and Δt_k represents the length (in days) of each of the N_s time steps in the simulation. In the following we describe the different constraints for the well placement part of the problem.

3.4.3 Methodology

In this section we describe the non-linear constraints implemented for the well placement part of the optimization. These constraints are based on projecting the well placement coordinate vector onto the feasible space. At the end, we use pseudo-code to offer a concise description of the implementation of the approaches. We provide two sets of pseudo-code that explain the progression of the procedure with and without the implementation of the well placement constraints.

Non-linear constraint handling

In this section we define the projection operator P to describe the handling of non-linear constraints for well placement coordinates in our application. We first define P and the constituent non-linear constraints, and then describe how P is approximated as an iterative sequence of projection operators. The different implementations of the well-length constraint, and the application of the well-distance and reservoir-bound constraints are then described in detail.

Projection operator P . Below we introduce the operator $P(\cdot)$ that includes all projections performed on well placement coordinates during optimization. P applies to well

heel and toe coordinates that are out of specified bounds, and enforces inter-well distance and well-length constraints. The projections are applied to each well coordinate vector \mathbf{x}_p^0 generated by the well placement part of the optimization procedure. Resulting feasible well heel and toe coordinates are presented by \mathbf{x}_p . We reiterate that the computation of $P(\mathbf{x}_p^0)$ does not require a reservoir simulation, and has negligible cost when compared to the evaluation of the cost function. Other constraints of this type (i.e., not involving reservoir flow simulation) may be added to P though solving for feasible well configurations may be more demanding. $P(\mathbf{x}_p^0)$ is given as:

$$\begin{aligned}
 P(\mathbf{x}_p^0) &= \underset{\mathbf{x}_p \in \mathbb{R}^n}{\operatorname{argmin}} \|\mathbf{x}_p - \mathbf{x}_p^0\|^2 & (3.3) \\
 &\text{subject to} \\
 &\mathbf{c}_{wl}(\mathbf{x}_p) \leq \mathbf{l}_{max} \\
 &\mathbf{c}_{rb}(\mathbf{x}_p) \leq 0 \\
 &\mathbf{c}_{wd}(\mathbf{x}_p) \geq d_{min},
 \end{aligned}$$

where $\mathbf{c}_{wl}(\mathbf{x}_p) \in \mathbb{R}^{N_w}$ and $\mathbf{c}_{rb}(\mathbf{x}_p) \in \mathbb{R}^{2 \cdot N_w}$ represent the well-length constraint and the heel-toe bound constraint, respectively. The heel-toe bound constraint defines a feasible area for each heel and toe independently. These areas are shown in Figure 3.13; red and blue areas bound heel and toes, respectively. Together, these bounds ensure that wells have a reasonable alignment with respect to the platform. The inter-well distance constraint is given by $\mathbf{c}_{wd}(\mathbf{x}_p) \in \mathbb{R}^{N_w}$. Here $\mathbf{l}_{max} \in \mathbb{R}^{N_w}$ refers to the maximum lengths for each horizontal well, and d_{min} specifies the minimum distance between any two well trajectories.

Iterative sequence of projection operators. In this work, the projection operator P is implemented as an iteratively sequence of three projection operators, each of them related to the three types of constraints in (3.3). The iterative implementation of constraints is partly a by-product of the collaboration work with the industry operator. In this regard, code development progressed in stages. At each stage, we would present preliminary results and obtain feedback, in this particular case during problem definition, with respect to which constraints that were important to implement within the optimization. For example, the well-length projection was added to an existing implementation that only included the projections associated with the heel and toe bounds, and the minimum distance between wells. Subsequent code development, such as the addition of the well-length constraint, was then built on top of the existing implementation. Here we denote the projections associated with $\mathbf{c}_{wl}(\mathbf{x}_p)$, $\mathbf{c}_{rb}(\mathbf{x}_p)$, and $\mathbf{c}_{wd}(\mathbf{x}_p)$, as P_{wl} , P_{rb} and P_{wd} , respectively. In our implementation, P_{wl} is solved first, then P_{rb} and P_{wd} . This sequence of operators is repeated if the projected coordinates for all wells violate any of the constraints (although the iterative process is not guaranteed to yield a feasible solution, in our application, constraint satisfaction was achieved in all cases).

Implementations of well-length constraint. Additionally, we have applied three different implementations of P_{wl} that account for reasonable engineering techniques to deal with the maximum well-length constraint. These implementations are denoted here by




Figure 3.13: Heel and toe circular constraint regions implemented for each of the four wells in this work.

P_{wl}^A , P_{wl}^B and P_{wl}^C . Our first implementation, P_{wl}^A , solves for P_{wl} using the norm of $\mathbf{x}_p - \mathbf{x}_p^0$ as cost function. P_{wl}^A is given by:

$$P_{wl}^A(\mathbf{x}_p^0) = \operatorname{argmin}_{\mathbf{x}_p \in \mathbb{R}^n} \|\mathbf{x}_p - \mathbf{x}_p^0\| \quad \text{s.t.} \quad \mathbf{l}_w(\mathbf{x}_p) \leq \mathbf{l}_{max}, \quad (3.4)$$

where $\mathbf{l}_w(\mathbf{x}_p) \in \mathbb{R}^{N_w}$ refers to the length of each horizontal well. The second implementation, P_{wl}^B , approximates P_{wl} by finding feasible well configurations within the subspace of well lengths. P_{wl}^B is given as:

$$P_{wl}^B(\mathbf{x}_p^0) = \operatorname{argmin}_{\mathbf{x}_p \in \mathbb{R}^n} \|\mathbf{l}_w(\mathbf{x}_p) - \mathbf{l}_w(\mathbf{x}_p^0)\| \quad \text{s.t.} \quad \mathbf{l}_w(\mathbf{x}_p) \leq \mathbf{l}_{max}. \quad (3.5)$$

The third implementation, P_{wl}^C , is an approximation of P_{wl} where the maximum well length constraint from (3.4) and (3.5) is applied in a straightforward manner. In the implementation of P_{wl}^C , any wells longer than the maximum length are set equal to \mathbf{l}_{max} by moving the toes closer to the heels while keeping the heel positions fixed. As with the other projections, the P_{wl}^C implementation is applied iteratively together with the well-distance and bound constraints.

In a special configuration, no well length implementation is applied during the optimization procedure. Rather, for this particular configuration, P_{wl}^C is applied once after

the optimization routine has arrived at a solution, i.e., maximum well lengths are enforced only on the final well configuration, such that $\mathbf{x}_p^* = P_{wd}^C(\mathbf{x}_p^*)$. (When using this configuration, the a posteriori enforcement of the well-length constraint did not affect the satisfaction of the other constraints in our implementation.)

Well-distance constraint. We approximate the distance between any two wells by computing the minimum distance between a well heel or toe, and the other well trajectory. (Future code development will compute the minimum distance between any two points along each of the well trajectories.) The implemented projection P_{wd} is illustrated in Figure 3.14 as follows. If the distance from the toe/heel to a neighboring trajectory is smaller than the minimum inter-well distance d_{min} , then the toe is perpendicularly moved away from that other trajectory (i.e., along the red line in Figure 3.14) until the distance becomes equal to d_{min} . This operation is performed iteratively for all wells. The iterative process continues until all inter-well distances are larger or equal to d_{min} . (This process is set to stop after a maximum number of iterations. Again, this procedure is not guaranteed to converge, though, in our implementation, all well coordinates achieved feasibility before the given number of maximum iterations was reached, i.e., all solutions are feasible.)




Figure 3.14: Example of P_{wd} applied to toe of well B. P_{wd} moves the toe along the direction perpendicular to the well A trajectory (red line), until the distance between toe and trajectory is equal to d_{min} .

Reservoir-bound constraint. The third projection approximation, P_{rb} , is based on not allowing the heel and toe to move outside given circular regions (not necessarily centered at the heel and toe) which contain the heel and toe coordinates for the initial configuration \mathbf{x}_p^0 . If a given solution for the heel and toe coordinates lies outside the circular region, the operator projects heel and toe coordinates onto their respective bounds. This operation is exemplified in Figure 3.15 for an infeasible well toe. The actual heel and toe bound areas for each of the four wells in our implementation are shown in Figure 3.13. These bounds were designed in collaboration with the field operator such that the optimization routine would generate solutions with a realistic heel-toe orientation with respect to the planned location of the platform. In our model, the reservoir characteristics and shale distribution yield an irregular reservoir boundary with many grid blocks inactive. The projection shown in Figure 3.15 takes this into account and only projects infeasible heels or toes onto active grid blocks inside the corresponding bound area.




Figure 3.15: Example of projection of infeasible well toe coordinate onto corresponding feasible bound area.

Pseudo-code for optimization procedure. This section describes the optimization procedure applied in this work using pseudo-code presentations. In Algorithm 2 we describe the implementation of the joint and sequential fixed approaches without non-linear constraint handling. Algorithm 3 expands the routine presented in Algorithm 2 by adding the sequence of projection operators that handle non-linear constraints on the well placement variables. Recall that the addition of non-linear constraints to treat well coordinates is computationally inexpensive since the enforcement of these constraints does not require any reservoir flow simulations. For computationally expensive constraints one can use approaches such as the filter method (Echeverría Ciaurri et al., 2011a). For clarity, on some occasions, the well position and control variables \mathbf{x}_p and \mathbf{x}_c are denoted together as \mathbf{x} in algorithms 2 and 3. We end this section with brief descriptions of the solution processes represented by these two algorithms.

As mentioned before, Algorithm 2 describes the overall solution procedure of the joint and sequential fixed approaches without the non-linear constraint implementation. As such, the differences between Algorithm 2 and 3 represent the algorithmic development of the work from Chapter 2 to the current Chapter 3 implementation. In Algorithm 2, the well placement part of the procedure is represented by the *while* loop operating between lines 1 and 7. The body of this *while* loop contains the conditional statement that either performs a control optimization at the given well placement iterate \mathbf{x}_p^i , i.e., solves for

$$\mathbf{x}_c^* = \operatorname{argmin}_{\mathbf{x}_c \in R^m} - \operatorname{FOPT}(\mathbf{x}_p^i, \mathbf{x}_c),$$

or uses the fixed control settings \mathbf{x}_c^0 to compute $\operatorname{FOPT}(\mathbf{x}_p^i, \mathbf{x}_c^0)$. Obviously, actuating the first alternative of this statement means we are implementing the joint approach, while engaging the second alternative during the well placement search constitutes the implementation of the sequential approach. (The conditional treatment when using a reactive control strategy is homologous to the one described here for fixed controls.) For the sequential approach, lines 8 to 9 describe the additional optimization of controls performed after a solution has been obtained from the well placement routine.

Algorithm 3 details the progression of the optimization procedure subject to the well-

length, well–distance and reservoir–bounds constraints. Recall that in this algorithm, the iterative sequence of projection operators is represented by $P(\mathbf{x}_p^i)$. Which implementation of the well–length constraint that is applied during optimization is given before the procedure starts. In Algorithm 3, the conditional statement previously given in lines 1 to 7 in Algorithm 2 is now present between lines 6 and 10. Now, before entering this statement, we have our projection operator P acting upon the well placement iterate \mathbf{x}_p^i . This activity is described by line 3 of Algorithm 3,

$$\mathbf{x}_p^i = P(\mathbf{x}_p^i),$$

which represents the iterative sequence of projection operators applied onto the well coordinates in \mathbf{x}_p^i advanced by the well placement procedure (functioning from line 1). Line 3 is applied to each well coordinate iterate \mathbf{x}_p^i until all constraints are satisfied, or a maximum number of loops has been reached.

Next, in Section 3.5, we present the results obtained from the application of Algorithm 3 when using the work model developed for the Martin Linge oil reservoir.

Algorithm 2 Optimization procedure for both joint and sequential approach. Well position and control parts of iterative trial solution given as \mathbf{x}_p^i and \mathbf{x}_c^i , respectively.

Require: Specify whether running joint or sequential approach, provide initial well position and controls: $\mathbf{x}^0 = (\mathbf{x}_p^0, \mathbf{x}_c^0)$.

Ensure: Improved solution for well position and control: $\mathbf{x}^* = (\mathbf{x}_p^*, \mathbf{x}_c^*)$.

{Optimal controls at position iterate \mathbf{x}_p^i obtained by solving
 $\mathbf{x}_c^* = \operatorname{argmin}_{\mathbf{x}_c \in R^m} - \text{FOPT}(\mathbf{x}_p^i, \mathbf{x}_c)$ using the SNOPT solver.}

{Embedded optimization:}

- 1: **while** searching for \mathbf{x}^* **do**
- 2: **if** joint approach **then**
- 3: solve for \mathbf{x}_c^* at \mathbf{x}_p^i , use FOPT $(\mathbf{x}_p^i, \mathbf{x}_c^*)$
- 4: **else if** sequential approach **then**
- 5: compute FOPT $(\mathbf{x}_p^i, \mathbf{x}_c^0)$
- 6: **end if**
- 7: **end while**

{Additional optimization of controls:}

- 8: **if** sequential approach **then**
 - 9: solve for \mathbf{x}_c^* at \mathbf{x}_p^* , use FOPT $(\mathbf{x}_p^*, \mathbf{x}_c^*)$
 - 10: **end if**
-

Algorithm 3 Optimization procedure for both joint and sequential approach subject to non-linear constraints on well placement variables. Well position and control parts of iterative trial solution given as \mathbf{x}_p^i and \mathbf{x}_c^i , respectively. The iterative sequence of projection operators for non-linear constraint handling is represented by $P(\mathbf{x}_p^i)$.

Require: Specify whether running joint or sequential approach, provide initial well position and controls: $\mathbf{x}^0 = (\mathbf{x}_p^0, \mathbf{x}_c^0)$, specify choice of well-length constraint during optimization: $P_{wl} \in \{P_{wl}^A, P_{wl}^B, P_{wl}^C\}$, or only P_{wl}^C applied at end of procedure.

Ensure: Improved solution for well position and control: $\mathbf{x}^* = (\mathbf{x}_p^*, \mathbf{x}_c^*)$, \mathbf{x}^* satisfies all non-linear constraints given in (3.3).

{Optimal controls at position iterate \mathbf{x}_p^i obtained by solving
 $\mathbf{x}_c^* = \underset{\mathbf{x}_c \in R^m}{\operatorname{argmin}} - \operatorname{FOPT}(\mathbf{x}_p^i, \mathbf{x}_c)$ using the SNOPT solver.}

{Embedded optimization with non-linear constraints:}

- 1: **while** searching for \mathbf{x}^* **do**
- 2: **for** well coordinate iterate \mathbf{x}_p^i **do**
- 3: $\mathbf{x}_p^i = P(\mathbf{x}_p^i)$
- 4: until all constraints satisfied or maximum number of loops reached
- 5: **end for**
- 6: **if** joint approach **then**
- 7: solve for \mathbf{x}_c^* at \mathbf{x}_p^i , use $\operatorname{FOPT}(\mathbf{x}_p^i, \mathbf{x}_c^*)$
- 8: **else if** sequential approach **then**
- 9: compute $\operatorname{FOPT}(\mathbf{x}_p^i, \mathbf{x}_c^0)$
- 10: **end if**
- 11: **end while**

{Additional optimization of controls:}

- 12: **if** sequential approach **then**
 - 13: solve for \mathbf{x}_c^* at \mathbf{x}_p^* , use $\operatorname{FOPT}(\mathbf{x}_p^*, \mathbf{x}_c^*)$
 - 14: **end if**
-

3.5 Optimization results

This section describes the results obtained from the application of the optimization procedure. It starts with a description of the different optimization runs performed with the AD-GPRS work model (recall this model uses the 1200 day production time frame). In total, nine different joint and sequential optimization runs have been performed using different configurations for the well-length constraint. One table and two sets of function evolution graphs are used to present the main results from these optimization runs. The solutions are described in terms of their final objective function value and performance. Furthermore, comparisons are made between the results obtained using the joint and the sequential approaches. We then implement the solutions obtained using the work model on the Eclipse field case model. The results from this transfer of solutions are presented in a table where both the entire solution (\mathbf{x}^*), i.e., the well placement (\mathbf{x}_p^*) and the well control part (\mathbf{x}_c^*), and then only the well placement part, are tested on the field case model. We select the best-performing solution from the transfer table, in terms of final objective function value, and plot the saturation maps for this solution at different times using the commercial simulator.

3.5.1 Optimization runs

Our optimization procedure has developed a total of nine well placement and control solutions using the AD-GPRS work model. Table 3.5 shows the final objective function values for these solutions, in addition to the corresponding well lengths. Function evolution graphs for sequential and joint solutions are presented in figures 3.16 and 3.17 (these figures will be described in further detail at a later point). The graphs presented in these figures correspond to joint and sequential optimization runs using different configurations of projection operators. That is, these runs are the result of launching different configurations of the optimization procedure described in Algorithm 3 (page 90). Of the nine solutions, four are obtained using the sequential approach, while five solutions were developed using the joint approach. Joint and sequential runs are denoted by names "JNT" and "FXD", respectively. Additionally, the name for each individual solution indicates which configuration of projection operators (see Section 3.4.3: Non-linear constraint handling, page 84) that was applied during the optimization. While the use of well-bound and inter-well distance operators P_{rb} and P_{wd} is the same for all solutions, there are four possible implementations for the well-length operator P_{wl} : "OPT2", "OPT", and "CUT" correspond to P_{wl}^A , P_{wl}^B , and P_{wl}^C being applied during the optimization procedure, respectively, while "M1" means that a P_{wl}^C projection (with an added lower bound) is applied only once, at the end of the optimization iteration.

Description of results shown in Table 3.5. Table 3.5 shows the final objective function values (FOPT), total number of reservoir simulations (nsims), and corresponding well lengths for each of the nine solutions. Well lengths for all solutions are shown in columns 4 to 7. Solution values associated with the sequential and joint approaches are shown in the upper and lower half of the table, respectively. For comparison, values associated with the initial well configuration, referred to as "BASECASE", are given in the first row in Table 3.5. All objective function values in this table are normalized rel-

Table 3.5: FOPT and well lengths corresponding to the well placement solutions for joint and sequential runs. The total number of reservoir simulations run by at each optimization is also given. The upper level title “AD-GPRS₁₂₀₀” refers to the fact that the results in this table are obtained using the AD-GPRS work model approximation, run using the 1200 day production time frame. Initial base case values for FOPT and well lengths are provided for comparison.

Solution	AD-GPRS ₁₂₀₀						
	FOPT		nsims	Well lengths [m]			
	$(\mathbf{x}_p^*, \mathbf{x}_c^0)$	$(\mathbf{x}_p^*, \mathbf{x}_c^*)$		WL.-A	WL.-B	WL.-C	WL.-D
BASECASE	1.000	1.011	-	1439	1247	1409	874
FXD1M1	1.193	1.208	708	1500	1500	1500	1200
FXD2OPT	1.298	1.315	1144	1364	1139	1299	1358
FXD2OPT2	1.268	1.272	1343	1129	1202	1295	1362
FXD2CUT	1.225	1.234	905	1129	1202	1373	1069
JNT2M1	-	1.334	5692	1500	1500	1500	1200
JNT1M1	-	1.308	2873	1500	1500	1500	1200
JNT2OPT	-	1.329	9727	1088	1474	1299	1316
JNT2OPT2	-	1.307	9633	1129	1142	1302	1155
JNT2CUT	-	1.294	6798	1129	1034	1299	1271

ative to the initial FOPT value for BASECASE, i.e., $\text{FOPT}(\mathbf{x}_p^0, \mathbf{x}_c^0)$. For the sequential approach, we also provide the objective function values before the final optimization of well controls, i.e., $\text{FOPT}(\mathbf{x}_p^*, \mathbf{x}_c^0)$, and after control optimization, i.e., $\text{FOPT}(\mathbf{x}_p^*, \mathbf{x}_c^*)$. Note that runs JNT2M1, JNT1M1 and FXD1M1 were developed without restrictions on well length during optimization. However, once these iterations finished, the well lengths corresponding to these solutions were projected onto upper and lower bounds of 1500 and 1200 meters. The objective function values for these solutions were then recalculated using the new well lengths. In the following, we will compare the joint and sequential runs. (The JNT1M1 run is not included in this comparison because this solution is an alternative version of the JNT2M1 run, which will be discussed separately at a later point.)

Joint vs. sequential approach: Comparison of mean FOPT values from Table 3.5.

For the sequential runs, the second column values, i.e., $\text{FOPT}(\mathbf{x}_p^*, \mathbf{x}_c^*)$, in the upper half of Table 3.5, result from an optimization of controls performed after a well placement configuration has been found using fixed controls. Taking the mean of these values, and comparing them to the initial well placement and control configuration (first-column FOPT value corresponding to BASECASE), we obtain a mean increase in FOPT of almost 26% for the sequential runs. Continuing towards the bottom half of Table 3.5, we see that the corresponding joint solutions have higher FOPT increases than their sequential counterparts. The $\text{FOPT}(\mathbf{x}_p^*, \mathbf{x}_c^*)$ values for the joint solutions result in a mean increase in FOPT of close to 33% over the initial configuration. In summary, we have that for this particular problem case, the solutions obtained using the joint approach yield, on average, a 7% higher FOPT increase compared to the solutions obtained using the sequential approach.

Influence of control optimization in each approach. For each of the sequential optimization runs, the increase due to the final optimization of controls is found by comparing the values in the $\text{FOPT}(\mathbf{x}_p^*, \mathbf{x}_c^0)$ column against those in the $\text{FOPT}(\mathbf{x}_p^*, \mathbf{x}_c^*)$ column. From Table 3.5 we have that the mean increase these runs receive due to their sole optimization of controls at the end is somewhat above 1% (Roughly, the increases are 1.5% for FXD1M1, 1.7% for FXD2OPT, 0.4% for FXD2OPT2 and 0.9% for FXD2CUT.) These increases are comparable to the increase obtained when optimizing the controls using the initial well configuration (which is 1.1%, see first row). For this particular problem case then, we have that the yields from control optimization when using the initial well configuration, or any of the final well configurations obtained from the sequential runs, are modest. Interestingly, from function evolution data for the joint runs, we observe that the control routine when embedded within the well placement search, yields, on average, considerable higher increases when applied over whole ranges of different well placement trial solutions. For all the trial solutions in each of the joint runs, Table 3.6 shows the mean increases in FOPT associated with the control optimization part only. Corresponding standard deviations are also given. Roughly, the data show that when using the joint approach, each well placement trial solution obtains a general improvement in objective function value due to the inner-loop control procedure of, on average, 7.5%. This could be an indication that the control routine is compensating for less promising well locations, as also noted for the much simpler case in Section 2.4.1 (page 26). A final note is that the mean FOPT increase due to the embedded optimization of controls is maintained throughout the entire well placement search for all joint runs. However, we see from the relatively large standard deviations for all runs in Table 3.6 that the magnitude

of the contributions is unevenly distributed among the trial solutions.

Table 3.6: Mean increases in FOPT resulting from embedded control optimization during joint runs. $\langle \Delta \text{FOPT} \rangle_{ctrl}^{tsolns}$ represents the mean FOPT increase associated with the control optimization part only, averaged over all trial solutions (*tsolns*) that are performed during a joint run.

Solution	$\langle \Delta \text{FOPT} \rangle_{ctrl}^{tsolns}$	
	[%]	σ
JNT2M1	7.6	3.8
JNT2OPT	7.6	3.2
JNT2OPT2	8.5	4.0
JNT2CUT	6.2	3.7
Mean	7.5	3.7

Cost of joint vs. sequential approach. The gains obtained from the implementation of the joint approach must be balanced against the computational cost involved in performing the additional control optimizations. Due to the embedded control routines, the joint approach is significantly more costly in terms of total number of reservoir simulations that need to be performed during the optimization procedure. From Table 3.5 we have that the mean total number of reservoir simulations required by the joint runs is almost 7 times higher than the mean number of reservoir simulations required by the sequential runs. This is an important factor for real field applications where reservoir simulations are very time consuming. In such a case, a sequential approach might be a better choice to optimize for well locations.

Influence of different well-length constraint implementations. Different implementations of the well-length constraint have been used in this work (see Section 3.4.3: Non-linear constraint handling, page 84). A pair of joint and sequential optimization runs have each been launched using one of four possible implementations for the well-length operator P_{wl} . Using Table 3.5, we confirm that, as expected, the joint solutions are higher than their sequential counterparts within each of these joint-sequential pairs of runs. This result is interesting because, in practical applications, users may launch the optimization procedure using very different well-length constraint formulations.

Trade-off between constraints and approaches. For expensive problems, an attractive option to accelerate the optimization process might be to launch an optimization run using the configuration that only imposes the well-length constraint at the end of the iteration, i.e., the “M1” implementation. From Table 3.5 we see that the M1 solutions require the least number of reservoir simulations for each type of approach, i.e., 5692 and 708 for the joint and sequential M1 solutions, respectively. However, while the joint M1 solution yields the highest FOPT among all solutions in our case, the M1 solution corresponding to the sequential approach yields the lowest final objective function value. Not enough data is available to support a choice, but based only on the two current data points, it appears the simpler implementation of the well-length constraint is more amenable for

use when performing an optimization using the joint approach. Alternatively, the more sophisticated well-length constraint implementations (e.g., “OPT2” or “OPT”) could be chosen using either approach. From a practical point of view, even though using either one of these constraint implementations increases the cost of the optimization, there seems to be no major difficulty in applying them with the less costly sequential approach to achieve acceptable results.

Function evolution graphs. The trade-off between greater FOPT and computational cost is apparent when comparing the function evolution graphs for both approaches. Figures 3.16 and 3.17 both show the objective function evolution graphs for all the joint and sequential runs. In each figure, the joint and sequential function evolution curves are plotted in separate graphs (curves corresponding to the sequential solutions are plotted on the left graphs while curves from the joint runs are plotted on the graphs to the right). In Figure 3.16, all function evolution curves are plotted with respect to the number of objective function evaluation calls. In Figure 3.17, the same function evolution data is plotted but this time with respect to the cumulative number of reservoir simulations performed during the optimization procedure (recall that in the joint approach an objective function evaluation is equal to a control optimization that typically requires several reservoir simulations). Consequently, Figure 3.16(a) is equal to Figure 3.17(a), since a cost function evaluation for the sequential approach only requires one reservoir simulation. (Though equal, we plot both graphs to complement the joint-sequential figure array.) Taking particular note of the x-axis scaling in these figures, we confirm the substantially higher cost of the optimization procedure, in terms of total number of reservoir simulations, when the procedure uses the joint approach compared to when it uses the sequential approach.

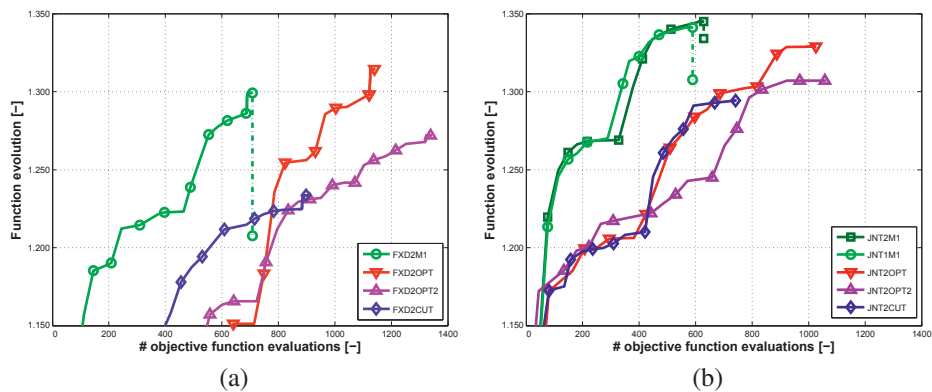


Figure 3.16: Function evolution as a function of number of function calls for 3.16(a) sequential and 3.16(b) joint runs.

Performance cost function evolution graphs. From the graphs in Figure 3.16 we can compare the function evolutions curves corresponding to the sequential solutions, shown in Figure 3.16(a), with the function evolution curves for the joint solutions, shown in Figure 3.16(b). These curves are plotted with respect to the number of objective function evaluations performed by the well placement optimization part of the procedure. The comparison is then made in terms of cost function calls for each of the two approaches.

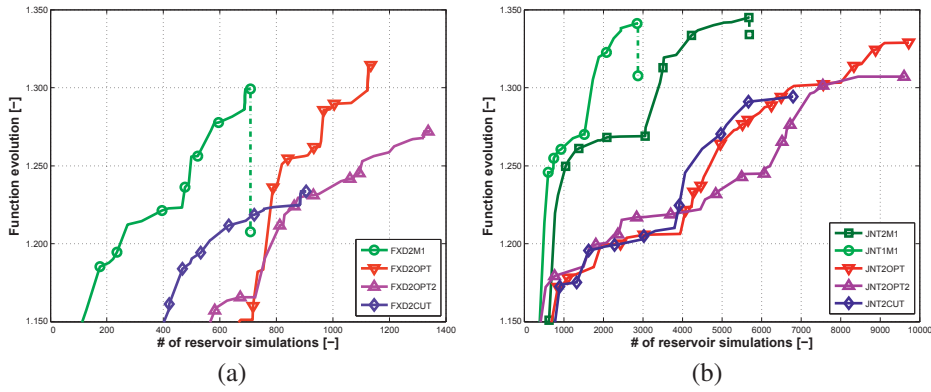


Figure 3.17: Function evolution as a function of total number of reservoir simulations for 3.17(a) sequential and 3.17(b) joint runs.

Comparing these graphs we notice each joint solution displays a better-performing function evolution curve, in terms of quicker progression and higher final FOPT, than their corresponding sequential cost function curve.

Decrease in FOPT drop due to M1 well-length constraint implementation. From figures 3.16(a) and 3.16(b) we can also observe the drop in final FOPT caused by imposing the well-length restriction after the optimization iteration has ended (as mentioned, this is the “M1” enforcement of the maximum well-length). The M1 configuration is applied in optimization run FXD2M1 in Figure 3.16(a), and in run JNT2M1 in Figure 3.16(b). (Since run JNT1M1 is a different version of JNT2M1, the general discussion regarding the latter also applies to the former. However, the JNT1M1 run is discussed separately below). Notice that a new optimization of controls is performed using the projected well configuration with shorter well lengths. From the application of “M1” we notice the close link between well length and oil recovery, in that the reduction in well lengths causes a significant drop in FOPT for both solutions. Interestingly, the drop in FOPT is less for the joint JNT2M1 run than for the sequential FXD2M1 run. (Several more data points would be needed to further extrapolate based on these results⁷.)

Comparison of runs JNT2M1 and JNT1M1. Among the joint runs, JNT1M1 is a special run in that it was launched using a setting of only 3 major iterations in the SNOPT solver used to perform the embedded control optimization. All other joint runs, on the other hand, were run using 6 major iterations in the embedded control routine. By comparing the JNT1M1 and JNT2M1 curves in Figure 3.16(b), we confirm that the two runs require close to the same number of objective function evaluations. Moreover, we have from Table 3.5 that the JNT2M1 solution, as expected, requires almost twice as many reservoir simulations, 5692 in total, as the JNT1M1 run, which performs 2873 reservoir simulations. This relationship is clearly seen in Figure 3.17(b).

The JNT1M1 run was launched to test whether a setting of only 3 major iterations in SNOPT was enough to optimize for controls at various well configurations. (Control optimizations embedded within the well placement search usually stop due to the tight major iteration limit imposed on the SQP implementation.) In this regard, we see the

Table 3.7: Performance, in terms of normalized FOPT, of solutions obtained using the AD-GPRS work model (column one) when transferred to the Eclipse field case model (columns two and three). Columns two and three result from the first and second type of transfer, respectively. The first type of transfer refers to the implementation of both optimized well placement and controls, while the second refers to the implementation of only the solution well configurations alongside original field case simulator settings (\mathbf{x}_c^S). Results in the second and third column are normalized with respect to the FOPT obtained when running the BASECASE well configuration using field case simulator settings for controls (this value is found in the first row in column three of that table). The upper level title “FOPT₁₂₀₀” refers to the fact that the results in this table show the FOPT obtained using the AD-GPRS work model approximation and the Eclipse field case model, both run for the 1200 day production time frame.

Solution	FOPT ₁₂₀₀		
	AD-GPRS	ECLIPSE	
	$(\mathbf{x}_p^*, \mathbf{x}_c^*)$	$(\mathbf{x}_p^*, \mathbf{x}_c^*)$	$(\mathbf{x}_p^*, \mathbf{x}_c^S)$
BASECASE	1.011	0.984	1.000
FXD1M1	1.208	1.125	1.133
FXD2OPT	1.315	1.229	1.235
FXD2OPT2	1.272	1.195	1.219
FXD2CUT	1.234	1.156	1.176
JNT2M1	1.334	1.249	1.255
JNT2OPT	1.329	1.243	1.249
JNT2OPT2	1.307	1.230	1.233
JNT2CUT	1.294	1.215	1.215

lower setting performs sufficiently well to yield a function evolution curve comparable to the one from JNT2M1. However, the decrease in FOPT after the well-length constraint is imposed, is significantly larger for the JNT1M1 run than for the JNT2M1 solution. More data would be needed from this test case to make further associations regarding how this setting may affect joint runs with different constraint implementations. Notice, however, that this discussion is related to our previous observation in Section 2.4.2 (page 34), regarding how an insufficient maximum number of major iterations may yield suboptimal solutions. Finally, a setting of 6 major iterations was chosen for the application of the joint approach because this setting was considered more stable with respect to the different implementations of the well-length constraint and the range of different well placement trial solutions.

Transfer of results to Eclipse field case model

As stated in the beginning of this chapter, Section 3.1: Targets and strategy for application development (page 42), one of the main targets for this optimization effort has been to test the application of our methodology on a real field case. At this point, we therefore shift our attention to focus on transferring the solutions obtained using the work model on to the original Eclipse field case model. As treated previously in Section 3.3.3: Field model transfer and validation (page 69), various approximations to original simulator functions and model properties have been introduced into the work model. Consequently, these approximations are an inherent part of the different solutions obtained using the optimization framework. Due to the differences in the work model, a general decrease in gains might be expected once the solutions are tested on the more complex, original field case model running on the commercial simulator.

On a separate note, we have up to now described the results obtained from the optimization procedure only in terms of performance, but have not yet treated the work model solutions in reservoir engineering terms. Rather, in this work we perform this type of analysis (e.g., presentation of solution well trajectories with relevant production profiles) only after the different solutions have been implemented on the original field case model. On this point, we have decided based on our second strategy component that emphasizes making as much of the current application effort as possible, accessible to the field development work process of the operator. For this purpose, we devote the entire next chapter, Chapter 4, to testing and analyzing all solutions subject to various considerations important within the perspective of field development operations. In the following, we present test results from the transfer of solutions on to the original Eclipse field case model. As a prelude to the next chapter, at the end of this section we also present saturation maps corresponding to the best-performing solution from the transfer process.

Description of transfer table. In Table 3.7 we compare how solutions obtained using the AD-GPRS work model perform when transferred to the Eclipse field case model, again validated for the 1200 day production scenario. As briefly commented, it is important to note that because the AD-GPRS work model is an approximation to the field case model, some of the gains achieved by the optimization procedure using the work model, are likely to decrease once the solutions are transferred to the Eclipse model. For reference, in column one of Table 3.7, we again show the cost function values obtained using the optimization procedure (these values have previously been presented in column two of Table 3.5). The solutions obtained from the optimization procedure are implemented in the field case model in two ways, or in two different types of transfer. In the first type of transfer, both the well placement part of the solution, i.e., \mathbf{x}_p^* , and the optimized well control settings, i.e., \mathbf{x}_c^* , are implemented in the Eclipse field case model. In a second type of transfer, only the well placement part of the solution is implemented in the field case model. For both types of transfers, to implement the well placement part of the solutions, we used the Petrel software to rebuild the well trajectories, i.e., to find the completions that correspond to the solution well coordinates, and to compute the associated well transmissibility factors. For the control part in the second type of transfer, the Eclipse test simulations were run using the same simulator settings, i.e., no approximations, as those used for the original field case model. These simulator settings

are hereby referred to as \mathbf{x}_c^S . Specifically, implementing the original setting means managing the production from the new well trajectories using the same standard simulator well and group control functions as the original field case model, including the exact same liquid and gas target/limit rates as before. Moreover, among other parameters that were modified but are now in their original state, the implementation of original settings also means that no modifications to aquifer pressure support are present, and that the simulator gas rate allocation routine for artificial lift is used. Normalized FOPT values resulting from the first and second type of transfer are shown in columns two and three in Table 3.7, respectively. In Table 3.7, all results in the second column are normalized with respect to the FOPT obtained when running the BASECASE well configuration using \mathbf{x}_c^S for controls (this value is found in the first row in column three of that table). All results are obtained using the same 1200 day production time frame that was used in the optimization procedure. In the following, we describe these results in detail.

Overall transfer of solutions. In Table 3.7, we see a general decrease in cost function values once the solutions are transferred to the Eclipse field case model. Comparing the first AD–GPRS column with the second Eclipse column, we notice that in this application, the solutions obtained through the optimization procedure using the work model yield lower field oil production totals once these solutions are implemented within the field case model. As previously noted, the work model approximations introduced during the validation process are likely the main reason why optimization gains are not fully transferred from the work model solutions to the field case application.

An additional note is that, for the BASECASE well configuration in column two, the application of optimized controls \mathbf{x}_c^* yields an almost 2% drop in FOPT compared to if we run the Eclipse simulation using original simulator settings, i.e., \mathbf{x}_c^S . (Observe that column two values in Table 3.7 are normalized with respect to this latter value.) On the other hand, we see from Table 3.7 that \mathbf{x}_c^* , which was obtained from the gradient-based routine, originally yields about a 1% increase in FOPT when using the work model. Recall also that the work model prediction for total oil production using the equivalent simulator settings, validated using the 5174 day time frame, differs with the field case model prediction by about 6%. Overall, these values give a measure of the discrepancy that can occur when using an approximation during optimization, and then utilizing the obtained solutions on the original model.

Second column values in transfer table. Notably, we see that joint runs still outperform sequential runs after the transfers of solutions to the field case model. From column two, the joint solutions yield a mean increase in FOPT of slightly more than 25%. In comparison, the sequential solutions yield, on average, an increase in FOPT of somewhat less than 20%. (Note mean increases of values within each column are computed relative to their respective BASECASE values, e.g., for the means above the BASECASE value of 0.984 is used.) Compared to the results obtained using the work model, the mean FOPT gain associated with the joint solutions has dropped by 8%, from 33% to 25%, for the first type of transfer to the Eclipse field case model. Similarly, for the same transfer process, the mean FOPT gain associated with the sequential solutions has dropped by about 6%, from 26% to 20%.

Third column values in transfer table. FOPT values presented in the third column are obtained using only the well placement part of each solution, i.e., \mathbf{x}_p^* . Joint solutions in

this column yield a mean increase in FOPT of about 24%, while sequential runs yield a mean FOPT increase of 19%. An important point to consider for further testing is that the joint solutions in this column, even though for this type of transfer only the well placement part of the solutions is implemented, still out-perform their sequential counterparts by, on average, 5% higher FOPT. The best joint and sequential solutions from this column yield increases in FOPT of 25.5% and 23.5%, respectively.

Further testing of solutions. At this point, our second strategy component encourages us to move forward with the testing of the obtained solutions on the field model to further examine the applicability of the results. Additionally, an important goal is to assess how the solutions perform once tested for a greater set of field development considerations besides those specified in the optimization scope. These issues will be explored extensively in the next chapter.

Here we take a first step in this direction by further examining the results from column three of Table 3.7 (in addition to the best solution from this set). At this point, and for further testing, we choose to use the type two transfer solutions, i.e., those that only include the implementation of x_p^* . Two main reasons for this choice are that results from these solutions are higher when considering the original field case implementation, as seen in Table 3.7, and that this testing setup uses the production schedule (Eclipse terminology for well control strategy) originally devised for the Martin Linge field development. Importantly, we consider this latter point a strong argument for achieving the main second target (Section 3.1.1) of this entire application effort. (Also, keep in mind that the joint runs in this set of test solutions is still significantly higher than the sequential runs, so arguably some optimality from the embedded nature of the approach is still retained even for this partial transfer.)

Extension of column three results from Table 3.7. In Table 3.8, the FOPT results from the third column in Table 3.7 are presented along with the differences in well oil production total (WOPT) compared to the production from the wells in the initial configuration. Table 3.8 is here reproduced from Chapter 4 because we wanted to connect the treatment of optimization solutions in this chapter to the further testing of these solutions in Chapter 4. The same applies for the associated saturation maps presented next. Results from Table 3.8 and associated maps are therefore only commented briefly here, and will be further described in Chapter 4.

Though field-wise increases for the different runs in Table 3.8 are similar, the contributions from each of the wells varies substantially between solutions. In general, the four solutions with the highest FOPT value (in decreasing order: JNT2M1, JNT2OPT, FXD2OPT and JNT2OPT2), include relative WOPT increases of more than 50% and 150% at both their B and D wells (compared to corresponding BASECASE wells). From Table 3.8 we see that, overall, the increases from the B wells do not, to any large degree, diminish the production from their neighboring A wells. The increases from the D wells, on the other hand, do significantly influence the production of their adjacent C wells. The different well trajectories, as well as their relative positioning, are shown in Figure 3.18.

Introduction to saturation maps. In Figure 3.18 we present the saturation maps at different times that correspond to the best solution from Table 3.8, i.e., JNT2M1. These saturation data are obtained from the field case Eclipse model, and plotted using Petrel. The saturation maps shown in Figure 4.2 are contour-height maps created by multiplying

Table 3.8: Transfer results given as percentage increases in field and well oil production total for the field case model, ΔFOPT and ΔWOPT , respectively. Values correspond to the 1200 day production horizon used for optimization. Field increases and well changes are given for all solutions relative to corresponding base case total oil production values.

Solution	ECLIPSE ₁₂₀₀ ($\mathbf{x}_p^*, \mathbf{x}_c^S$)				
	ΔFOPT	ΔWOPT			
		WL.-A	WL.-B	WL.-C	WL.-D
	[%]	[%]	[%]	[%]	[%]
FXD1M1	13.3	2.4	56.6	1.2	44.5
FXD2OPT	23.5	4.5	57.9	-16.7	171.4
FXD2OPT2	21.9	0.9	46.6	-16.9	177.2
FXD2CUT	17.6	0.9	44.0	-12.5	133.7
JNT2M1	25.5	-0.1	99.8	-12.5	150.6
JNT2OPT	24.9	3.2	67.1	-13.5	167.5
JNT2OPT2	23.3	3.0	63.1	-12.7	156.0
JNT2CUT	21.5	0.9	47.6	-13.6	162.1

porosity with oil saturation at a given time t . (This representation is further explained in Chapter 4.) Notice that for this figure, we also let the simulation run for the entire production time frame of 5174 days, which is the time horizon originally planned for the field case development. Oil saturations are therefore given at times $t = 0, 1200$ and 5174 days. In total, we then have six oil saturation maps, three that correspond to recovery using the wells from the initial configuration (top row), and three maps that show production using the wells from the JNT2M1 solution (bottom row).

Saturation maps for best solution: JNT2M1. In Figure 3.18, well trajectories corresponding to the JNT2M1 solution are shown in red, while initial well trajectories are shown in black. In Figure 3.18 we see how the B well from the JNT2M1 solution targets a somewhat isolated oil accumulation in the eastern part of the reservoir. The increased recovery produced by this re-positioning is a significant factor contributing to the overall increase in FOPT for this solution. We also see that the D well in this solution is longer and has a trajectory that is much closer to the C well. The greater drainage area now available to the solution D well accounts for much of its significant increase in oil production. However, its new location causes a decrease in oil production from the C well.

Before we proceed to the further testing of solutions in Chapter 4, we discuss some particular topics arising from this work.

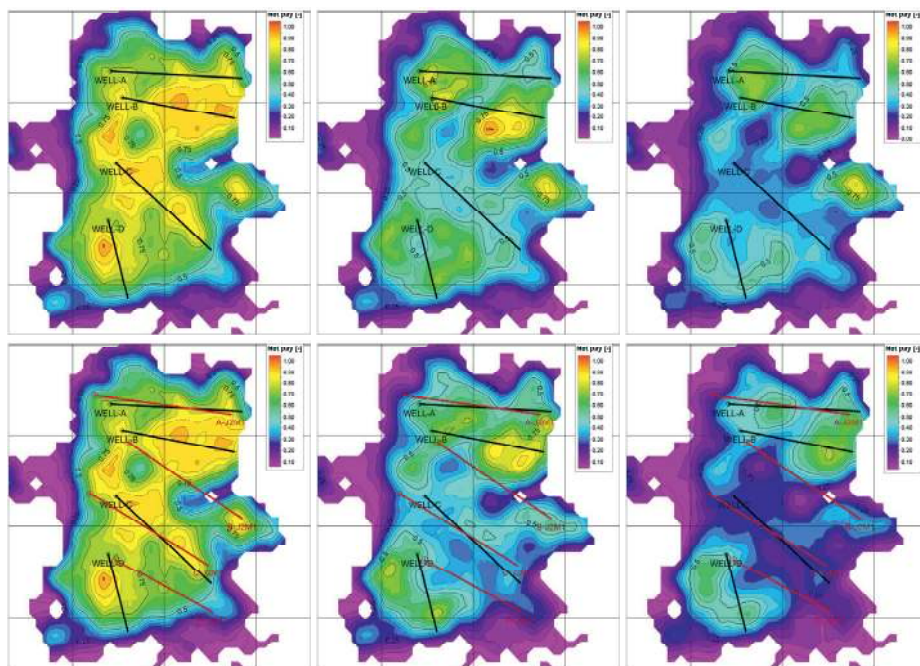


Figure 3.18: Oil saturation maps (HuPhiSo) at 0, 1200 and 5174 days of production. Above: Base case wells; below: JNT2M1 solution wells. Scales have been normalized for confidentiality reasons.

3.6 Discussion and suggestions for further work

In this final section of this chapter, we treat some of the topics arising from the optimization work and application process. The following topics are discussed: role of embedded control routine during well placement search; role of non-linear constraints during well placement optimization; critique of sequential constraint handling; and collaboration work.

Role of embedded control routine during well placement search

Results from embedded optimization. We have seen that the mean difference between final cost function values obtained using the joint and sequential approaches is of about 7% (see page 93). Furthermore, we have that embedded control optimization provides each well placement iterate an average increase in cost function of somewhat more than 7% (see page 93). A question that will be treated in further work is whether the difference between joint and sequential final cost function values will increase if the embedded control optimization problem is of a more complex nature than the ones dealt with in this work. If the embedded problem is more complex, we might expect a greater gain at each well placement trial solution due to the embedded optimization compared to using fixed controls or a reactive control strategy. Below we briefly discuss the topic of more complex formulations for the control optimization problem, and a possible implementation of inflow control valves for the development of the Martin Linge oil reservoir.

Increased gain from embedded control routine due to greater problem complexity.

The topic to be explored in the future is whether there will be a greater benefit from solving the embedded problem through optimization, if the problem involved is more complex. The idea is that the added complexity is likely to make fixed-control settings, or the use of heuristic control strategies, much less effective at a greater number of well placement trial solutions. The most obvious source for increased complexity to the type of continuous control optimization problems treated in this work is the inclusion of non-linear production constraints. Dealing with a more complex problem is not a point in itself, but rather, it is the result of more interesting, i.e., realistic, production scenarios often requiring more advanced configurations. A single aspect of problem configuration that may become more advanced and require the use of optimization techniques is the formulation of the objective function. For instance, the control problem formulation used in the latter work in this thesis is relatively straightforward, aimed at increasing cumulative oil production only. More interesting problem formulations could simply mean replacing the cost function definition from FOPT to net present value (NPV) as the objective for control optimization. NPV formulations involve computing revenue from total oil production (often with a discount factor), and may include one or several cost parameters associated with facilities and the production (and injection, if present) of water. As discussed in Chapter 2, these problem configurations are harder to solve for optimally using reactive control procedures. Overall, production problems that involve complex formulations are likely to decrease the effectiveness of simpler, e.g., heuristic, techniques, and this may increase the advantage of implementing a joint approach for well placement and control optimization.

Role of non-linear constraints during well placement optimization

Different well-length constraint implementations. The solutions from this chapter have been developed using different configurations of the well-length constraint. Most of these configurations have implemented the well-length constraint during the optimization procedure, while an additional configuration imposes the maximum well-length constraint only after the optimization procedure is completed. The different implementations of the constraint have been applied for both the joint and sequential approaches. From the function evolution curves (see figures 3.16 and 3.17 on page 95), we noticed that the decrease due to the a posteriori well-length constraint enforcement is somewhat less severe for the two joint runs (runs JNT2M1 and JNT1M1) compared to the drop observed for the sequential solution (FXD1M1). (Note that a new control optimization is performed after the well lengths in these runs are reduced. However, for each graph in question, the increase due to this new control optimization is small compared to the drop in cost function.)

At this point it is important to note that we do not consider the above results to be in any sense sufficient to make any further claims regarding the different constraint implementations. However, these results give us an idea for further work, which we describe below. It could be these results are an indication that cost function values obtained using the joint approach are less susceptible to certain changes in only one type of variable, in this case, changes corresponding to well length. This apparent robustness may be the result of joint solutions effectively integrating well trajectories with individual well control settings. As discussed in Chapter 2, page 15, the joint approach searches the space of control-optimized well locations, and a joint solution can therefore claim (local) optimality with respect to both types of variables. Consequently, the optimal control part of a joint solution may help mitigate a drop in the final cost function value, if this drop is primarily caused by only a relatively minor change in the well placement part of the solution, i.e., a decrease in well length. The above discussion is solely based on a few data points. Further work is necessary to properly test whether this property is present.

Critique of sequential constraint handling

In this work, the series of projections dealing with non-linear constraints on well placement variables has been implemented as a sequence of projection operators (see Section 3.4.3: Methodology, page 84). Handling projection operators in this manner may not be efficient since a sequential handling of constraints cannot ensure the final projection is orthogonal to the common solution space. An alternative way is to handle the feasibility constraints concurrently. If we treated all constraints associated with well placement coordinates, i.e., bounds, well-distance and well-length constraints, as one projection task, we could possibly improve the performance of the constraint-handling procedure. In our case, this would mean including all constraints into a single problem formulation that solved for the minimum distance to the feasibility bound. This reformulation will be the subject for further work.

Collaboration work

Problem translation. To achieve the stated goals for this application (see Section 3.1: Targets and strategy for application development, page 42), an effective collaboration between research partner and industry operator is key. During the course of this work we established a close collaboration with the group of engineers from Total E&P Norge AS assigned with planning the development of the Martin Linge oil reservoir. This collaboration enabled us to have ample access to the reservoir model (e.g., both as Eclipse model and as numerous Petrel projects), and facilitated information transfer and quality feedback. Moreover, the open-ended interaction was important to efficiently set up and treat specific design issues, e.g., we held various meetings and received clear information about the type of constraints that should be applied to the well placement coordinates. Furthermore, input from the group was instrumental for the work model validation process. However, despite the steadfast commitment to collaboration work from the operator team, the process of settling on various specific parameters that ultimately define the optimization problem was challenging. The challenges, from our research point of view, were often linked to not having sufficient understanding about the reservoir and/or knowledge about underlying assumptions and motivations regarding the development of the asset⁸. In the following we offer a broad outline of the tasks of problem definition and knowledge translation, based on the accomplishment of these tasks in this work, and then propose a collaboration procedure that may improve the performance of these tasks in future applications.

A very broad background for the application work conducted in this thesis is that, to foster innovation within the petroleum industry, research work needs to be challenged with realistic definitions of operation (i.e., real-life) problems (Lægreid, 2001). The general issue we focus on here is that these operation problems need to be specified as precisely as possible for the application of research to be effective. Based on this background, and our experience during the course of this work, we argue that obtaining a useful solution from the application of research on such a problem, requires that we perform a functional translation of the operation problem and the knowledge embedded in it. Importantly, with a functional translation we mean a work process that not only transfers the technical description, but also attempts to incorporate the intent and purpose of the operation problem into the end-formulation of the application problem. A primary goal of any collaboration work then, should be that this end-formulation of the application problem, stated using standard notation for mathematical programming, embodies those fundamental aspects mentioned for the operation problem. In our opinion, to achieve this end, a perceptive collaboration effort is required that combines specialized knowledge contributions from both the operator and research side, and that can also facilitate the flow of expert knowledge from one side to the other.

Further work in this regard could be the development of a test procedure where the research and reservoir team would work together on a rough visualization of the (well placement) problem search space. The idea is that over several work iterations the visualization will become a customized problem formulation. We end this section with a suggestion for such an iterative test procedure.

Test procedure for problem definition. The test procedure we are suggesting targets the translation process. The overriding motive is to meet the reservoir team halfway and to

enable them to re-express their conceptualizations within an optimization context. The visual representation is a simple way to make the end-formulation accessible to the engineers. For example, for the well placement problem, the visualization would show bounds and non-linear constraints. Importantly, through a graphical user interface, for example, it would enable the engineering team to directly manipulate the shape and parameters of these end-formulation concepts. Using the visualization, well placement constraint parameters, in this example, can then be overlapped with operation problem specifications.

The procedure is thought to be applied in a stepwise manner. Again for the well placement problem, once search regions, individual well and inter-well length and distance relationships, in addition to feasible depth intervals, have all been manipulated graphically by the reservoir team, trial optimization runs using the specified constraints would be performed. These runs would not run reservoir simulations. Rather, the optimization algorithm would use random number generators as cost functions, or possibly an analytic function based on the particular problem (without the ambition of being a surrogate). At this point, the idea is to launch a large number of trial optimization runs to resolve any issues that may be linked to the feasible space currently defined, and if possible, make the optimization process more effective, e.g., by tuning. Importantly, the very quick turn-over for each trial runs allows for several instances of the configuration-test cycle to be performed, which enables the developing problem formulation to be re-assessed and updated after each instance. This means one testing phase of the procedure would be followed by a new work iteration with the operator team where the current problem formulation would be evaluated. After each instance, the reservoir team would provide feedback, add to and further tune the visualized problem formulation following their own technical specifications, until the knowledge that they consider the most important about the operation problem, including some of the original intent and purpose, are sufficiently represented. At this point we would launch the optimization proper using reservoir simulations. A dedicated collaboration test procedure as the one described above could be a way to develop end-formulations that are highly tuned to the business needs of the operator.

The next chapter deals with the testing of obtained solutions on the real field case model.

Testing of Solutions on Field Case Model

This chapter tests the well placement solutions developed in the previous chapter on the original base case Eclipse model for the Martin Linge oil reservoir. We start by introducing the eight solution cases that have resulted from implementing various solutions within the field case model. Each simulation case is then studied individually. For each case, we plot the corresponding well placement configuration against the base case wells. Oil recovery at the different time horizons used in this work (i.e., 1200 and 5174 days) are presented using saturation maps. Also, for the different configurations, production profiles are given to study the performance of each well. Results from the simulation cases are then studied collectively. For the two production time frames, we provide tables presenting field and well oil production totals for each case. We then plot the increases in field oil production totals, and the changes in well oil production totals, corresponding to each of the solution cases, against the sum and individual well lengths associated with each well configuration, respectively. Finally, all solution cases are tested over a multiple realization scenario. From these tests, we provide tables presenting expected values (represented by $\langle \ \rangle$ delimiters), along with associated standard deviations, of field and well oil production totals for each case. We end this chapter with trying the application of simple heuristic rules to modify the well configurations obtained using the optimization procedure. In this context, the application of heuristic rules (thought to be created using expert knowledge about the development) is seen as a way to modify and/or adapt the obtained solutions according to engineering considerations that were not part of the optimization process. In the following section we start our treatment of the individual solution cases.

4.1 Test results from solution cases

Results from eight simulation cases are presented. Importantly, the results from these solution cases are obtained using the exact model and simulation parameters (including well control settings) as those applied in the original base case model used for validation (see Section 3.3.3: Field model transfer and validation, page 69). This means that for the eight solution cases, the only parameters that are different from the base case settings are the well names and the well connection factors associated with the new well grid blocks (i.e.,

only the `WELSPECS` and `COMPDAT` Eclipse keywords are modified).

Each simulation case has a well configuration that corresponds to one of the well placement solutions from the optimization effort (Section 3.5, page 91). More specifically, each well placement solution obtained from our methodology provides all the heel and toe coordinates needed to make the well configuration for a simulation case. To create the actual well trajectories for the new configuration, we transfer the coordinates to Petrel (Schlumberger, 2012c). Petrel is an industry-standard software platform used for regular petroleum engineering tasks such as grid parameter visualization and well design. Once heel and toe coordinates are defined within the reservoir grid specified by Petrel, we are able to design the well bore, determine which grid cells are traversed by the well trajectories, and to calculate the well connection factors associated with each of the well grid blocks. Finally, the entire reservoir model with the new well configuration is exported from Petrel as an Eclipse reservoir simulation case.

Tables presenting simulation case parameters. Eight well placement solutions have been adapted to Eclipse simulation cases. The parametrization of some of the Eclipse keywords used in these cases is shown in tables 4.1 and 4.2. We have selected these keywords for presentation because they have the greatest relevance to the simulation cases and results presented in this section. The specificity of some of the data shown in the tables is meant to help calibrate the results from this section to other simulations that may have been run by Total E&P Norge AS with different model parameters. We reiterate that the parameters shown in tables 4.1 and 4.2 are the original parameters without any of the approximations introduced from the validation effort, and that all the results presented in this section are based on reservoir simulations using these parameters.

Table 4.1 shows the main grid, simulation and fluid properties for the simulation runs. The table presents production times, grid dimensions, number of active grid cells, and the general dimension and set of keywords used in the configuration of the aquifer. We notice that production from the B well starts 584 days after production start from the remaining A, C and D wells, and that the hysteresis option is *on*, which requires separate tables for drainage and imbibition processes for both the water-oil and gas-oil saturation functions (*SWOF* and *SGOF*, respectively).

Importantly, in this and subsequent sections, we consider two production time frames when analyzing and comparing the results both between simulation cases and against the base case. The first production time frame is the 1200 day production horizon used for the optimization effort in Section 3.5, while the second production time frame is the planned field model production horizon of 5174 days. Both these production horizons are introduced in Table 4.1. During our discussion we should be aware that, when testing a solution from the optimization effort in a field model simulation case, the well placement configuration developed for this case is optimal with respect to the 1200 day production time frame, and that some of this optimality is necessarily lost when testing the same case for the larger production horizon. (During problem design, we attempted to reduce this expected loss by selecting an optimization time frame that contained most of the reservoir dynamics, e.g., a time after which production rates were observed to be near constant when using the base case configuration; see Section 3.1, page 42, for further discussion on this topic; and gas, oil and water well production rates for peak values in figures 6.1, 6.2

and 6.3 in Appendix A, page 209.) Through the analyses in this and subsequent sections we want to obtain a proper measure of the different increases achieved from the optimization effort, while also find how well the production time frame approximation performs on the intended field scale. For this reason, in the following we treat all simulation results using both the reduced production time frame of 1200 days, and the planned field model production horizon of 5174 days.

Table 4.2 shows the main Eclipse keyword parameters for well description and production strategy. `WELSPECS` and `COMPDAT` are the main keywords used to define a well. The first keyword specifies the preferred phase for the well and which well block that should be used for bottom-hole pressure measurement. The second keyword specifies which grid blocks the wells have perforated, and the well connection factors associated with each of the perforations. A well definition is realized by specifying the state of the well (`OPEN` or `SHUT`) using the `WCONPROD` keyword. If `OPEN`, this keyword furthermore specifies the main control parameter for the well, e.g., gas, oil, water, liquid, bottom-hole or tubing-head pressure, or if the well is to be subordinated to the controls of the group the well belongs to (`GRAT`, `ORAT`, `WRAT`, `BHP`, `THP`, `LRAT` or `GRUP` control settings, respectively). The control of wells as groups or subgroups, each with specific target rates that also function as constraints for fluid production, is an important feature of the original base case production strategy represented in Table 4.2. The grouping of wells is specified in the `GRUPTREE` keyword, while the definition and enforcement of group controls is handled by the `GCONPROD` keyword. We note that the information specified in `WELSPECS` regarding the preferred phase of a well is used by the `GCONPROD` keyword for well-group control handling during simulation (e.g., determining worst offending well once a group production rate limit is exceeded; Schlumberger, 2012a). Furthermore, the keywords `LIFTOPT` and `WLIFTOPT` are used to tune the rate of lift gas injected to boost production from the individual wells. Resulting flow rates due to lift gas injection are calculated using `VFP` tables (which `VFP` table is used at a well is given by the number in curly parenthesis in the `WCONPROD` parameter specification). Further detailed descriptions of keyword function and settings can be found in the Eclipse reference manual (Schlumberger, 2012a).

Section structure, individual and collective treatment of results. The first part of this section, Section 4.1.1, treats the results from each of the simulation cases independently. In this segment, the well placement configuration and production profiles for each simulation case are compared against the well trajectories and production rate profiles from the base case (see below). We refer to each simulation case using the name of its corresponding well placement solution. Also, from hereon, the terms "solution" and "simulation case" are used interchangeably. Solution results are analyzed using saturation maps given at start of production (i.e., saturation data at 0 days), and after 1200 and 5174 days of production. Production profiles from the individual wells are plotted against base case production curves. For confidentiality reasons, each of these performance profiles is normalized with respect to the peak value of its corresponding base case profile. The solutions covered are `JNT2M1`, `FXD1M1`, `JNT2OPT`, `FXD2OPT`, `JNT2OPT2`, `FXD2OPT2`, `JNT2CUT` and `FXD2CUT`. (We cover the `JNT1M1` solution within the discussion of the `JNT2M1` solution, since, except for only minor variations, the `JNT1M1` and `JNT2M1`

Table 4.1: The table shows main parameters for grid, simulation and fluid properties used for all field test solution results in this work. The specificity of the data shown in the table is intentional. This table, and Table 4.2, are meant to help compare the results from this work with other simulation results that may have been obtained using different simulator settings.

ECLIPSE model simulations parameters	
Grid properties	
Grid dimensions (DIMENS)	[59, 82, 76]
Number of active grid cells (ACTNUM)	51486
Analytic aquifer function	Use of AQUCT, AQUANCON
Aquifer dimensions (AQUDIMS)	[1, 1, 1, 36, 1, 9652, 1, 1]
Simulation setup	
Production start (A, C, D wells) [date, days]	[01 NOV 2016, 0]
Production start (B well) [date, days]	[01 JUN 2018, 584]
Production end, optimization case [date, days]	[14 FEB 2020, 1200]
Production end, field case [date, days]	[31 JAN 2031, 5174]
Fluid properties	
<i>PVT functions</i>	
PVTO table (first line)	[0, 1.0000, 1.0175, 5.652]
PVDG table (first line)	[1, 1.1784, 0.0131]
<i>Relative permeability</i>	
SATOPTS (RUNSPEC option)	HYSTER
SWOF SGOF tables (first and last two lines):	
SWOF Drainage table	SWOF Imbibition table
[0.0800 0.0000 1.0000 10]	[0.0800 0.0000 1.0000 1E-5]
[0.1830 0.0046 0.7228 3]	[0.1830 0.0046 0.7228 3E-6]
⋮	⋮
[0.8090 0.2300 0.0000 1*]	[0.8090 0.2300 0.0000 7E-7]
[1.0000 0.2300 0.0000 0]	[1.0000 0.2300 0.0000 0]
SGOF Drainage table	SGOF Imbibition table
[0.0000 0.0000 1.0000 0]	[0.0000 0.0000 1.0000 0]
[0.0300 0.0000 0.8900 0]	[0.2440 0.5800 1.0000 0]
⋮	⋮
[0.8320 0.8020 0.0003 0]	[0.7658 0.5362 0.0000 0]
[0.9200 0.9200 0.0000 0]	[0.9200 0.9200 0.0000 0]

Table 4.2: The table shows main parameters for Eclipse keywords related to well description and production strategy (i.e., group and individual well controls and gas lift settings). VFP table numbers are given in curly brackets in WCONPROD keyword settings.

ECLIPSE model simulations parameters	
Well descriptions	
<i>Well specification (WELSPECS)</i>	
A and D wells (preferred phase)	[GASPROD (GAS)]
B and C wells (preferred phase)	[OILPROD (OIL)]
B well (preferred phase, from 01 AUG 2018)	[OILPROD (LIQ)]
<i>Completion data (COMPDAT)</i>	
Well connection factor, skin factor	[By Petrel, 15]
<i>Well groups (GRUPTREE)</i>	
[1: FIELD (2GROUPS)]	[3: DUMMY2 (GASPROD)]
[2: 2GROUPS (DUMMY2)]	[4: 2GROUPS (OILPROD)]
Production strategy	
<i>Group well-controls (GCONPROD)</i>	
For wells in group OILPROD:	
[ORAT, 5500, 1*, 1.0E6, 5500]	[RATE, 4*, RATE, RATE]
For wells in group GASPROD:	
[GRAT, 10600, 1*, 2.0E6, 10600]	[RATE, 4*, RATE, RATE]
For wells in group 2GROUPS:	
[ORAT, 10600, 1*, 2.1E6, 10600]	[RATE, 4*, RATE, RATE]
<i>Well-controls (WCONPROD)</i>	
For wells A, C, and D:	
[OPEN, GRUP, 1*, 1*, 1.5E6]	[5500, 1*, 70, 30, {3,2,4}, 4*]
For B well:	
[OPEN, LRAT, 1*, 1*, 1.5E6]	[5500, 1*, 70, 30, {1}, 4*]
WECON (all wells, from 01 NOV 2019)	[2*, 0.9800, 2*, WELL]
<i>Gas lift</i>	
LIFTOPT	[1000, 0.0, 0.0]
WLIFTOPT (all wells)	[YES, 500000, 1*, -1.0]
Other runtime parameters	
DRSDT	[0.0001]
GEFAC	[FIELD, 0.92]

solutions are very similar to one another.)

The second part of this section, Section 4.1.2, treats the results from the solutions in a collective manner. Tables are given that present the relative increases in field and well oil production totals at both the reduced production time frame of 1200 days, and at the field model production time frame of 5174 days. Also in this segment, total field and well oil production curves are compared to base case field cumulative profiles. In Section 4.1.3, we plot the increases in field oil production total corresponding to each of the solutions, against the difference in total production drains length between the given solution and the base case well configuration. This final analysis represents the important trade-off between higher recovery against the increased uncertainty associated with drilling longer wells.

Presentation of results. Results are presented both in the form of saturation maps and in terms of selected production profiles (see below for further description of oil saturation HuPhiSo maps and selection of production curves). Both the configuration of the oil saturation maps and the selection of production profiles presented in this section are based on general guidelines from Total E&P Norge AS about what types of information content are common and useful from an industry perspective. Emphasis has been put on making the results from this work in general, and the results from the field testing of solutions in particular, into pieces of information that can be readily accessed by an industry work process. Our ultimate communication goal has been to effectively shape and channel the information contained in our optimization results such that they can easily be studied and further analyzed within the development process of the reservoir management plan.

Well placement solutions and oil saturation maps. Here we briefly explain the general characteristics of the saturations maps and the display of well configurations within these maps. To make a HuPhiSo saturation map we take the oil saturation (S_o) at a given time t , i.e., S_o^t , multiply it with the porosity (ϕ), and then sum the product over all reservoir layers in order to create a map of total oil in place. The result is a map with depth-contours showing the spatial distribution of oil volumes at time t . For all maps, base case well trajectories and name labels are plotted in black, while solution wells and labels are drawn in dark red color.

Production profiles (WOPR, WGOR, WCUT) and tables of well length and depth. For all wells in each solution, we plot production profiles for well oil production rates (WOPR), well gas-oil ratios (WGOR) and well water cuts (WWCT). The curves for the solutions are drawn as thick lines while the corresponding base case profiles are drawn as thin lines in the same graphs. WOPR and WGOR profiles from solutions are normalized using the maximum values from the WOPR and WGOR base case profiles, respectively. Production profiles are shown over the production time frame used in our optimization, i.e., the 1200 day production horizon. The main reason for selecting this time window in our graphs is that, since tail-production has already ensued before 1200 days have passed in practically all cases, the more interesting comparisons in rate and ratio changes to document are those that occur in the first few years of production. Consequently, our description of dynamic well performance for the different well placement configurations has

Table 4.3: Final well lengths for base case and solutions. Notice these well lengths have previously been introduced together with their corresponding (work model) FOPT values in Table 3.5.

Solution	Well lengths [m]			
	WL.-A	WL.-B	WL.-C	WL.-D
BASECASE	1439	1247	1409	874
FXD1M1	1500	1500	1500	1200
FXD2OPT	1364	1139	1299	1358
FXD2OPT2	1129	1202	1295	1362
FXD2CUT	1129	1202	1373	1069
JNT2M1	1500	1500	1500	1200
JNT1M1	1500	1500	1500	1200
JNT2OPT	1088	1474	1299	1316
JNT2OPT2	1129	1142	1302	1155
JNT2CUT	1129	1034	1299	1271

a perspective based on this time horizon. Still, saturation maps are shown and commented at both 1200 and 5174 days. Furthermore, we will consistently use both time frames in all well and field analyzes in Section 4.1.2. In that section we will compare all of the solutions against the base case in terms of increases in total cumulative oil produced for the whole field and for individual wells.

In the following section, Section 4.1.1, we will start our individual treatment of the base case and the different simulation cases. Since these descriptions readily use final base case and solution well length and depth information, we choose to already here present this information in tables 4.3 and 4.4, respectively.

4.1.1 Individual analysis: Final well configurations

BASECASE solution

A general strategy for base case well design is to seek maximum recovery while minimizing drilling risks related to uncertainties in both reservoir structure and extent of hydrocarbon accumulation¹. The base case well design for the Martin Linge oil reservoir consists of four horizontal production wells (WELL-A, WELL-B, WELL-C, and WELL-D – these wells will commonly be referred to as base case A, B, C and D wells, respectively) with lengths of around 900 to 1450 meters (see Table 4.3). Current operational strategy is to drill and initiate production from wells A, C and D in a concurrent manner, while delaying production from the B well until about a year and a half after production start. This operational sequence is part of the production strategy and has been implemented in all simulation cases (see Table 4.1 for details regarding production start for each well).

Table 4.4: Final well depths for base case and solutions.

Solution	Well depths [m]			
	WL.-A	WL.-B	WL.-C	WL.-D
BASECASE	1742	1739	1740	1740
FXD1M1	1744	1739	1740	1740
FXD2OPT	1744	1739	1740	1744
FXD2OPT2	1742	1739	1740	1744
FXD2CUT	1742	1739	1740	1744
JNT2M1	1746	1739	1740	1744
JNT1M1	1746	1739	1740	1744
JNT2OPT	1742	1739	1740	1742
JNT2OPT2	1742	1739	1740	1740
JNT2CUT	1742	1739	1740	1744

The Martin Linge oil reservoir has a relatively thin oil column of about 20 meters which causes the development of the reservoir to have a high probability of early water breakthrough. Well depth–positioning is therefore a significant concern for the development of the reservoir. To delay water breakthrough as much as possible, the reservoir development plan places current base case wells close or right below the gas–oil contact (GOC)². However, these base case wells are still expected to experience a rapid water breakthrough due to the strong aquifer support and high mobility of water with respect to oil (reservoir flow properties were introduced in Section 3.1, page 68). Well water cut profiles are therefore an important measure of performance when considering production from optimized well positions for this particular development case. In particular, when considering the performance of well placement solutions in the following segments, we will focus on whether the new locations show delayed water breakthrough while at least producing the same amount of hydrocarbon volumes (though oil production is our main concern).

The development plan for the Martin Linge oil reservoir accounts for the presence of a relatively small accumulation of free gas in the reservoir. The planned depth–locations for the base case wells in this development are expected to yield high gas production rates (also compared to oil rate, leading to a high gas–oil ratio; GOR) during the first months of production. We see this trend in the GOR for the base case A, C and D wells in Figure 4.3 (thin purple line). For each of these wells, the GOR peaks within six months of production, and then drops to about or less than a third of its peak value after 300 days.

Well lengths and depth–positioning are fundamental design parameters for overall drainage strategy and have direct influence on drilling cost, production efficiency and ultimate recovery³. In an operational context, decisions regarding well lengths and depth are constantly being checked against reservoir uncertainties, e.g., uncertainties associated

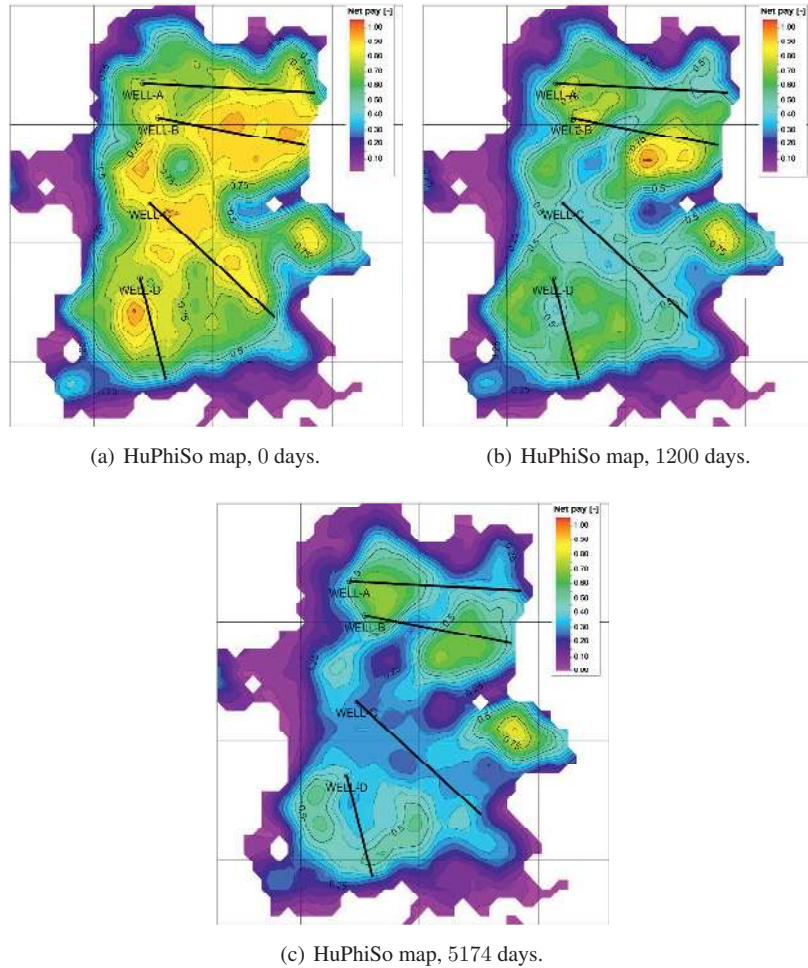


Figure 4.1: HuPhiSo oil saturation maps at three production times for BASECASE solution. Scales have been normalized for confidentiality reasons.

with the curvature and depth of the top reservoir surface. For this type of uncertainty, a decision may be included in the field development plan to drill pilot holes before the actual base case production wells. For the current development plan for the Martin Linge oil reservoir a choice has been taken to drill pilot holes for each well in order to increase the probability that the horizontal well bores both hit their determined reservoir entry points and are drilled within the specified depths⁴. Another way to deal with this general uncertainty is to plan for a sequential drilling and operation of wells. For the Martin Linge base case, performance information from the first batch of wells will help de-risk the drilling and ultimate placement of the base case B well. Topics such as pilot well drilling and sequential operation of wells are not treated in our optimization routine. However,

these topics are still commented here because they show the larger industry context within which the original base case configuration will be compared against the solutions found by the optimization procedure. Similar contextual topics and considerations related to the drilling of the individual base case wells are briefly discussed below.

WELL-A. The pilot well planned for the WELL-A well has a significant influence on the reservoir entry point of the WELL-A horizontal well-bore. At one point during planning, it was considered that a neighboring Brent well trajectory passing by the oil reservoir could serve as a pilot hole. Current plan, however, is to drill a dedicated pilot hole for this well.

WELL-C. The current placement of WELL-C and its pilot well are expected to produce information to significantly reduce the uncertainty around the top structure saddle area in the central part of the reservoir (see topic on reservoir structure, see Section 3.1, page 67).

WELL-D. Current base case trajectory specifies the drilling of a pilot hole for this well. However, for some time during planning, an option was considered where the trajectory was thought to be shifted counter-clockwise facing the eastern part of the reservoir. The advantage of such a configuration would have been that an existing discovery hole could then be used for calibration purposes. This option would have saved a pilot hole but lost the additional appraisal information regarding the reservoir saddle point that is likely to be gained from the original pilot well position. Eventually, it was decided a dedicated pilot hole would be drilled for this well, as for the other wells.

We end here our discussion of the general development context behind the base case, and start the descriptions of each of the simulation case results. Notice that each description is placed in a text-saturation map-production profiles arrangement, which leaves some blank spaces on some of the pages, but allows us to present the results in an ordered fashion. We start by describing the JNT2M1 solution.

JNT2M1 solution

Compared to base case results, the JNT2M1 solution yields increases in field oil production total of 26% and 13% at 1200 days and after 5174 days of production, respectively (see tables 4.5 and 4.6 for full details). At 1200 days, increases in well cumulative oil production are 0%, 100%, -13% and 151% for wells A, B, C and D, respectively. At 5174 days, increases in well cumulative oil production are -5%, 61%, -21% and 86% for the same wells, respectively. HuPhiSo saturation maps at 0 days and at production times 1200 and 5174 days are given next. Following this we given individual descriptions of location and performance for the wells in the JNT2M1 solution.

A-J2M1 and C-J2M1 wells. The general trajectories of the A-J2M1 and C-J2M1 wells closely resemble those of the base case A and C wells. However, the lengths of A-J2M1 and C-J2M1 are somewhat longer, by 61 and 91 meters, respectively, and their heel positions are significantly shifted westward. Moreover, the C-J2M1 trajectory is somewhat rotated counter-clockwise relative to its base case counterpart, while the A-J2M1 well

has a slight clockwise rotation and a heel shifted toward the northern border of the accumulation. The A-J2M1 well lies at a depth of 1746 meters, four meters deeper than the base case A well (see Table 4.4). This heel shift and depth change may impact pilot well decision making.

In terms of performance we observe the shift and rotation have removed the peak gas production for this well (see GOR at upper left graph, Figure 4.3). It is also worth noting that water breakthrough for this well starts at a significantly earlier time. Finally, while the A-J2M1 well toe is moved away from the border towards the interior of the reservoir, the heel is brought closer to the flank on the other side. The former movement may possibly reduce some risk of drilling beyond good reservoir sands, but it is clearly counteracted by the latter repositioning which brings the well heel dangerously close to the structural flank with more risk to enter into a non-reservoir zone.

B-J2M1 and D-J2M1 wells. Both the B-J2M1 and D-J2M1 well trajectories are significantly different than their base case counterparts. Importantly, the B-J2M1 well stretches over to a small pocket at the eastern part of the reservoir. While production rates (see Figure 4.3) for the A and C wells in the JNT2M1 solution are comparable to those from the base case wells, the B and D wells show substantial oil rate increases compared to their base case counterparts. The oil rate increase from the B-J2M1 well may be attributed to its longer well bore (this well is 1500 meters long, 253 meter longer than its base case counterpart), and to its position tapping the small height at the eastern lobe.

The D-J2M1 well is 326 meters longer than the D base case well, and has an eastward direction that intrudes into a reservoir area drained exclusively by the C well in the base case. The JNT2M1 solution accommodates the C-J2M1 and D-J2M1 wells in a close-to-parallel array, with the C-J2M1 well somewhat shifted to the north (thus honoring inter-well distance constraint). Even though the drainage areas of the C-J2M1 and D-J2M1 wells both cover the same eastern area of the reservoir, the individual rates from these wells are equal to or greater than the rates from their base case counterparts. The depth of the D-J2M1 well is 1744 meters, compared to 1740 meters for its base case well analog. Interestingly, water cuts for both the B-J2M1 and D-J2M1 wells increase at a significantly lower pace and are markedly delayed.

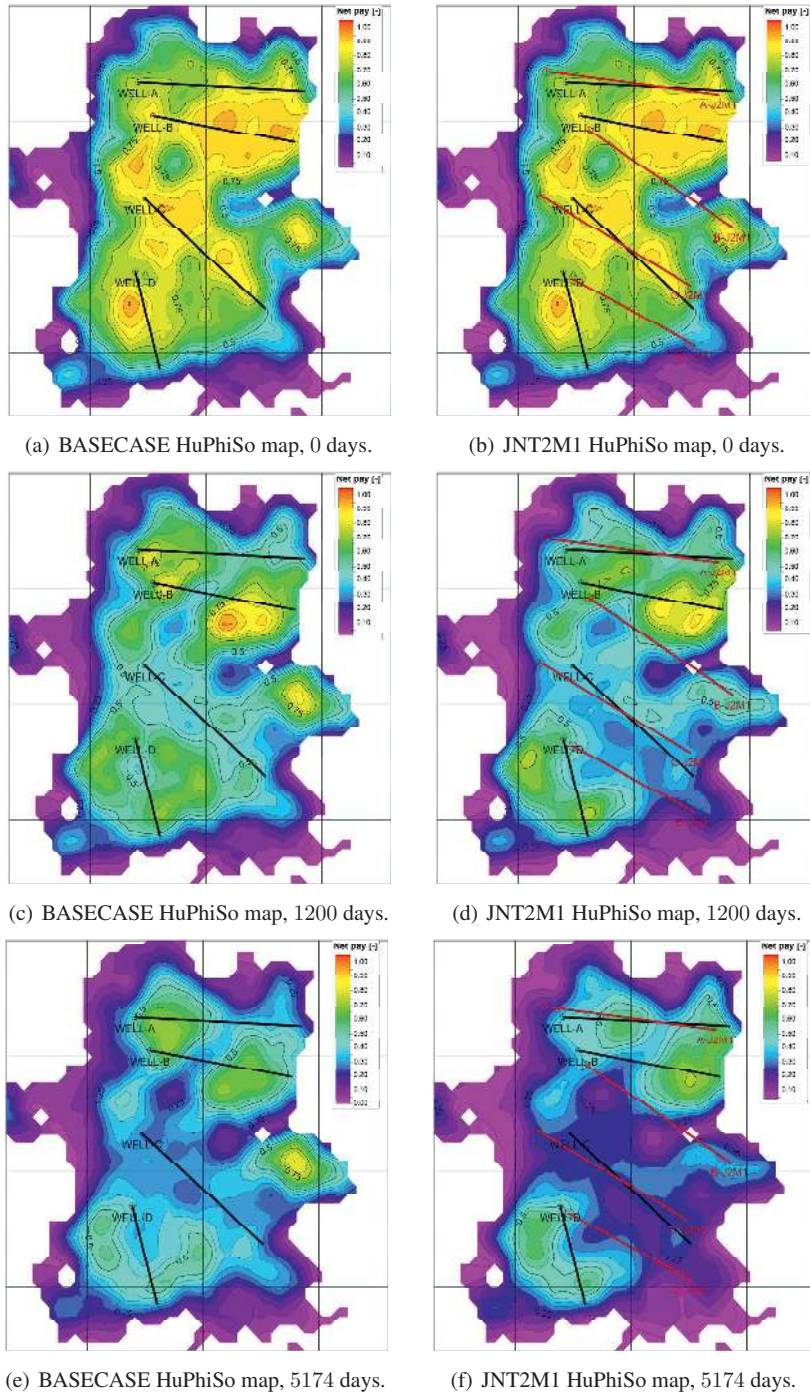


Figure 4.2: HuPhiSo oil saturation maps at three production times for BASECASE (left column) and JNT2M1 solution (right column). Scales have been normalized for confidentiality reasons.

4.1 Test results from solution cases

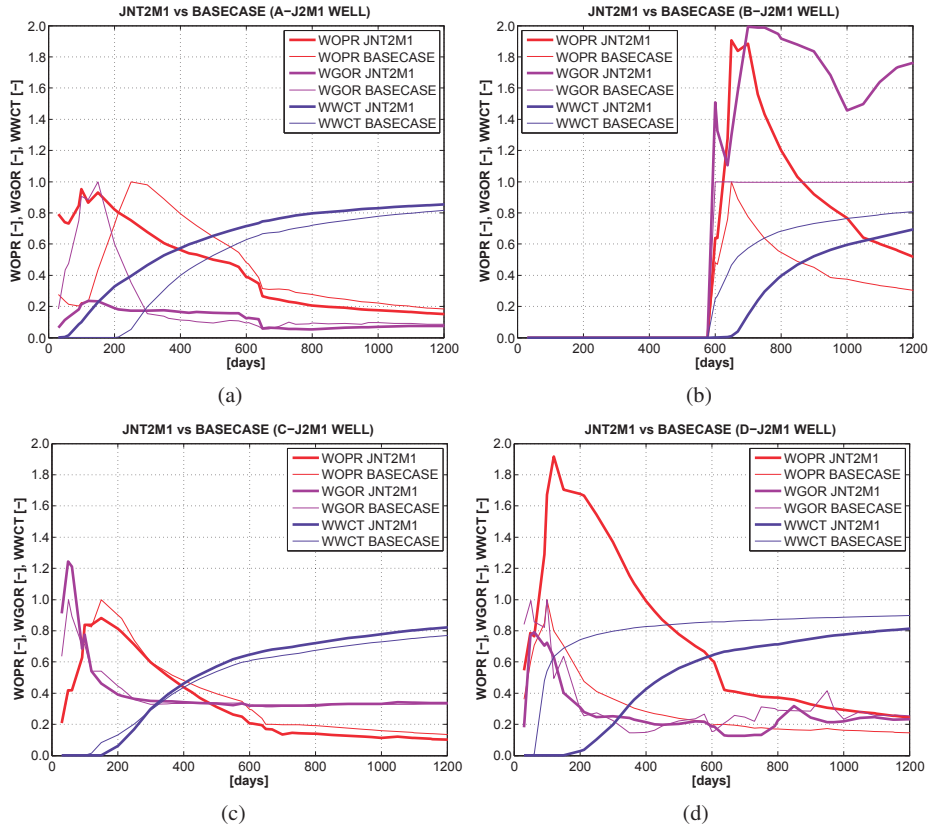


Figure 4.3: Production profiles for JNT2M1 solution and base case wells: Well oil production rates (WOPR), well gas-oil ratios (WGOR) and well water cuts (WWCT).

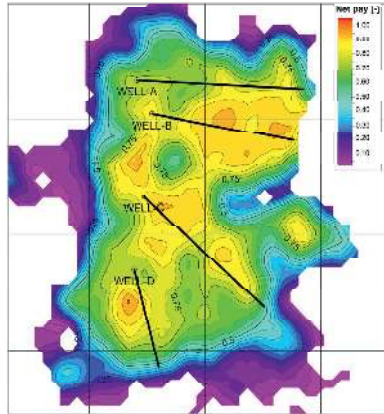
FXD1M1 solution

Compared to base case results, the FXD1M1 solution yields increases in field oil production total of 13% and 8% at 1200 days and after 5174 days of production, respectively (see tables 4.5 and 4.6 for full details). At 1200 days, increases in well cumulative oil production are 2%, 57%, 1% and 45% for wells A, B, C and D, respectively. At 5174 days, increases in well cumulative oil production are 3%, 18%, 4% and 17% for the same wells, respectively. HuPhiSo saturation maps at 0 days and at production times 1200 and 5174 days are given next. Following this we give individual descriptions of location and performance for the wells in the FXD1M1 solution.


A-F1M1 and B-F1M1 wells. A and B wells in the FXD1M1 solution, named A-F1M1 and B-F1M1, are 61 and 253 meters longer than their analog A and B base case wells. Moreover, the trajectories for both these wells have a substantial clockwise rotation compared to their base case counterparts. Their depths are 1744 and 1739 meters, which means that A-F1M1 lies four meters deeper and B-F1M1 lies one meter shallower than the base case A and B wells, respectively. Also, the reservoir entry points for the A-F1M1 and B-F1M1 wells, i.e., their heel positions, are moved substantially from their original location. For the A-F1M1 well, the new entry point is significantly shifted north, while the entry point for the B-F1M1 well is moved westward. However, production profiles for the A-F1M1 well roughly match those for the base case A well, while, compared to the base case B well, production profiles for the B-F1M1 well show a similar but higher oil production rate and a delayed water breakthrough.

C-F1M1 and D-F1M1 wells. The C-F1M1 well trajectory resembles the base case C well trajectory, but is 91 meters longer and its entire well bore is shifted west. The production profiles for the C-F1M1 well closely match those from its base case analog. The depths for both the C-F1M1 and D-F1M1 wells remain unchanged from base case depths. The D-F1M1 well is positioned in the south-west flank of the reservoir, and is 326 meters longer⁵ than the base case D well. This well has a substantially larger oil production rate, and even though water breakthrough time is the same, the water cut has a slower increase than that of the base case D well.


4.1 Test results from solution cases




(a) BASECASE HuPhiSo map, 0 days.




(b) FXD1M1 HuPhiSo map, 0 days.




(c) BASECASE HuPhiSo map, 1200 days.



(d) FXD1M1 HuPhiSo map, 1200 days.



(e) BASECASE HuPhiSo map, 5174 days.



(f) FXD1M1 HuPhiSo map, 5174 days.

Figure 4.4: HuPhiSo oil saturation maps at three production times for BASECASE (left column) and FXD1M1 solution (right column). Scales have been normalized for confidentiality reasons.

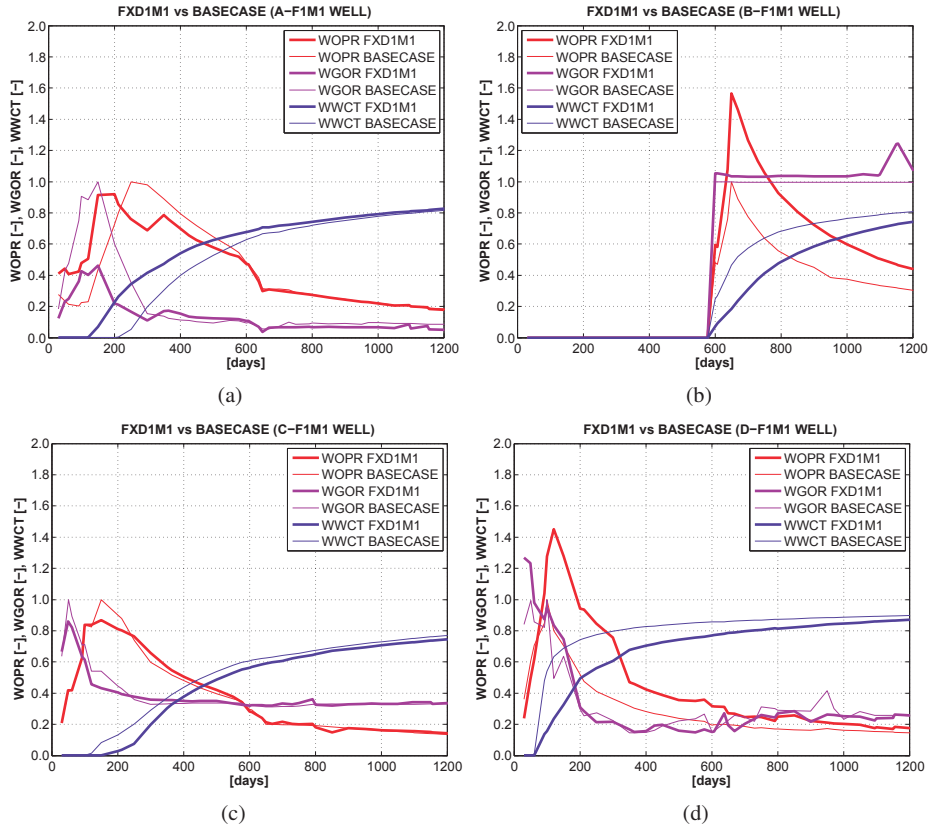


Figure 4.5: Production profiles for FXD1M1 solution and base case wells: Well oil production rates (WOPR), well gas–oil ratios (WGOR) and well water cuts (WWCT).

JNT2OPT solution

Compared to base case results, the JNT2OPT solution yields increases in field oil production total of 25% and 13% at 1200 days and after 5174 days of production, respectively (see tables 4.5 and 4.6 for full details). At 1200 days, increases in well cumulative oil production are 3%, 67%, -14% and 168% for wells A, B, C and D, respectively. At 5174 days, increases in well cumulative oil production are 9%, 28%, -26% and 104% for the same wells, respectively. HuPhiSo saturation maps at 0 days and at production times 1200 and 5174 days are given next. Following this we given individual descriptions of location and performance for the wells in the JNT2OPT solution.

A-J2OT well. All trajectories in this solution, except the one of the A-J2OT well, vary sharply from the base case configuration. The A-J2OT well trajectory closely matches the A well base case trajectory, and the depth is an exact match. Interestingly, the production profiles for the A-J2OT well and the base case A well are almost the same, even though the length of A-J2OT is only 1088 meters, compared to 1439 for the base case A well, a 351 meters difference (or almost 25% decrease in well length).

B-J2OT well. The heel of the B-J2OT well is considerably further south compared to the heel of the base case B well. And with a length of 1474 meters, compared to 1247 meters for the base case B well, the B-J2OT well yields an effective sweep of both the north-central area and the east lobe accumulation. Furthermore, this well offers a robust increase in oil production rate and a modest delay in water breakthrough time.

C-J2OT well. At 1299 meters, the C-J2OT well is 110 meters shorter than its base case analog. In particular, the C-J2OT well heel is moved close to the west boundary of the accumulation, while the toe of the well is pulled away from the east boundary towards the central region of the reservoir. Despite the dissimilarities, the C-J2OT well has production profiles that roughly resemble those of its base case well analog.

D-J2OT well. Similarly to the D well in the JNT2M1 solution, the D-J2OT well is positioned in a close to parallel trajectory relative to the C-J2OT well. This positioning avoids most of the south-east area targeted by the original base case D well. Still, the D-J2OT well yields a substantial increase in oil production, a somewhat delayed water breakthrough time and a gradual rise of the water cut.

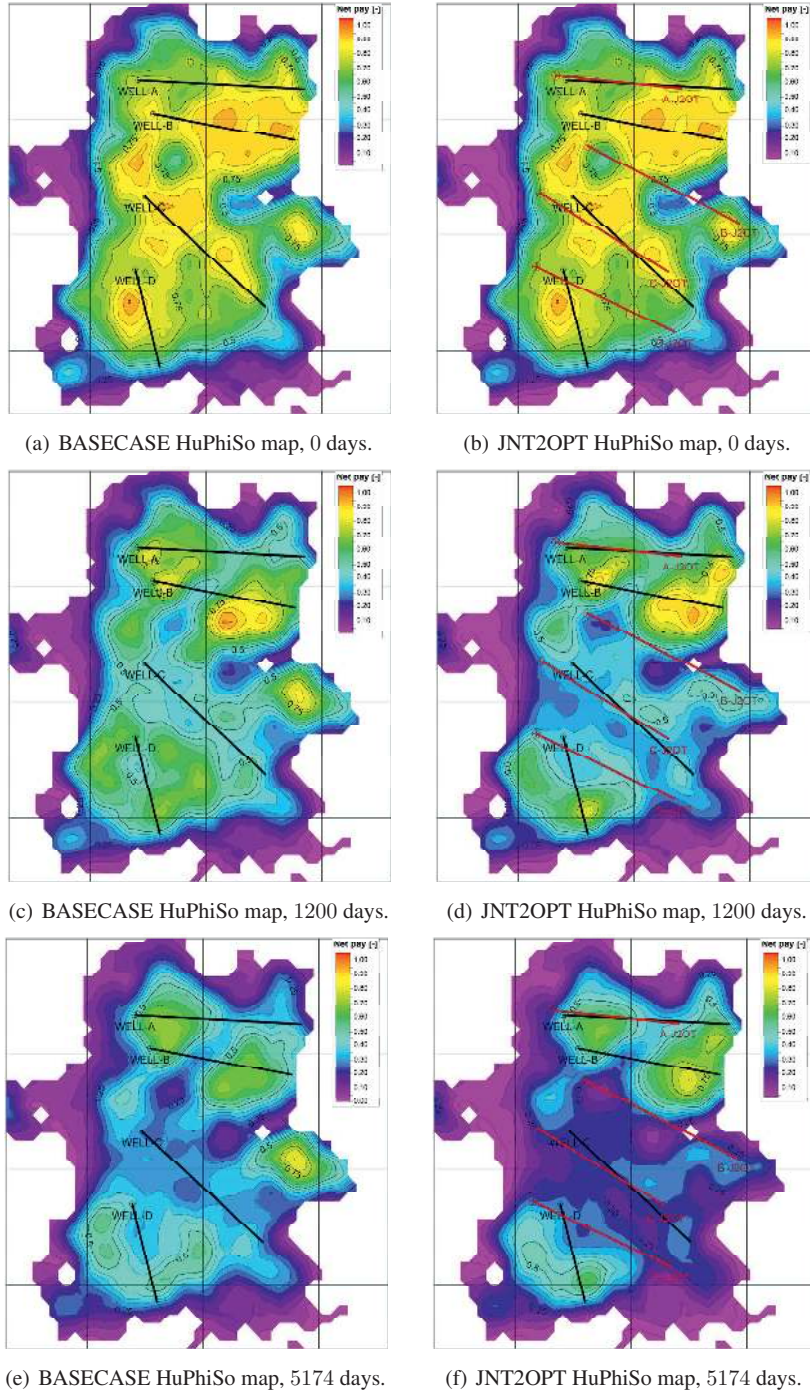


Figure 4.6: HuPhiSo oil saturation maps at three production times for BASECASE (left column) and JNT2OPT solution (right column). Scales have been normalized for confidentiality reasons.

4.1 Test results from solution cases

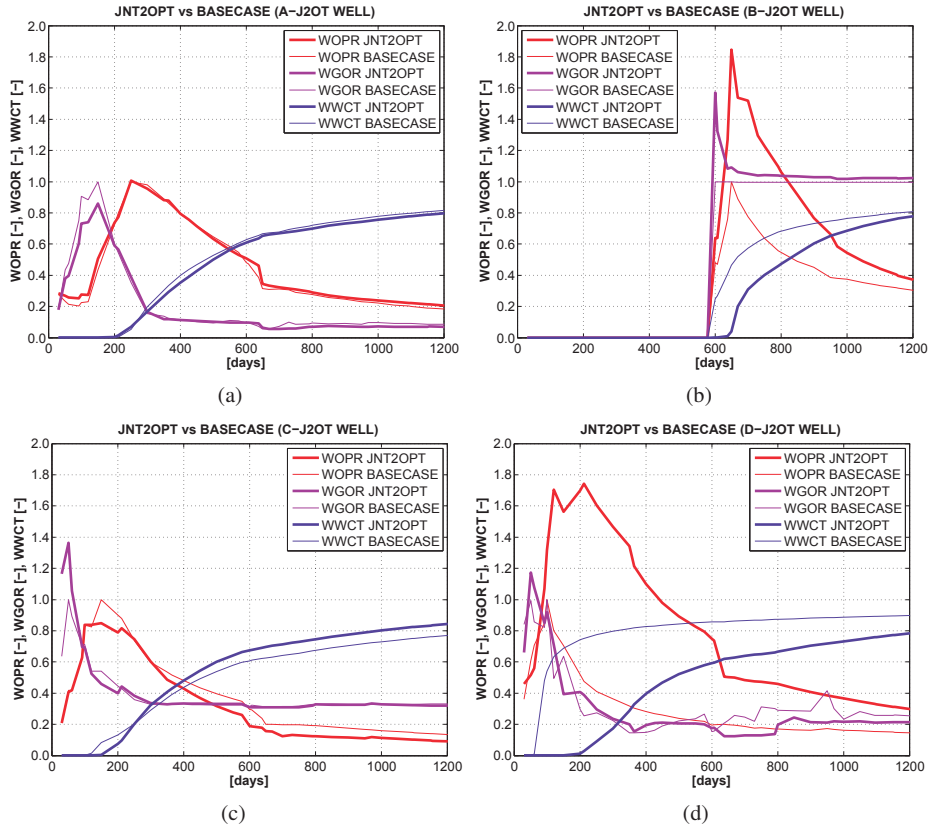


Figure 4.7: Production profiles for JNT2OPT solution and base case wells: Well oil production rates (WOPR), well gas-oil ratios (WGOR) and well water cuts (WWCT).

FXD2OPT solution

Compared to base case results, the FXD2OPT solution yields increases in field oil production total of 24% and 10% at 1200 days and after 5174 days of production, respectively (see tables 4.5 and 4.6 for full details). At 1200 days, increases in well cumulative oil production are 5%, 58%, -17% and 171% for wells A, B, C and D, respectively. At 5174 days, increases in well cumulative oil production are 0%, 30%, -28% and 103% for the same wells, respectively. HuPhiSo saturation maps at 0 days and at production times 1200 and 5174 days are given next. Following this we give individual descriptions of location and performance for the wells in the FXD2OPT solution.

A-F2OPT well. Unlike any of the A wells in the previous solutions, the A-F2OPT well has a counter-clockwise rotation compared its base case equivalent. This well is 75 meters shorter than its base case equivalent, and has a reservoir entry point positioned close to the west side of the accumulation boundary, and a heel moved north. The A-F2OPT well has the same peak oil production as its base case analog, but reaches this rate, and starts water production, at an earlier time.

B-F2OPT well. The B-F2OPT well does not reach the east accumulation lobe, but rather attains a trajectory that is shifted southward, parallel to the base case B well. This well is 108 meters shorter than that base case B well, but yields a substantially larger oil production rate and high GOR (the latter somewhat compensating for the diminished gas production from the A-F2OPT well).

C-F2OPT and D-F2OPT wells. Entry points for the C-F2OPT and D-F2OPT wells are moved eastward and positioned along the eastern accumulation border. The C-F2OPT well is 110 meters shorter while the D-F2OPT well is 484 meters longer than the corresponding C and D base case wells, respectively. Again we see the D-F2OPT well move away from the south-eastern region toward the central part of the accumulation, and a realignment of the C-F2OPT and D-F2OPT wells in a parallel manner. It is important to notice that while production profiles from the C-F2OPT well are similar to those of the base case C well, the oil production from the D-F2OPT well is substantially higher than its base case counterpart, and its water cut profile shows only a gradual increase.

4.1 Test results from solution cases

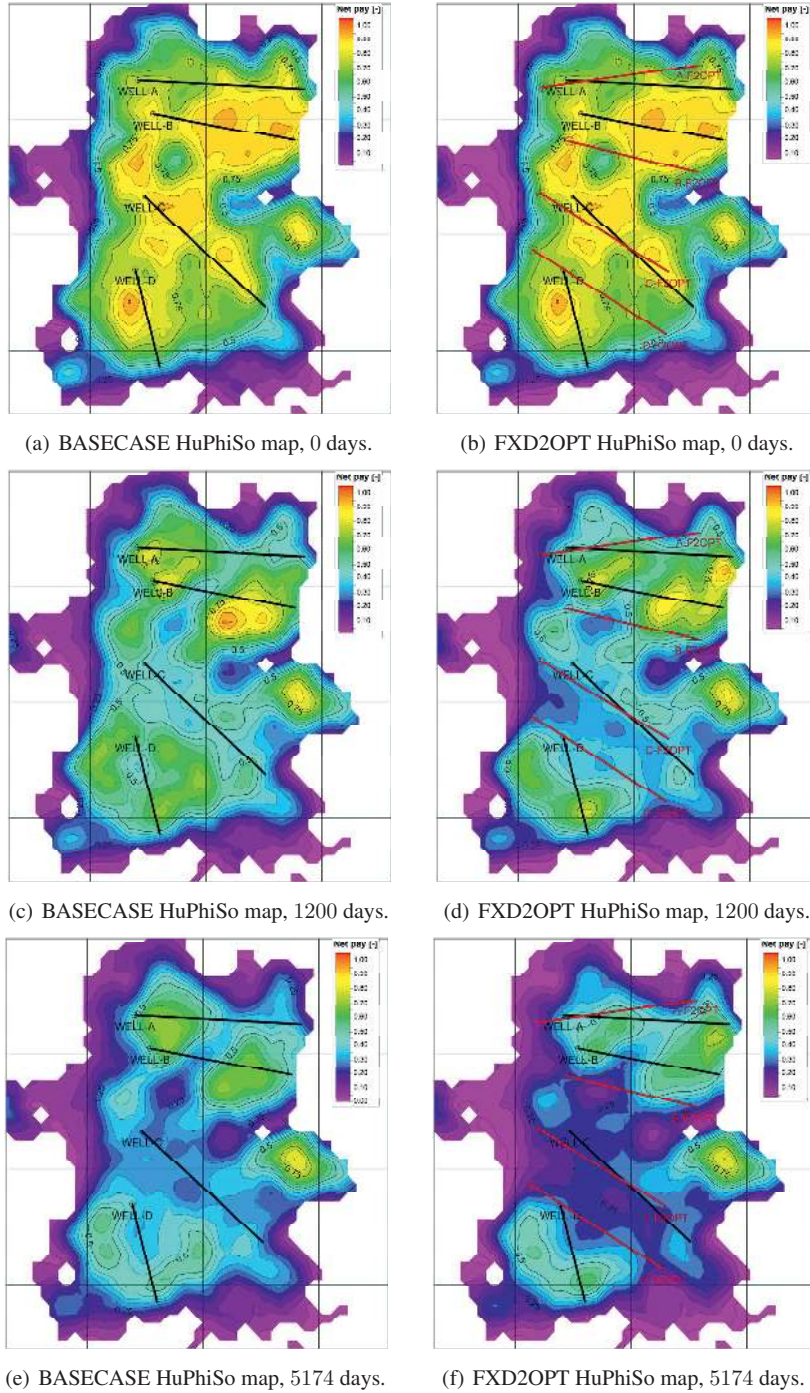


Figure 4.8: HuPhiSo oil saturation maps at three production times for BASECASE (left column) and FXD2OPT solution (right column). Scales have been normalized for confidentiality reasons.

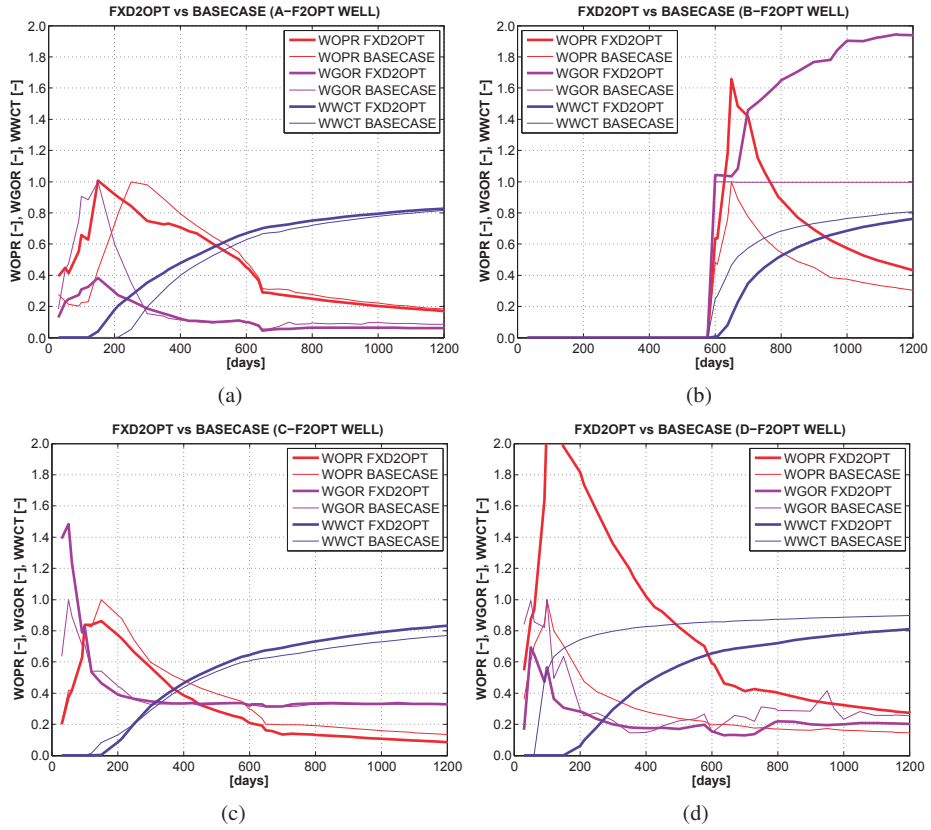


Figure 4.9: Production profiles for FXD2OPT solution and base case wells: Well oil production rates (WOPR), well gas–oil ratios (WGOR) and well water cuts (WWCT).

JNT2OPT2 solution

Compared to base case results, the JNT2OPT2 solution yields increases in field oil production total of 23% and 13% at 1200 days and after 5174 days of production, respectively (see tables 4.5 and 4.6 for full details). At 1200 days, increases in well cumulative oil production are 3%, 63%, -13% and 156% for wells A, B, C and D, respectively. At 5174 days, increases in well cumulative oil production are 10%, 25%, -24% and 101% for the same wells, respectively. HuPhiSo saturation maps at 0 days and at production times 1200 and 5174 days are given next. Following this we give individual descriptions of location and performance for the wells in the JNT2OPT2 solution.

A-J2OT2 and B-J2OT2 wells. The A well in the JNT2OPT2 solution, i.e., A-J2OT2, resembles the A well in previous joint solutions in that its trajectory is similar but shorter than its base case analog (1129 meters compared to 1439 meters for the base case A well). As the B wells in solution JNT2M1 and JNT2OPT, the toe of B-J2OT2 is also prominently placed on the eastern lobe of the reservoir.

In this case though, B-J2OT2 is shorter⁶ than other B wells from other solutions that also produce from the eastern lobe area, e.g., B-J2OT2 is 1142 meters compared to 1474 meters for B-J2OT in the JNT2OPT solution. The difference in length allows this well to have a reservoir entry point that is significantly more south, and to the center of the reservoir, compared to the base case B well. It also means production from this well is more focused on the eastern accumulation. Even though shorter than its base case analog, the B-J2OT2 well has a higher oil rate and a delayed water breakthrough. The A-J2OT2 well shows practically the same profiles as its base case counterpart.

C-J2OT2 and D-J2OT2 wells. Production profiles for the C-J2OT2 well are very similar to production profiles from its base case counterpart, even though this well has a substantial counter-clockwise rotation, and is 107 meters shorter. The D-J2OT2 well is 281 meters longer than the base case D well, and has a trajectory pointing south-east. This new orientation yields both a significant increase in oil rate production and substantial delay in water breakthrough.

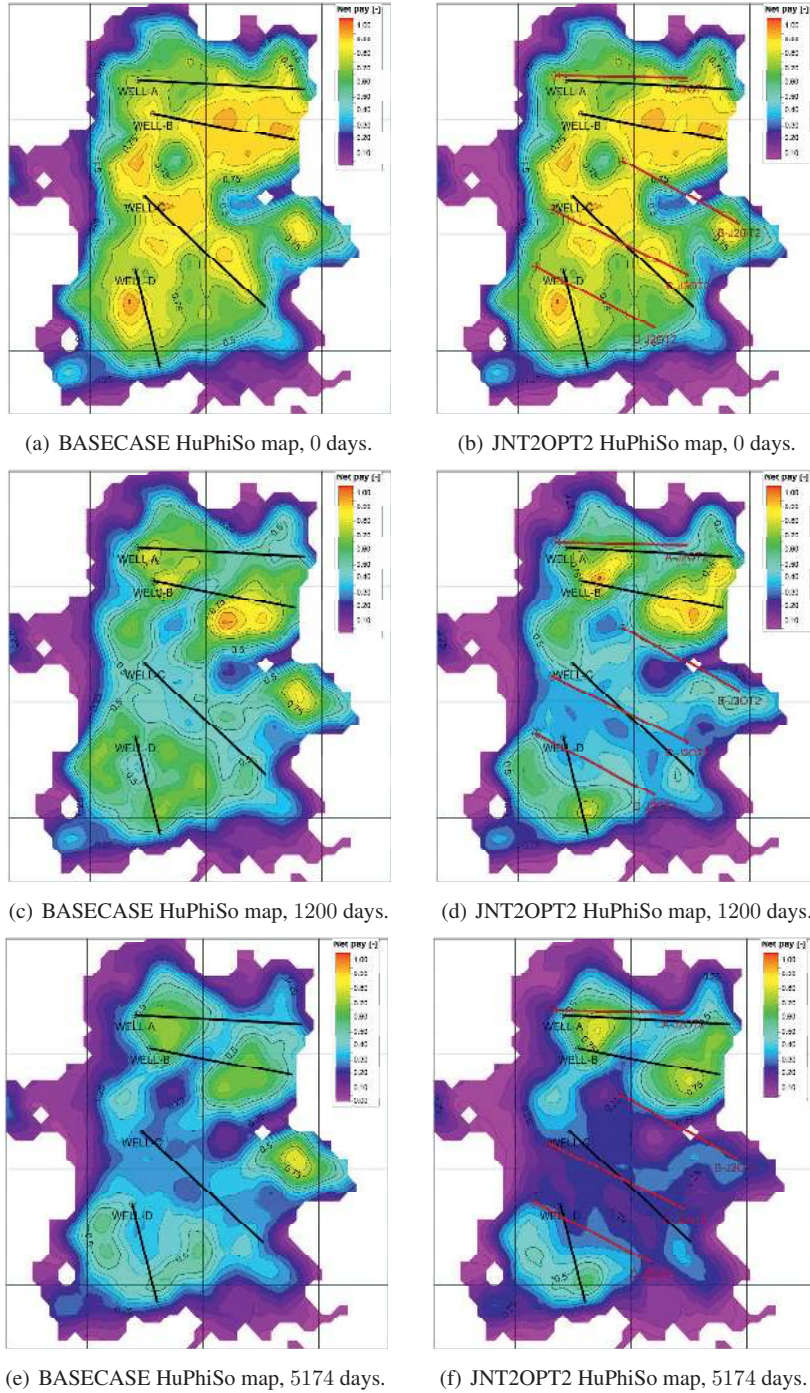


Figure 4.10: HuPhiSo oil saturation maps at three production times for BASECASE (left column) and JNT2OPT2 solution (right column). Scales have been normalized for confidentiality reasons.

4.1 Test results from solution cases

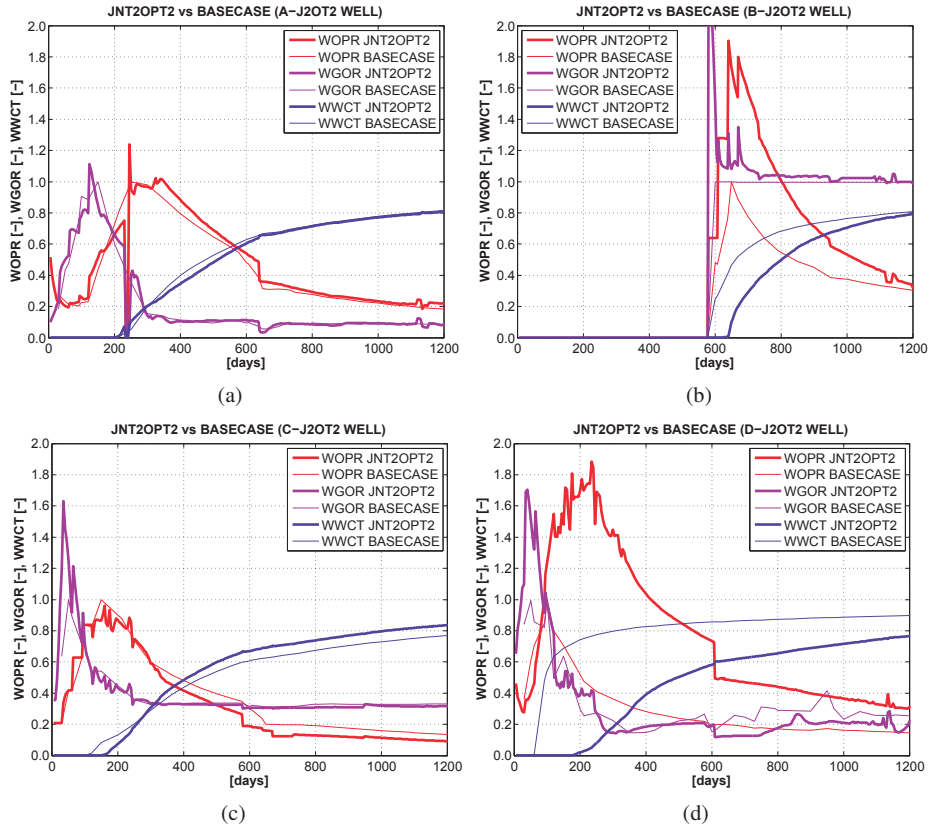


Figure 4.11: Production profiles for JNT2OPT2 solution and base case wells: Well oil production rates (WOPR), well gas-oil ratios (WGOR) and well water cuts (WWCT).


FXD2OPT2 solution

Compared to base case results, the FXD2OPT2 solution yields increases in field oil production total of 22% and 10% at 1200 days and after 5174 days of production, respectively (see tables 4.5 and 4.6 for full details). At 1200 days, increases in well cumulative oil production are 1%, 47%, -17% and 177% for wells A, B, C and D, respectively. At 5174 days, increases in well cumulative oil production are 3%, 21%, -28% and 107% for the same wells, respectively. HuPhiSo saturation maps at 0 days and at production times 1200 and 5174 days are given next. Following this we given individual descriptions of location and performance for the wells in the FXD2OPT2 solution.


A-F2OPT2 and B-F2OPT2 wells. We see that this solution, together with the other sequential solutions (FXD1M1, FXD2OPT, and FXD2CUT), and JNT2CUT, keep away from production of the somewhat isolated eastern lobe of the reservoir. For this solution we have that both the A-F2OPT2 and B-F2OPT2 wells have a similar configuration as in the FXD2OPT solutions. The A-F2OPT2 and B-F2OPT2 wells are 310 and 45 meters shorter than their base case counterparts, and we observe once again that a shorter A-F2OPT2 well yields similar production profiles as the longer base case A well.

C-F2OPT2 and D-F2OPT2 wells. As in the FXD2OPT solution, the C-F2OPT2 and D-F2OPT2 wells are aligned parallel to each other, though at a greater distance. In this configuration, the oil production rate from the C-F2OPT2 well is slightly lower compared to its base case analog. However, we see a substantial increase in oil rate and decrease in water cut for the D-F2OPT2 well, similar to the one observed in the production profiles for the F2OPT well in the FXD2OPT solution.


4.1 Test results from solution cases




(a) BASECASE HuPhiSo map, 0 days.




(b) FXD2OPT2 HuPhiSo map, 0 days.




(c) BASECASE HuPhiSo map, 1200 days.



(d) FXD2OPT2 HuPhiSo map, 1200 days.



(e) BASECASE HuPhiSo map, 5174 days.



(f) FXD2OPT2 HuPhiSo map, 5174 days.

Figure 4.12: HuPhiSo oil saturation maps at three production times for BASECASE (left column) and FXD2OPT2 solution (right column). Scales have been normalized for confidentiality reasons.

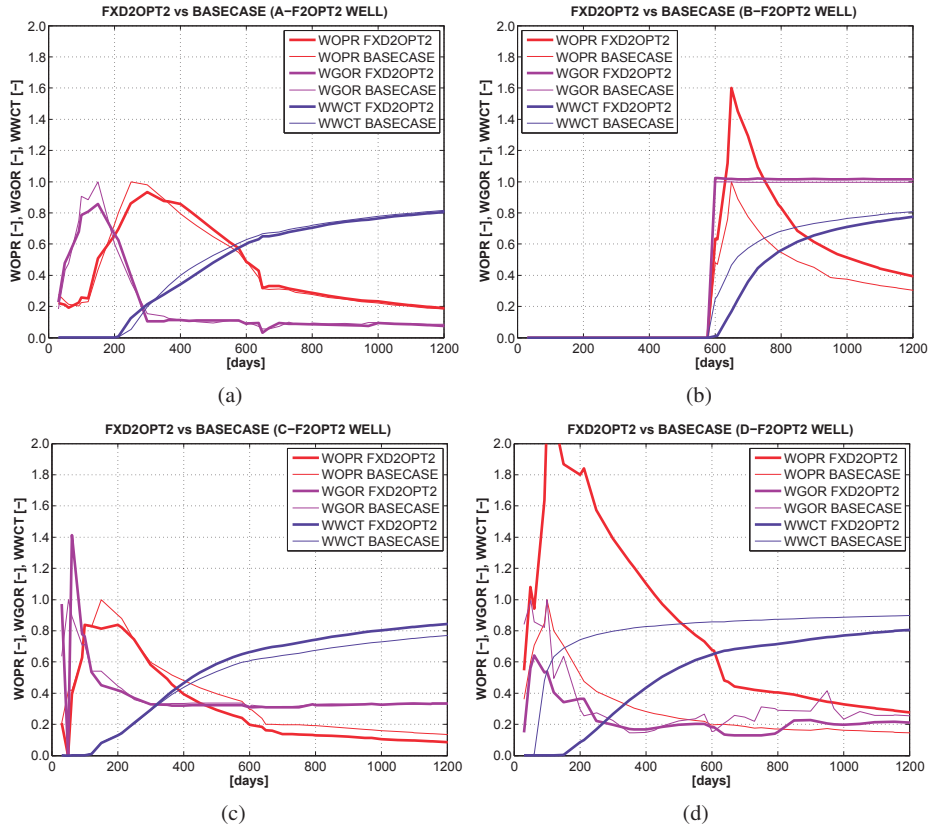


Figure 4.13: Production profiles for FXD2OPT2 solution and base case wells: Well oil production rates (WOPR), well gas–oil ratios (WGOR) and well water cuts (WWCT).

JNT2CUT solution

Compared to base case results, the JNT2CUT solution yields increases in field oil production total of 22% and 9% at 1200 days and after 5174 days of production, respectively (see tables 4.5 and 4.6 for full details). At 1200 days, increases in well cumulative oil production are 1%, 48%, -14% and 162% for wells A, B, C and D, respectively. At 5174 days, increases in well cumulative oil production are 3%, 19%, -25% and 96% for the same wells, respectively. HuPhiSo saturation maps at 0 days and at production times 1200 and 5174 days are given next. Following this we give individual descriptions of location and performance for the wells in the JNT2CUT solution.

A-J2CT and B-J2CT wells. Overall, there are only minor differences between the JNT2CUT and the previous FXD2OPT2 solution (see page 4.1.1). Since production profiles from this solution are similar to those from the FXD2OPT2 solution, we refer to the general descriptions given for the FXD2OPT2 solution (see page 132). As in the FXD2OPT2 solution, the A-J2CT well in the JNT2CUT solution is also 1129 meters, i.e., significantly shorter than its base case analog, by 310 meters. Still, the production profiles for this well closely match those from the base case well. Compared to the FXD2OPT2 solution, the B-J2CT well is 168 meters shorter (reduced at its heel-end), and slightly shifted southward, though with similar production profiles.

C-J2CT and D-J2CT wells. Compared to the FXD2OPT2 solution, both the C-J2CT and D-J2CT wells have a slight counter-clockwise rotation, and the heel of the C-J2CT well is positioned close to the eastern accumulation border. Compared to the C and D wells in the FXD2OPT2 solution, the heel of the D-J2CT well is moved slightly southward, while its toe is close to the area where the heel of the base case C well is located. Production profiles for this well, and the C-J2CT well, are similar to those for the FXD2OPT2 solution.

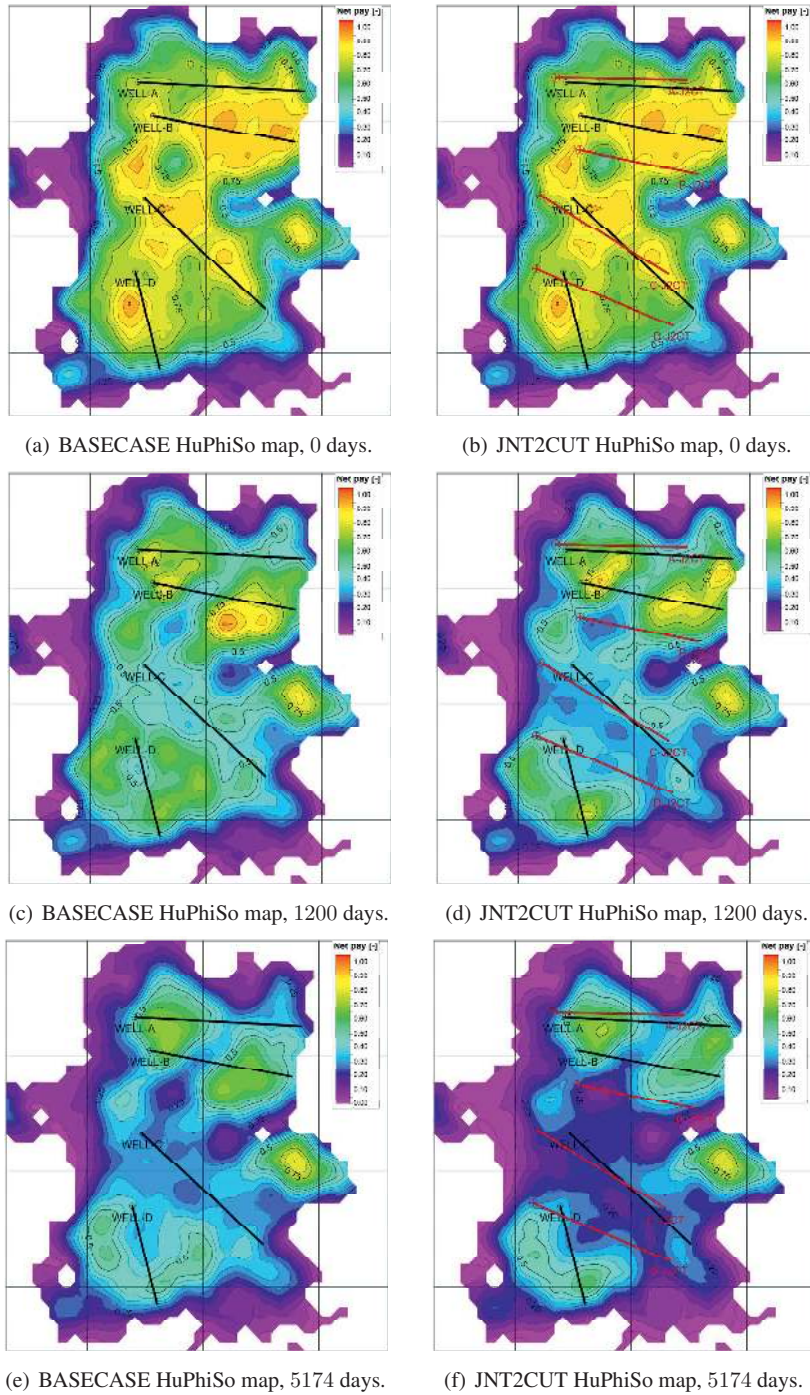


Figure 4.14: HuPhiSo oil saturation maps at three production times for BASECASE (left column) and JNT2CUT solution (right column). Scales have been normalized for confidentiality reasons.

4.1 Test results from solution cases

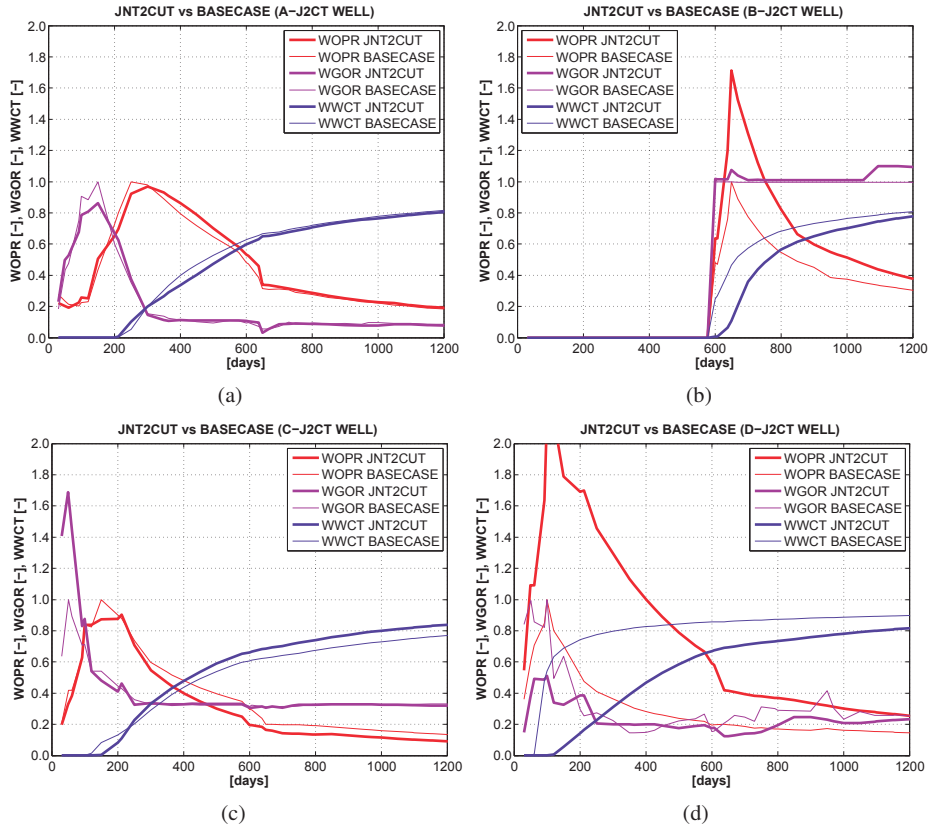


Figure 4.15: Production profiles for JNT2CUT solution and base case wells: Well oil production rates (WOPR), well gas-oil ratios (WGOR) and well water cuts (WWCT).

FXD2CUT solution


Compared to base case results, the FXD2CUT solution yields increases in field oil production total of 18% and 8% at 1200 days and after 5174 days of production, respectively (see tables 4.5 and 4.6 for full details). At 1200 days, increases in well cumulative oil production are 1%, 44%, -13% and 134% for wells A, B, C and D, respectively. At 5174 days, increases in well cumulative oil production are 3%, 20%, -22% and 82% for the same wells, respectively. HuPhiSo saturation maps at 0 days and at production times 1200 and 5174 days are given next. Following this we give individual descriptions of location and performance for the wells in the FXD2CUT solution.

A-F2CUT and B-F2CUT wells. As in the JNT2CUT solution, the A-F2CUT and B-F2CUT well trajectories and production profiles for the FXD2CUT solution are similar to those of the A and B wells in the FXD2OPT2 solution.


C-F2CUT well. The trajectory of the C-F2CUT well is almost the same as the trajectory of the base case C well, though with a 36 meter shorter well bore, and with the well slightly moved in the north-west direction. Production profiles for this well, compared to profiles for its base case analog, show similar evolutions, but with a slightly lower oil production rate, and a small increase in water cut.

D-F2CUT well. Compared to its base case analog, the D-F2CUT well is rotated counter-clockwise and forms a close to parallel configuration with the C-F2CUT well. The D-F2CUT well is 195 meters longer, and has a heel positioned slightly northward, compared to its base case analog. Its production profiles are roughly similar to those of the D well in the JNT2CUT solution, i.e., we observe a large increase in oil production rate in addition to a delayed water breakthrough time and a gradual, rather than a steep, increase in water cut.


4.1 Test results from solution cases



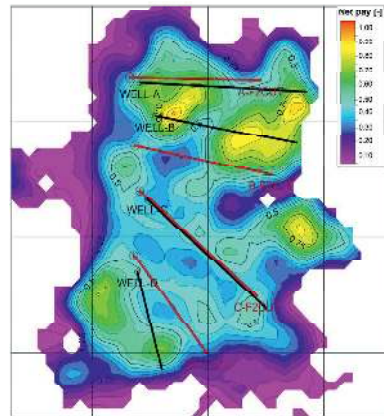
(a) BASECASE HuPhiSo map, 0 days.




(b) FXD2CUT HuPhiSo map, 0 days.




(c) BASECASE HuPhiSo map, 1200 days.



(d) FXD2CUT HuPhiSo map, 1200 days.



(e) BASECASE HuPhiSo map, 5174 days.



(f) FXD2CUT HuPhiSo map, 5174 days.

Figure 4.16: HuPhiSo oil saturation maps at three production times for BASECASE (left column) and FXD2CUT solution (right column). Scales have been normalized for confidentiality reasons.

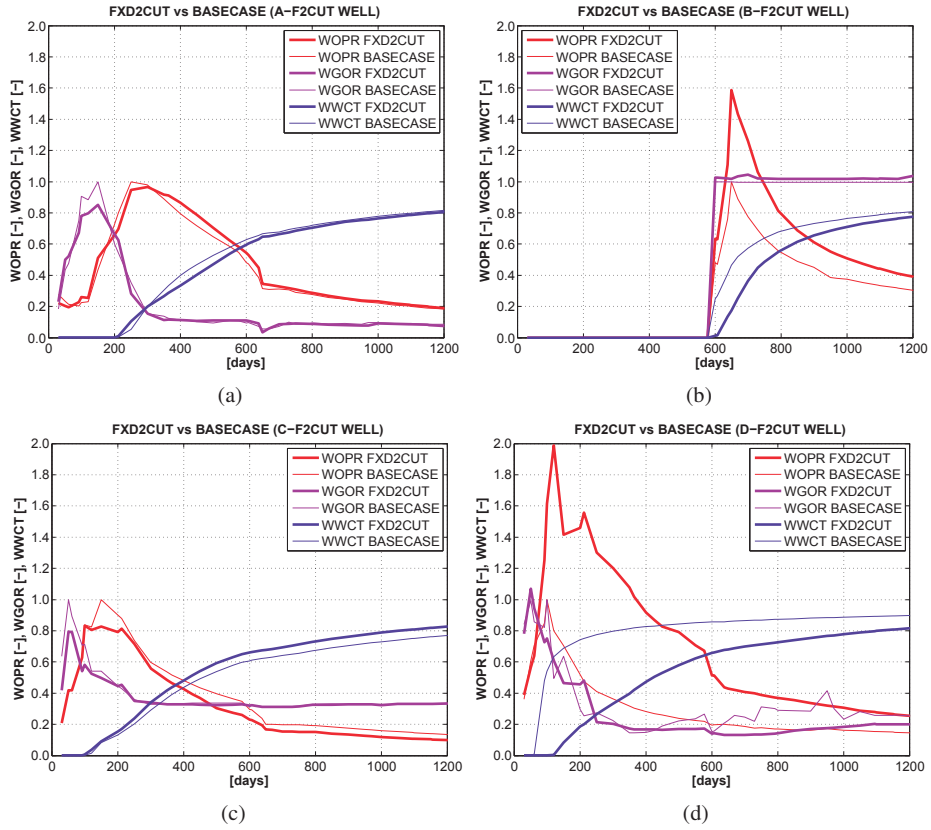


Figure 4.17: Production profiles for FXD2CUT solution and base case wells: Well oil production rates (WOPR), well gas–oil ratios (WGOR) and well water cuts (WWCT).

Summary of individual analysis of final well configurations

We have presented oil saturation maps and production profiles for eight different well configurations developed using our optimization framework. The results may be viewed from different perspectives, e.g., the results can be analyzed in terms of whole configurations, or particular wells can be singled out for individual treatment. In this summary we compare the individual well trajectories and production profiles across the different configurations and summarize the most interesting features.

- Several solutions (i.e., JNT2OPT, JNT2OPT2 and FXD2OPT2, JNT2CUT and FXD2CUT) arrive at A wells that have close to, or practically the same, positioning as the base case A well, but with a much shorter well bore. Still, these shorter wells have production profiles that closely match those of the base case A well.
- Those configurations that yield the greatest increases in field oil production total for both the 1200 day and the 5174 day time frames (i.e., JNT2M1, JNT2OPT and JNT2OPT2) have B wells that aggressively target the eastern lobe oil accumulation.
- Most configurations have only slight variations of the base case C well trajectory. In most cases the variations consist of modest rotations and/or shifts in position. Overall, these changes seem to have little effect on performance, and production profiles from the different solution C wells closely match those of the base case C well.
- In all solutions, except one (FXD1M1) where the increase is moderate, the oil production rates obtained from solution D wells are significantly larger than the oil production rate from the base case D well.
- Except for the FXD1M1 solution, all other solutions have arrived at D well trajectories that are significantly rotated counter-clockwise compared to the original base case D well trajectory, and that align the D well bore pointing toward the south-east direction.
- In all cases, we confirm an expected increase in base case D well length since this well has a particularly short initial length of 874 meters. Overall, the lengths of the solution D wells are 200 – 300 meters longer than the length of the base case D well. For the remaining wells, the proposed trajectories have many different lengths compared to base case well lengths, with no clear pattern (see Table 4.3).
- Well depths show very little variation from base case values, except for the D well position, where most of the solution D wells suggest a deeper positioning of the well by up to 4 meters (see Table 4.4).

The next section presents the increases in cumulative oil production from all the well placement solutions for both the entire reservoir and for the individual wells. Results are presented for both the 1200 day and the 5174 day production time frame.

4.1.2 Collective analysis: Total field and well oil production values

In this section we give a collective presentation of the total field and well oil production values for the various well placement solutions. We present these results at both the reduced production time frame of 1200 days, and at the field model production time frame of 5174 days.

Tables 4.5 and 4.6 summarize the changes in cumulative oil production for each solution at 1200 and after 5174 days of production, respectively. The tables show the percentage changes in total oil production for each well and for the entire field (our optimization objective has been to maximize total field oil production). The main percentage change in total oil production for each well is given relative to the cumulative oil produced by its counterpart base case well. Next to this value, an additional percentage change in cumulative oil production is given in parenthesis. This change is equal to the main percentage value, but normalized with respect to the increase in total field oil production, i.e., the percentage values in parenthesis show how much each well contributes to the total increase in cumulative field oil production.

Profiles for field and well oil production totals are also plotted for each of the two production time frames. Field and well oil production profiles for the 1200 day time frame are given in Figures 4.18 and 4.19, respectively. Figures 4.20 and 4.21 show field and well oil production profiles for the time period between 1200 and 5174 days. As discussed in the beginning of Section 4.1, the results shown in this section are obtained using the original control strategy (see Section 3.5 for further discussion of other control strategies). In tables 4.5 and 4.6 we refer to the original production strategy as x_c^S .

We present the results for the 1200 day production time frame first. A discussion of the results for the 5174 day production time frame starts on page 145. In Section 4.1.3 we present further analysis treating the trade-off between increased oil recovery and changes in total well drain length.

Field and well recovery: 1200 days production time frame

Table 4.5 shows the differences in final field and well cumulative oil production for each solution compared to base case values after a 1200 day production time frame. Figures 4.18 and 4.19 show the corresponding field and well cumulative oil profiles, respectively. These results have been discussed previously in terms of performance, e.g., between joint and sequential solutions and for different non-linear constraint handling techniques (see Section 3.5, page 91). Here we limit our discussion to the obtained differences in both field and well cumulative results relative to base case values.

Field-wise comparison of cumulative oil after 1200 days. At this time frame, we see from Table 4.5 that the highest increase in field oil production (ΔFOPT), is 25.5% obtained by the JNT2M1 solution, closely followed by the JNT2OPT solution with an increase of 24.9%. The JNT1M1 and FXD2OPT solutions are close behind with increases of 23.6% and 23.5%, respectively. The FXD2M1 and FXD2CUT solutions offer relatively poor increases with 13.3% and 17.6%, respectively.

The total field oil production profiles for this time frame (normalized by base case FOPT) are shown in Figures 4.18. In this figure we observe the increases are realized

Table 4.5: Percentage increases in field and well oil production total for a 1200 day production horizon, ΔFOPT and ΔWOPT , respectively. Field increases and well changes are given for all solutions relative to base case total oil production values. Two values are given for each well. The main value for a well is the percentage change in cumulative oil production relative to its base case counterpart well. The second value in parenthesis is the same value but in addition normalized with respect to the field increase in total oil production. The value in parenthesis thus shows the individual well contribution to the overall field increase. Notice a simpler table (Table 3.8) showing only the main values in this table was given as a prelude at the end of Section 3.5.

Solution	ECLIPSE ₁₂₀₀ ($\mathbf{x}_p^*, \mathbf{x}_c^S$)									
	ΔFOPT		ΔWOPT ($\Delta\text{WOPT} \cdot (\frac{\text{WOPT}}{\text{FOPT}})_{\text{Solution}}$)							
			WL.-A		WL.-B		WL.-C		WL.-D	
	[%]	[%]	[%]	[%]	[%]	[%]	[%]	[%]	[%]	[%]
FXD1M1	13.3	2.4	(0.8)	56.6	(5.9)	1.2	(0.5)	44.5	(6.1)	
FXD2OPT	23.5	4.5	(1.4)	57.9	(6.1)	-16.7	(-7.5)	171.4	(23.5)	
FXD2OPT2	21.9	0.9	(0.3)	46.6	(4.9)	-16.9	(-7.6)	177.2	(24.3)	
FXD2CUT	17.6	0.9	(0.3)	44.0	(4.6)	-12.5	(-5.6)	133.7	(18.4)	
JNT2M1	25.5	-0.1	(-0.0)	99.8	(10.5)	-12.5	(-5.6)	150.6	(20.7)	
JNT1M1	23.6	-3.0	(-0.9)	89.5	(9.4)	-12.3	(-5.6)	150.5	(20.7)	
JNT2OPT	24.9	3.2	(1.0)	67.1	(7.0)	-13.5	(-6.1)	167.5	(23.0)	
JNT2OPT2	23.3	3.0	(0.9)	63.1	(6.6)	-12.7	(-5.7)	156.0	(21.4)	
JNT2CUT	21.5	0.9	(0.3)	47.6	(5.0)	-13.6	(-6.1)	162.1	(22.3)	

throughout the span of the time frame. In particular, we see that two of the best solutions, JNT2M1 and FXD2OPT, show substantial increase in FOPT already early in production (the JNT2M1 curve somewhat overlaps the curve of the FXD2OPT solution during these early production times).

Well-wise comparison of cumulative oil after 1200 days. Though several final results are similar when seen field-wise, how much each well contributes to the total increase varies substantially between solutions. In general, we have that field results from the four best solutions (in decreasing order: JNT2M1, JNT2OPT, JNT1M1 and FXD2OPT), have substantial contributions from both the B and D wells.

The contributions from the B and D wells in these solutions have ranges of ($\min = 57.9, \max = 99.8$)%, and ($150.6, 177.2$)%, respectively. Whenever the increases




Figure 4.18: FOPT for all solutions, 1200 days production time.

from either the B or D wells, or both, fall short, we obtain solutions with a significantly lower FOPT increase, e.g., FXD2CUT and FXD2M1. This is evident when we observe the columns for the A and C wells. For all solutions, increases from the A wells are in the low range of $(-3.0, 4.5)\%$, while, except for the FXD2M1 solution, all C well contributions are negative within a range of $(-16.9, -12.3)\%$.

From the results in parenthesis in Table 4.5, we see that the most substantial contributions to the total FOPT increase come from the D wells. We stipulate that the general increase from the D wells is derived, up to a point, from the repositioning of most of the solution D wells within the original drainage area of the C well (as can be seen in several of the maps in Section 4.1.1, the solution D wells often intrude on the drainage area of the base case C well). If true, this overtake could possibly be one of the main reasons behind the negative contributions from the C wells (though further analyzes that detail specific fluid flow patterns would be required to determine this effect). After the D wells, the most important contributions come from the B wells, while the A wells have marginal contributions to overall FOPT. We discuss the B and D well contributions in further detail below.

For the B wells, it is important to notice that the substantial increases attained by these wells do not, to any large degree, diminish the production from their neighboring A wells. This appears to be the case for all B well trajectories, even those that do not specifically target the eastern oil accumulation (and that are generally closer to the A well trajectories). As mentioned, this cannot be said about the D wells, which do appear to significantly influence the production of their adjacent C wells. Based on our results, it appears the D well contributions to overall increase in recovery, though substantial (in the order of 20%; see parenthesis column for D wells in Table 4.5), is partly based on a diminished production from the C wells (somewhat below 10%). In the end, one has to choose a configuration of C and D wells that combined yields an increase in oil production. But which configuration of well trajectories to choose, i.e., how the total increase should be distributed between the wells is a decision that would have to be made at the well strategy

4.1 Test results from solution cases




Figure 4.19: WOPT for all solutions, from 0 to 1200 days production time.

level (taking into account other factors such as the level of uncertainty surrounding the individual drainage areas of the wells, e.g., with respect to structural uncertainty).

Figure 4.19 shows the total oil production for each well (WOPT; all solution well profiles are normalized relative to their corresponding base case well). For nearly all solutions, we observe the substantial increase in recovery from the B and D wells. In particular, we notice these wells start producing more oil right from their start of operation. It is worth noting that for the A well, FXD2OPT is the only solution that holds a larger WOPT for the entire time frame (for this solution, this well produces 4.5% more oil than its base case counterpart).

Field and well recovery: 5174 days production time frame

Table 4.6 shows simulation case results for each solution using a 5174 day production horizon. As for the 1200 day time frame, both field and well-wise results are presented. Figure 4.20 shows the field oil production total (FOPT) for all solutions against base case results. Similarly, for all solutions, Figure 4.21 presents the oil production total corresponding to each well (WOPT). To accentuate the difference between oil production

Table 4.6: Percentage increases in field and well oil production total for a 5174 day production horizon, ΔFOPT and ΔWOPT , respectively. Field increases and well changes are given for all solutions relative to base case total oil production values. As before, two values are given for each well. The main value for a well is the percentage change in cumulative oil production relative to its base case counterpart well. The second value in parenthesis is the same value but in addition normalized with respect to the field increase in total oil production. The value in parenthesis thus shows the individual well contribution to the overall field increase.

Solution	ECLIPSE ₅₁₇₄ ($\mathbf{x}_p^*, \mathbf{x}_c^S$)									
	ΔFOPT		ΔWOPT ($\Delta\text{WOPT} \cdot (\frac{\text{WOPT}}{\text{FOPT}})_{\text{Solution}}$)							
			WL.-A		WL.-B		WL.-C		WL.-D	
	[%]	[%]	[%]	[%]	[%]	[%]	[%]	[%]	[%]	[%]
FXD1M1	8.3	3.0	(0.9)	18.4	(2.9)	4.4	(1.8)	17.4	(2.7)	
FXD2OPT	9.7	-0.2	(-0.1)	29.7	(4.7)	-27.5	(-11.2)	103.2	(16.3)	
FXD2OPT2	9.6	3.2	(0.9)	20.7	(3.2)	-28.1	(-11.4)	106.7	(16.9)	
FXD2CUT	8.2	3.3	(0.9)	19.5	(3.1)	-21.5	(-8.7)	81.8	(12.9)	
JNT2M1	13.0	-5.4	(-1.5)	61.3	(9.6)	-21.4	(-8.7)	85.9	(13.6)	
JNT1M1	10.8	-13.0	(-3.6)	60.6	(9.5)	-21.4	(-8.7)	85.9	(13.6)	
JNT2OPT	13.0	9.2	(2.6)	28.1	(4.4)	-25.6	(-10.4)	103.9	(16.4)	
JNT2OPT2	12.8	10.3	(2.9)	24.6	(3.9)	-24.3	(-9.8)	100.5	(15.9)	
JNT2CUT	8.9	2.9	(0.8)	18.8	(2.9)	-24.5	(-9.9)	95.7	(15.1)	

profiles, we plot all graphs in Figures 4.20 and 4.21 from 1200 to 5174 days only, instead of for the complete 5174 day horizon. (Notice also that we have normalized all FOPT and WOPT solution graphs with respect to the corresponding final FOPT and WOPT values from the base case.)

Field-wise comparison of cumulative oil after 5174 days. For the 5174 day time frame, we see from Table 4.6 that both the JNT2M1 and JNT2OPT solutions yield the highest increases in field oil production (ΔFOPT) at 13% over base case. A close second is the JNT2OPT2 solution with an increase in FOPT of 12.8%. We also notice that, in general, the joint solutions yield higher FOPT increases than the sequential solutions.

Testing the solutions for the 5174 day production horizon gives us an indication about how fruitful the time-frame design decision has been, i.e., the decision of approximating the original time frame with a much smaller production horizon of 1200 days for opti-




Figure 4.20: FOPT for all solutions, 5174 days production time.

mization purposes. The basis for the approximation was that the majority of oil, for this case, is produced mainly during the first few years of production (the motivation for the approximation itself was to reduce simulation runtime to allow for hundreds or thousands of simulations to be launched during optimization). Because of the approximation, solutions from the optimization routine that are optimal with respect to the shorter horizon, are expected to be suboptimal for the larger time frame. This is in agreement with the general decline seen in FOPT values when going from the 1200 day results, shown in Table 4.5, to the results obtained using the 5174 day time frame, shown in Table 4.6. Our purpose here is to report on the overall magnitude of this general decline, and to document those cases, both field and well-wise, that still retain a significant part of the gain achieved in the 1200 day horizon when tested on the field case production scenario.

Overall, we observe that for the different solutions, the increase in FOPT at 5174 days is roughly 50% of what it was at 1200 days. However, the 5174 day increases are still significant, in the order of 10%, which suggests the approximation has been successful in focusing on main oil production⁷. A secondary point in this regard is that, even for the larger horizon, all of the joint solutions still outperform their sequential counterparts. A more detailed discussion of table results is given below.

Final FOPT increases for the different solutions shown in Table 4.6 range from 8.2% to 13%, with the best solutions being JNT2M1 and JNT2OPT (both with FOPT increases of 13%). Interestingly, even though these two solutions achieve the same increase in FOPT, there are clear differences in their individual well contributions. (The discussion also applies to JNT1M1 since its well configuration is very similar to the JNT2M1 solution.) For example, in Table 4.6 we see that the A well in the JNT2M1 solution has a decrease in WOPT of about 5%, while the A well in the JNT2OPT solution has an increase in WOPT of about 9%. Furthermore, the JNT2M1 solution B well has an increase in WOPT of about 61%, while the B well in the JNT2OPT solution only achieves an increase in WOPT of 28%. This is noteworthy because both these solutions have B wells that target the relatively isolated eastern lobe area of the reservoir, and their A well trajectories are




Figure 4.21: WOPT for all solutions, from 1200 to 5174 days production time.

quite similar (see Figure 4.2 on page 118 and Figure 4.6 on page 124 for JNT2M1 and JNT2OPT well configurations, respectively).

Taken together, the contributions from the A and B wells to the overall FOPT increase are about 8% and 7% for the JNT2M1 and JNT2OPT solutions, respectively. Similarly, the sum of the contributions from the C and D wells to the overall FOPT are close to 5% and 6%, respectively. Thus, even though the A and B wells for the JNT2M1 and JNT2OPT solutions have the same target regions, and the sum of their contributions is about the same, their relative positioning and length yield different individual performances. In terms of applicability of solutions, it is advantageous to have a set of solutions with varying configurations that all still achieve the same, or around the same, high FOPT increase in objective. This variation gives the operator the opportunity to select the configuration that best suits a broader well strategy that may include other considerations than those taken into account during the optimization process.

In Figure 4.20 we notice that at about 2000 days the cumulative oil curves for the three best solutions (JNT2M1, JNT2OPT and JNT2OPT2) continue their gradual increase while the increases for the remaining curves drop slightly. A continuous FOPT increase is an important trait given the final production time frame is not necessarily fixed, and

4.1 Test results from solution cases




Figure 4.22: WBHP for all solutions, from 0 to 5174 days production time.

a production scenario with a shorter horizon (though still several times larger than 1200 days) may eventually be designed by the operator.

Figure 4.21 shows the oil production totals for each of the wells. Throughout the discussion of the 5174 day data, it is important to remember that our main focus is on the effect of applying solutions created for a 1200 day production time horizon on a larger time frame. We are therefore focused on topics such as whether there is degradation in cumulative increases, and how the general production curves develop over the larger time frame.

In the following we compare the 5174 day total oil production curves shown in Figure 4.21, with the total oil production curves for the 1200 day horizon shown in Figure 4.19. (In comparing the two figures notice that the scales of the individual graphs are often different.) For the A wells, we notice a significant divergence in curves in the 5174 day data in Figure 4.21. From the individual production profiles shown for each simulation case in Section 4.1.1, we remember that most of the solution A wells give oil rates and production ratios very similar to those of the base case A well over a 1200 day production horizon. In Figure 4.21(a) we see that after 1200 days the total production curves for some of these solutions (notably FXD1M1 and JNT2M1) fall below base case values,

i.e., tail–production oil rates for the A well in these solutions go from being similar to becoming lower than their base case counterparts over the 5174 day time frame. However, the opposite is true for the A well in the JNT2OPT2 and the JNT2OPT solutions. For these A wells, total oil production increase is of about 10% over the 5174 day production horizon, even though these wells are at least 300 meters shorter than the base case A well. (Recall these increases are presented as final well production values in Table 4.6.) In summary, well A results shown in Figure 4.21(a) exemplify the difficulty in efficiently scaling optimization gains. These results show how well performance profiles that largely converge within a short range when tested over the intended time horizon (from -3 to 4.5% compared to base case), can significantly diverge when implemented in a production scenario using a larger time frame (from -13 to 10.5% compared to base case). Still, within this spread of solutions, we have several A well trajectories that yield substantial increases in performance over their base case counterpart (as mentioned, the JNT2OPT2 and JNT2OPT A well trajectories). Moreover, we will see below that other wells in the configuration do retain a significant portion of their increases over the extended time frame.

One of the main assumptions for using a reduced–time–frame approximation in our optimization was that the production curves for the 1200 day production scenario would largely continue their increase (regardless of whether they are lower or higher than their base case counterparts) over the extended production horizon. For the B wells the assumption holds for most of the curves, with the exception of FXD1M1 that decreases, and JNT2M1 (and JNT1M1) that yield positive increases. Furthermore, while C well curves show clear decreases in rates, total oil production curves for solution D wells have a much more gradual progression throughout the greater time horizon.

In summary, the discussion in this section has sought to analyze the performance of solution well configurations on the 5174 day field case production scenario. Collectively, we see scattering of oil production totals for the solutions A wells, a decrease for most solution C wells, and a significant retainment of optimization gains for the B and the D wells. Finally, interesting well trajectories or whole configurations may be selected for further study based on individual and case performance data presented in previous sections.

Corresponding well bottom–hole pressures. In Figure 4.22 we present the bottom–hole pressures for all wells in each of the simulation cases presented in this section. These pressures are presented to confirm that all solution wells for each simulation case ultimately operate within specified bounds. In this respect, we remind the reader that all simulation cases in this section have been run using the original production parameters summarized in tables 4.1 and 4.2. Bottom–hole pressures presented in Figure 4.22 are thus the result of solution well trajectories running base case operational parameters (e.g., minimum tubing–head pressures) and well–control targets/limits. As in the base case, specific lift gas and group control settings are determined by the simulator during runtime, following the parametrization given in Table 4.2. Solution C well pressures deviate by less than a few bars above or below the bottom–hole pressure of the base case C well. For the remaining A, B and D wells, all solution bottom–hole pressures are, at a maximum, either above or below the base case pressure by roughly 5 bars. We note, however, that even rel-

atively small-to-moderate variations in bottom-hole pressure can have significant impact on well production rates, e.g., in wells with high production index.

Summary of collective analysis for total field oil production values

A collective presentation of the total field and well oil production values for the various well placement solutions has been given in this section. Results at both the reduced production time frame of 1200 days, and at the field model production time frame of 5174 days have been presented. Below we summarize the main findings from this analysis and discussion.

- For the 1200 day production time frame, the JNT2M1 well placement solution yields the greatest increase in field oil production with a 25.5% increase in FOPT over base case.
- For the same time frame, solutions JNT2OPT, JNT1M1 and FXD2OPT are close behind with increases in FOPT of 24.9%, 23.6% and 23.5%, respectively.
- Though developed using the reduced time frame, solution well configurations still yield significant increases when tested on simulation cases running the substantially larger field case production horizon of 5174 days.
- For the 5174 day production time frame, both the JNT2M1 and the JNT2OPT solutions yield the highest increases in field oil production at 13% over base case. A close second is the JNT2OPT2 solution with an increase in FOPT of 12.8%.
- Several of the well placement solutions have increases in FOPT close to the highest achieved value, even though the individual contributions from the wells in these solutions vary substantially due to different trajectories and lengths. A greater set of high-performance solutions allows the operator to choose the well configuration that best fits the broader well strategy of the development plan.

4.1.3 Increases in FOPT versus changes in well length

In this section we study the correlation between well length changes due to new well trajectories, and the increases in total oil production obtained from these solutions. For the development of the Martin Linge oil reservoir we expect longer well bores to yield higher oil production rates. However, longer well bores are often more challenging to drill because of the uncertainty in reservoir knowledge, e.g., with respect to the extent of reservoir sands and the existence and layout of faults, and, they can be more difficult to operate due to well-performance issues such as well bore pressure losses caused by fluid flow friction. Taken together, these type of factors restrict how long wells can be within a planned well configuration, and are the main motivation behind our well-length constraint definition (see Section 3.5). In this section our concern is solely on the relationship between the well length parameter and production, even though other parameters e.g., inter-well distance and general well orientation towards platform location also have significant influence on final well configuration and ultimate recovery (and have been taken into account during the well placement optimization procedure). In the following, we discuss how the different solution well lengths vary relative to base case well lengths, and how these well-length differences correlate to the increases in total oil production achieved by the solution wells. We begin by describing the compromise between well-length and production, and which data we plot to study this trade-off.

Trade-off between well length changes and oil production totals. The general well-length versus production trade-off arises when we try to make wells longer, to increase drainage area and maximize recovery, but we also require well configurations to be practical. By practical we mean well configurations that have various design features that make them less likely to encounter problems during drilling and operation. For reference purposes, we consider the base case well configuration for the development of the Martin Linge oil reservoir to be a practical configuration consisting of wells with reasonable lengths. The proposal of new well trajectories establishes a trade-off between the potential increment in total oil production and the changes to base case well lengths required to realize this increment. In this section we study this trade-off based on the relationship between the following two points of information. The first point of information regards the lengths of the wells at each of the solutions and the base case. The second point regards the increases in well and field oil production total obtained from the different solutions compared to base case values. Information for the first point is obtained by simply computing the length difference between each solution well and its corresponding base case well. Additionally, for each solution, we also find the difference between the sum of all well lengths in that solution, and the total length of all base case wells. The second information point consists of the changes (relative to base case) in total oil production for each well, and collectively for each solution, as presented in tables 4.5 and 4.6, in Section 4.1.2.

Field and well-wise results. We present both field and well-wise results for the well-length versus production trade-off. For both the field and well-wise results, we present data for the 1200 day and the 5174 day production horizons (see page 108). For the field-wise results, we plot the production increases against the sum of all well lengths in a




Figure 4.23: Increase in field oil recovery compared to base case values (ΔFOPT), plotted against the difference in total drainage length (wells A, B, C, and D) computed between each solution and corresponding base case wells.

solution, i.e., we add together the lengths that correspond to wells A, B, C, and D, respectively.

Field results: FOPT increase vs. total well length increase

Figure 4.23 shows the increase in field oil recovery (FOPT) against the difference in total drainage length (using wells A, B, C, and D) for all solutions. Regular markers correspond to FOPT increases after 1200 days of production, while markers with crosses correspond to FOPT increases after 5174 days of production. The red line and arrow signify the trend of FOPT increases decreasing when going from the 1200 day to the 5174 day production horizon, as discussed in Section 4.1.2. The vertical black line represents the break–even point where the sum of base case well lengths is equal to the sum of solution well lengths. (Notice joint and sequential solutions are no longer grouped together, rather, solutions are now organized as joint–sequential pairs.)

We see from Figure 4.23 that solutions JNT2CUT and FXD2CUT are the only solutions with total drainage length less than base case (-4.7% and -4% for JNT2CUT and FXD2CUT, respectively). After 5174 days, these solutions yield roughly a 10% FOPT increase (see tables 4.5 and 4.6 for exact values). The FXD2OPT2 solution has only a .4% increase in total drainage length compared to base case, while yielding FOPT increases of more than 20% and 10% for the 1200 day and the 5174 day production time frames, respectively. Solutions JNT2OPT2, JNT2OPT and FXD2OPT yield FOPT increases in the neighborhood of 25% after 1200 days of production, and of around 12% for the 5174 day production horizon. These substantial FOPT increases are obtained by the solutions using only a roughly 4% longer total well bore drain length than the base case. This is about the same as drilling 200 meters more drain in the reservoir, distributed over four wells.

Finally, solutions JNT2M1, JNT1M1 and FXD1M1 in Figure 4.23 yield FOPT increases of roughly 25% and 10% for the 1200 day and 5174 day production horizons. However, these increases require a roughly 15% increase in total drainage lengths compared to the base case solution. (Recall that these solutions enforce well-length projections at the end of the optimization, and not during the routine as the other well-length constraint implementations; see Section 3.5 for further details.) For further reference, we notice, based on the individual well lengths for all solutions presented in Table 4.3, that the B and D wells have the greatest increases in well length across most solutions.

Well results: WOPT increase vs. individual well length increase

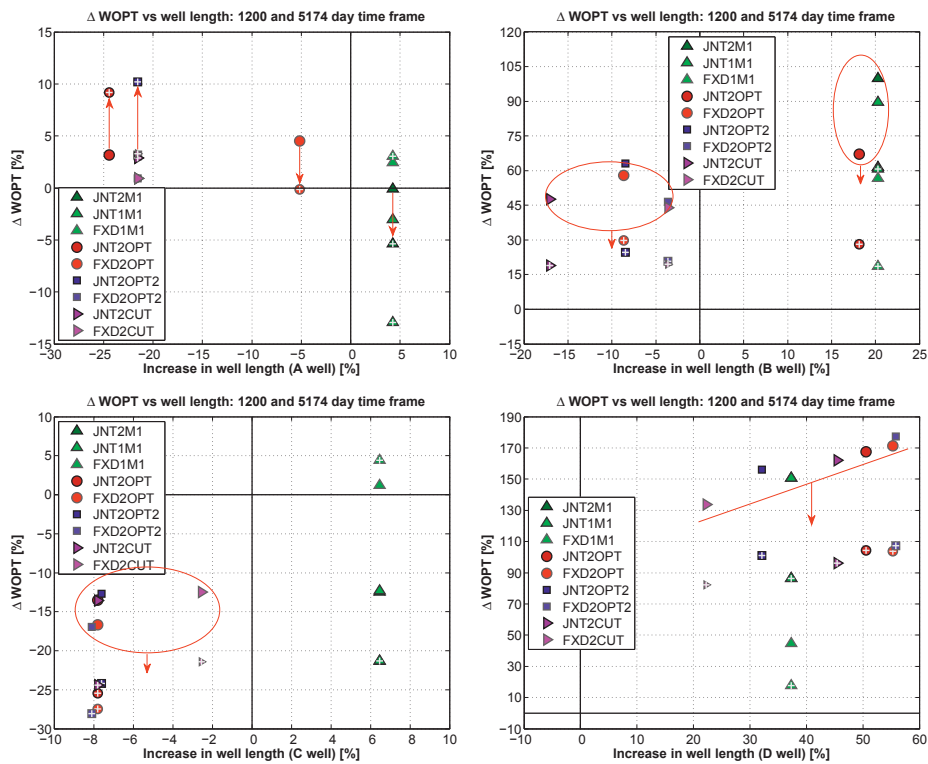


Figure 4.24: WOPT increases versus well length increases relative to base case values and length, respectively, for wells A, B, C, and D.

We now turn our attention to well-wise analysis. Figure 4.24 shows the change in well oil production total (WOPT) versus the increase (or decrease) in well length. This relationship is plotted for each solution well relative to the corresponding base case well WOPT values and lengths. As in previous figures, regular markers and markers with crosses correspond to 1200 day and 5174 day data, respectively. The vertical black line represents the well-length break-even point, while the horizontal black line represents the break-

even point for WOPT (since individual wells may have lower WOPT than their base case counterparts).

A well. Based on the A well graph in the upper left corner of Figure 4.24, we can distinguish between solution A wells that are significantly shorter (these are JNT2OPT, JNT2OPT2, FXD2OPT2, JNT2CUT and FXD2CUT), and wells that are close to or longer (namely FXD2OPT, JNT2M1, JNT1M1 and FXD1M1), than the base case A well. Interestingly, the shorter solution A wells increase their relative WOPT increases when operating in the 5174 day production horizon rather than in the 1200 day production time frame. The longer wells, on the other hand, have lower relative WOPT increases in the longer production time frame than in the shorter time horizon. These results indicate that shorter A well lengths are to be preferred if production is planned for a time frame comparable to the 5174 day production horizon, while longer well drains are recommended if the alternative is to achieve the greatest recovery within a short-term perspective.

B well. The graph in the upper right corner in Figure 4.24 shows WOPT versus well-length changes for the solution B wells. The two red ellipses with arrows show the general trend of WOPT decrease when we extend the production time frame from 1200 to 5174 days. At 5174 days, we have on the left side of the graph a group of B wells with WOPT increases ranging from 20% to 30%. These wells also have from about 5% to almost 20% shorter well bore lengths than the base case B well. The second group of B wells on the right hand side achieves similar or larger WOPT increases but with well bores that are almost 20% larger than the base case B well. One would expect that a significant part of the WOPT increases for the short wells in the first group (i.e., JNT2CUT, FXD2CUT, JNT2OPT2, FXD2OPT, and FXD2OPT2) comes from these well bores being placed at or close to the oil accumulation located in the eastern lobe of the reservoir. However, only two (JNT2OPT2 and FXD2OPT2) out of the five wells in this group actually reach into this area (see corresponding well configurations in figures 4.10 and 4.12, on pages 130 and 133, respectively). This is an important distinction because, even though production from the eastern lobe area using the B well is an important feature in several of the best-performing solutions (e.g., JNT2M1, JNT2OPT, JNT2OPT2; see sections 4.1.2 and 4.1.2), we see that this feature is not crucial either for an entire well configuration to do well (say, get an FOPT increase in the neighborhood of 10%, as seen e.g., in Table 4.6), or, as seen here, for the B well to achieve a significant increase in WOPT using a shorter well bore length.

C well. The bottom left graph in Figure 4.24 shows WOPT against well-length changes for the solution C wells. The red ellipse and arrow show the general trend of decreasing WOPT for increasing production time. As seen in the graph, this trend applies to all solution C wells (notice that to keep a reasonable shape, the ellipse does not extend over to the rightmost markers). From the graph we see that the solution C wells do not vary much in length, only between -8% and 6% , compared to the base case C well. Practically all solution C wells experience significant drops in WOPT of about -15% and of roughly

–25% for the 1200 day and the 5174 day production time frames, respectively. In Section 4.1.1 we saw that most of the solution C wells have well bore trajectories and lengths similar to that of the base case C well, and that the corresponding solution D wells are placed closer and often parallel to these trajectories. It is therefore likely that a substantial part of the WOPT drops seen for the solution C wells are caused by their respective solution D wells producing from close-by within the same reservoir area. We see some support for this interpretation based on the results from the FXD1M1 D well. This well is the only solution well placed west of the original base case D well location (the other solution D wells are placed east, or to the right, of this well), and can thus be said to have the least influence on its corresponding C well. In our graph we observe that this is the only well that obtains a positive WOPT difference compared to base case, which supports the reasoning of solution D wells overtaking solution C well drainage area.

D well. Finally, the bottom right graph in Figure 4.24 shows WOPT versus well-length changes for the solution D wells. As before, the straight red line and arrow indicate the general decrease in WOPT caused by increasing the production time frame from 1200 to 5174 days. All solutions wells (except one, discussed below) have large increases in WOPT in the range of 130% to 170% for the 1200 day production scenario. This range drops to between roughly 70% and 110% when using the 5174 production time frame. We see these results correlate with the previous description of the C well performances, in that the decreases observed for the C wells are here balanced by the increases in production from the D wells. Obviously, our general argument is that the solutions from the optimization procedure yield a combined improvement over the base case C and D wells, but, moreover our purpose here has been to discuss the difference in flow distributions between the solution wells compared to their base case counterparts. In the discussion for the C well, we contrasted these differences with the FXD1M1 solution that, as the base case, is single solution where the C and D wells far apart. Here we round off that discussion by pointing out that that single positive WOPT difference achieved by the FXD1M1 C well among all other solution C wells complements the relatively low increase in WOPT for the FXD1M1 D well. In summary, most solution C and D wells are placed closer to each other, compared to the base case configuration, and the solution D well trajectories are significantly longer than the base case D well. Combined, they yield higher recovery than their corresponding base case wells, but with different distributions of their individual productions.

Summary of FOPT increases versus changes in well length

In this section, we have plotted increases in field oil production against total drainage length differences to study the trade-off between longer well bore lengths and increased recovery. We have also shown how cumulative oil production for individual solution wells vary according to their relative well-length changes. We have provided these relationships so that they may be used to weight the recovery increases achieved through our optimization procedure against the uncertainties (outside of the optimization scope) associated with possibly drilling longer well drains. This allows the field operator to ultimately select a well configuration that balances recovery and drilling-risk according to the company's

risk attitude. We summarize the results from this section below.

- At total drainage lengths (for all wells) more than 4% shorter than base case, the JNT2CUT and the FXD2CUT solutions yield roughly a 10% increase in FOPT after 5174 days.
- At about the same total drainage length as the base case, the FXD2OPT2 solution yields FOPT increases of more than 20% and 10% for the 1200 day and the 5174 day production time frames, respectively.
- With only a 4% increase in total well drain length, the JNT2OPT2, JNT2OPT and FXD2OPT solutions yield FOPT increases of about 25% and 12% for the 1200 day and 5174 day production horizons, respectively.
- Well-wise results indicate that A wells with shorter well-lengths achieve greater relative oil recovery when applied within the longer 5174 day production horizon. On the other hand, longer well drains yield greater recovery if production is set for the shorter 1200 day production time frame.
- For the 5174 day time frame, solution B wells with about 5% to almost 20% shorter well bore lengths than the base case B well yield WOPT increases ranging from 20% to 30%.
- Overall, solution D well trajectories are significantly longer than the base case D well. Moreover, almost all solution C and D wells are placed closer to each other, and yield higher recovery than their corresponding base case wells, but with different distributions of their individual productions.

4.2 Field model tests on multiple realizations

In the previous section we tested all solution well configurations on a single base case model realization. In this section we expand this work and test each of the well placement solutions on a set of 10 realizations built around the base case model (these realizations were also provided by Total E&P Norge AS). The testing for the multiple realization case involves implementing each of the solutions on both the original base case model realization and on the 10 model realizations together, i.e., we test each solution on a total of $10+1$ model realizations. By a model realization we refer to an instance of the reservoir model that is associated with a unique set of reservoir data, e.g., porosity and permeability distribution. In addition to porosity and permeability, the multiple realization data set also accounts for changes in reservoir structure by having different distributions of active grid cells for the various realizations. Taken together, the set of model realizations serves as a rough representation of the geological uncertainty inherent in the description of the Martin Linge oil reservoir.

For each well configuration, we present the mean field oil production total ($\langle \text{FOPT} \rangle$) with standard deviation (σ) and compare it to the mean FOPT and standard deviation obtained when using the base case well configuration on the set of $10 + 1$ model realizations. We perform this comparison for both the 1200 day and the 5174 day production horizon.

Analog to testing well placement solutions for production time frames other than the one used during optimization, well placement solutions developed using only a single model realization are unlikely to yield optimal results when tested over a wide range of model realizations. Still, within a field development work process, information regarding solution robustness against reservoir uncertainty is important for the operator to decide if and how to further treat solutions obtained from the optimization procedure.

In the second part of this section we suggest a basic procedure the operator may apply to further treat solution data. The procedure consists of developing heuristic rules that work on the multiple realization solution data. For example, based on the obtained data (and possibly other criteria), rules could be developed that combine wells from the different solution well configurations, and also the base case, to produce hybrid well configurations. As an example in this work, we have applied a simple rule that replaces low-performing wells with high-performing wells from other solutions, and if necessary also reintroduces wells from the base case configuration. The main idea is to fuse current reservoir knowledge with the information obtained through the optimization procedure. We apply this rule on one out of the eight solution well configurations to obtain a small set of modified, or hybrid, well configurations. These hybrid well configurations have also been tested on the 11 model realizations. As for the original solutions, we also present the mean FOPT and standard deviations resulting from each of the hybrid solution tests, both for the 1200 day and for the 5174 day production horizons.

Hybrid cases are explored in Section 4.2.3. In the following we introduce the results from the multiple realization case using the original solutions. As in Section 4.1, we present results for the 1200 day and for the 5174 day production time frame separately. We introduce the results for the 1200 day production horizon first. Results for the 5174 day production time frame are presented in Section 4.2.2, on page 163. A summary of the

results presented in these first two sections is given in Section 4.2.3, on page 173.

4.2.1 Solution tests on multiple realizations: 1200 day production time frame

In this section we present results from each of the well placement solutions tested on the set of 11 realizations using the 1200 day production time frame. Results corresponding to the 5174 day production horizon are discussed in Section 4.2.2.

Table 4.7: Mean FOPT increase ($\Delta\langle\text{FOPT}\rangle$) and well mean WOPT increases ($\Delta\langle\text{WOPT}\rangle$) for all solutions tested over the multiple realization set. Each of the increases has a corresponding standard deviation from the mean (σ). Production data is obtained from Eclipse simulations running original base case production strategy (\mathbf{x}_c^S) over a 1200 day production horizon.

Solution	ECLIPSE ₁₂₀₀ (\mathbf{x}_p^* , \mathbf{x}_c^S)									
	$\Delta\langle\text{FOPT}\rangle$		$\Delta\langle\text{WOPT}\rangle$							
			WL.-A		WL.-B		WL.-C		WL.-D	
	[%]	$[\sigma]$	[%]	$[\sigma]$	[%]	$[\sigma]$	[%]	$[\sigma]$	[%]	$[\sigma]$
BASECASE	0.0	0.119	0.0	0.215	0.0	0.284	0.0	0.298	0.0	0.255
FXD1M1	1.9	0.183	-0.3	0.222	19.2	0.374	0.4	0.338	-3.3	0.164
FXD2OPT	1.9	0.226	-0.2	0.264	10.6	0.466	-11.5	0.233	16.9	0.343
FXD2OPT2	0.4	0.214	-2.2	0.247	5.2	0.437	-15.0	0.304	20.7	0.349
FXD2CUT	1.7	0.190	-4.3	0.242	-0.8	0.428	-2.9	0.250	16.1	0.223
JNT2M1	6.7	0.215	5.3	0.172	32.2	0.449	-3.2	0.233	6.8	0.313
JNT1M1	6.3	0.213	4.2	0.171	31.1	0.442	-3.0	0.231	6.6	0.315
JNT2OPT	2.9	0.208	-2.4	0.236	15.7	0.337	-9.2	0.239	17.7	0.333
JNT2OPT2	0.9	0.222	-1.6	0.246	5.0	0.443	-8.8	0.254	14.2	0.339
JNT2CUT	0.1	0.215	-2.3	0.237	-0.4	0.430	-7.7	0.235	13.2	0.354

Mean FOPT increase ($\Delta\langle\text{FOPT}\rangle$) and standard deviation data (σ) corresponding to each solution tested on the multiple realization set are given in Table 4.7. Well mean WOPT increases ($\Delta\langle\text{WOPT}\rangle$) with associated standard deviations are also given. As expected, we see that for each solution in this time frame, the mean FOPT increase over all realizations is significantly lower than the FOPT increase using only the single base case realization.

(It should be noted that, when using the base case well configuration, the FOPT curve corresponding to the original base case model realization clearly outperforms the FOPT curves corresponding to all the other realizations. I.e., when using base case wells on the original base case model realization we obtain the FOPT profile corresponding to the dashed line, upper-most gray curve in Figure 4.25. Also we see that for each solution well configuration tested over the set of multiple realizations, the best-performing FOPT curve is the one obtained when using the solution wells on the original base case model realization. This observation holds for all the multiple realization field production curves presented in this work, but not for the cumulative oil production profiles of the individual wells.)

Comparing the results in Table 4.7 with their counterpart single realization results shown in Table 4.5 (page 143), we see that the two best solutions from the single realization data, solutions JNT2M1 and JNT2OPT, decrease from 25.5% and 24.9% to 6.7% and 2.9%, respectively, when tested over multiple realizations (since JNT1M1 is a slight variation of JNT2M1, in this section the discussion of the JNT2M1 solution also applies to the JNT1M1 solution). The other solutions have steeper decreases, and are all below 2% for the multiple realizations case. Standard deviations corresponding to the mean FOPT of each solution are on average 1.8 times higher than the standard deviation for the base case. (Standard deviation tells us something about the spread of the data that make up the mean. In this context, it provides us with a measure of how robust a solution is with respect to the uncertainty represented by the multiple realization set.)

We notice the standard deviations corresponding to the solution A wells do not differ much from the standard deviation for the base case A well. Also, the standard deviations for all the solution C wells (except for the FXDM1 C well) are lower than the standard deviation for the base case C well. However, the standard deviations for the solution B and D wells show greater variation and are, on average, 1.5 and 1.2 times higher, than the standard deviations for their corresponding base case wells, respectively.

As discussed in the previous section, high increases in FOPT for several of the solutions are due to significant contributions from both the B and D wells. This is particularly the case for the JNT2M1 and the JNT2OPT solutions, which as mentioned, are the two solutions that have the highest FOPT increase in both the single and multiple realizations case. From the multiple realization data shown in Table 4.7, we see that the B well in these solutions has lost almost three-quarters of its gain compared to the single realization case (see Table 4.5). The decrease for the B wells in these solutions is due to the fact that in the base case map, these wells are draining an additional culmination to the south-east which may not be as high in other realization maps. Similarly, in the single realization case, the D well in these two solutions more than doubled its total production compared to base case, while in the multiple realization case, the increase in D well oil production for these two solutions is much less than 20%. For the JNT2M1 solution, standard deviations for these wells are close to average, i.e., B and D well standard deviations are 1.6 and 1.2 times higher, respectively, than their counterpart base case wells. For the JNT2OPT solution B well, standard deviation is lower than average at 1.2 times higher than base case B well, while D well standard deviation is 1.3 times higher than its base case analog.

Overall, both the JNT2M1 and JNT2OPT solutions show significant decreases when implemented over multiple realizations. In the multiple realization case, the mean well oil

production increases obtained from the B and D wells in the JNT2M1 and JNT2OPT solutions are much lower than the analog WOPT increases in the single realization case. However, in both the single and multiple realization case, these wells still drive much of the FOPT increases for these solutions.

In the following we plot the field and well cumulative oil profiles when using the JNT2M1 wells for each realization. For comparison, we also plot the production profiles obtained when using the base case wells over the same realizations. We choose the JNT2M1 solution because it is the solution with highest mean oil production total in Table 4.7.

Field and well oil production profiles for JNT2M1 solution

In figures 4.25 and 4.26 we plot the field and well cumulative oil production profiles for the JNT2M1 solution over the 1200 day production time frame. Production profiles corresponding to the solution are plotted as dark red curves, while base case production plots are drawn in gray. Curves representing average FOPT and WOPT production profiles for the solution and base case are plotted as thick red and gray lines, respectively.

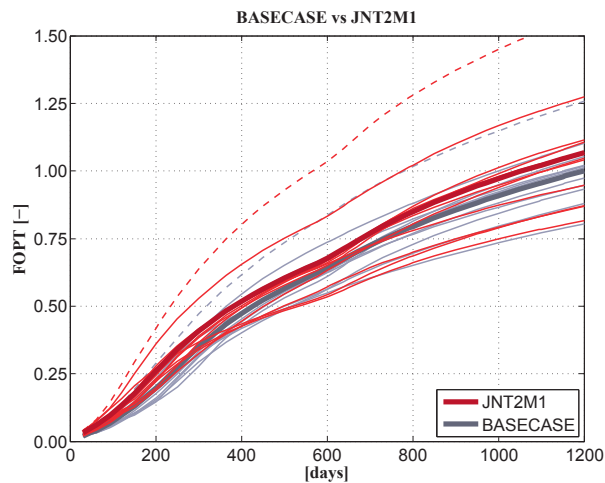


Figure 4.25: Field oil production totals for multiple realization case using the JNT2M1 solution (red) and base case well configuration (gray) over a production time frame of 1200 days. Dashed lines correspond to the original base case realization. Thin lines correspond to FOPT curves for individual realizations, while thick lines represent the average FOPT over all realizations.

Field oil production total plots. Figure 4.25 shows the field oil production total for each model realization running with the JNT2M1 solution. It also shows the field oil production totals for the same realizations when using the base case well configuration. In this figure we see that the mean FOPT curve for the JNT2M1 solution stays above the base case average for the entire 1200 day production time frame. As seen in Table 4.7,

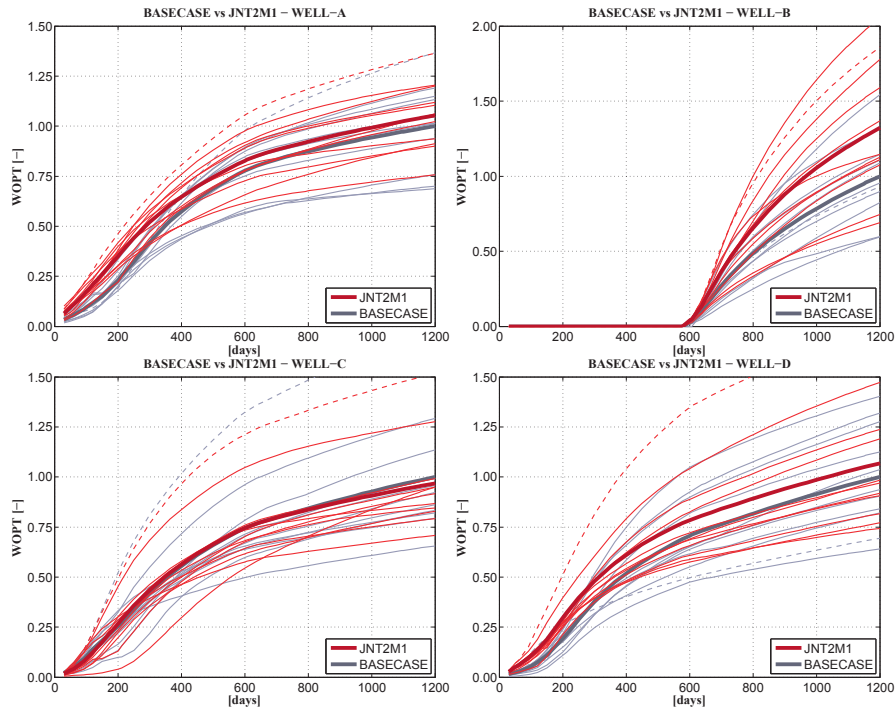


Figure 4.26: Well oil production totals for multiple realization case using the JNT2M1 solution (red) and base case wells (gray) over a production time frame of 1200 days. Dashed lines correspond to the original base case realization. For each well graph, thin lines correspond to WOPT curves for individual realizations, while thick lines represent the average WOPT over all realizations. Upper left and right graphs correspond to wells A and B, while production profiles for wells C and D are shown in the lower left and right graphs, respectively.

after 1200 days of production, the mean FOPT increase from the JNT2M1 solution is 6.7% larger than mean FOPT using the base case wells. For about the first 300 days of production, all FOPT curves for the JNT2M1 solution outperform their corresponding base case FOPT curves. From this point on, at least three solution FOPT curves veer off the general trend of the average FOPT curve. Covering the complete 1200 day period, two FOPT curves from the JNT2M1 solution (including the one corresponding to the base case model realization) greatly outperform their average.

Well oil production total plots. Figure 4.26 presents four separate graphs showing the cumulative oil profiles for each of the JNT2M1 wells over all model realizations. The graphs also show the oil production totals for the corresponding base case wells. In addition, for each of the four graphs, we can compare the mean production from each of the JNT2M1 solution wells to the mean performance of their corresponding base case wells.

Here we focus on comparing the performances for the multiple realization case against the well increases achieved in the single realization case. (Recall that single realization

WOPT plots for all solutions are shown in Figure 4.24, with detailed values given in Table 4.5.) In the single realization case, we have seen that the A well for the JNT2M1 solution shows no WOPT increase while its C well has a 13% decrease in WOPT compared to their respective base case wells. However, in the multiple realization case, we see from the upper left graph in Figure 4.26, that the A well for the JNT2M1 solution has a slight increase in mean WOPT compared to the base case mean, and from the lower left graph we notice that its C well mean WOPT is now only a few percent below its corresponding base case mean. Thus, relatively, the contributions from the A and C wells for this time frame increase in the multiple realization case. At the same time though, as we have noted before, the significant gains in mean total oil production achieved by the B and D wells in the single realization case are comparably much lower in the multiple realization case. Comparing the mean increases in the upper and lower right graphs in Figure 4.26 with the WOPT increases shown in Figure 4.24, we see how the gains achieved for these wells are significantly reduced once they are tested over the multiple realization set. In the following section we see these gains are harder to retain when solutions are tested over the larger field production horizon.

4.2.2 Solution tests on multiple realizations: 5174 day production time frame

In this section we present results from each of the well placement solutions tested on the set of 11 realizations using the 5174 day production time frame. The description of the 5174 day data, including how it is presented and treated, e.g., the layout of tables and graphs, is analogous to the description of the 1200 day data given in the previous section. This means that results in this section are organized into tables and figures in the same manner as in the preceding section, except that the tables and graphs in this section now present the production data obtained using the 5174 day , and not the 1200 day , production time frame. General descriptions of tables and graphs are therefore not repeated. Mainly, in this section, we will discuss results from the multiple realization case, and in particular, we will focus on comparing 5174 day results with results obtained using the 1200 day production time frame. We will begin the discussion with a description of main results given in Table 4.8, and finish with selecting the best solution from the result table and plotting its field and well cumulative oil profiles for each realization, and their corresponding means, against base case production profiles.

Table 4.8 shows a further overall reduction in mean FOPT results ($\Delta(\text{FOPT})$) when running all solutions over the set of multiple realizations using the 5174 day production horizon. At this production time frame, more than half of the solutions yield mean FOPT results lower than those achieved by the base case well configuration. Still, the JNT2M1 solution manages to retain a 2.9% increase in mean FOPT, and several wells show robust production increases in their individual performances (discussed further below).

On average, solution FOPT means have 1.5 times higher standard deviations than the standard deviation for the base case FOPT mean. While A wells for all solutions have, on average, about the same standard deviation as the base case A well, solution B, C and D wells have standard deviations that are, on average, roughly 1.2, .9 and 1.3 times the magnitude of their corresponding base case WOPT mean standard deviations.

Table 4.8: Mean FOPT increase ($\Delta\langle\text{FOPT}\rangle$) and well mean WOPT increases ($\Delta\langle\text{WOPT}\rangle$) for all solutions tested over the multiple realization set. Each of the increases has a corresponding standard deviation from the mean (σ). Production data is obtained from Eclipse simulations running original base case production strategy (\mathbf{x}_c^S) over a 5174 day production horizon.

Solution	ECLIPSE ₅₁₇₄ (\mathbf{x}_p^* , \mathbf{x}_c^S)									
	$\Delta\langle\text{FOPT}\rangle$				$\Delta\langle\text{WOPT}\rangle$					
			WL.-A		WL.-B		WL.-C		WL.-D	
	[%]	[σ]	[%]	[σ]	[%]	[σ]	[%]	[σ]	[%]	[σ]
BASECASE	0.0	0.110	0.0	0.235	0.0	0.313	0.0	0.258	0.0	0.213
FXD1M1	-1.8	0.160	-0.6	0.266	7.4	0.404	-0.9	0.343	-11.6	0.168
FXD2OPT	-1.0	0.168	13.8	0.258	-3.0	0.428	-20.0	0.203	10.1	0.294
FXD2OPT2	-1.5	0.162	13.7	0.245	-1.7	0.386	-25.8	0.242	14.3	0.325
FXD2CUT	-0.3	0.153	11.2	0.250	-8.9	0.376	-10.6	0.190	8.7	0.221
JNT2M1	2.9	0.163	9.8	0.191	16.5	0.408	-12.0	0.202	3.8	0.251
JNT1M1	2.5	0.158	7.1	0.178	17.2	0.398	-11.5	0.199	3.7	0.253
JNT2OPT	0.5	0.161	15.7	0.245	-0.8	0.323	-18.4	0.206	10.6	0.317
JNT2OPT2	0.2	0.171	20.8	0.273	-9.2	0.420	-15.7	0.215	8.1	0.298
JNT2CUT	-2.0	0.164	13.6	0.239	-8.4	0.375	-16.4	0.214	6.3	0.329

In the following we compare results between the 5174 day and 1200 day production horizons, both within the multiple realization case. In Table 4.8 we see that for all solutions, the A well yields larger relative WOPT mean increases for the 5174 day horizon than when using the 1200 day production time frame. Conversely, for most of the solutions tested over the larger production horizon, the contribution from the B wells is now below the base case WOPT mean. For the B well, most gains in the shorter horizon are positive, while for the longer horizon they are mostly negative, except for two solutions (solutions FXD1M1 and JNT2M1). We also notice that the solution with the highest FOPT increase, the JNT2M1 solution, is the only solution where both the B and D wells have increases over base case. As in the results for the shorter production time frame, the C well in all solutions performs poorly when tested over multiple realizations.

Next we plot the field and well cumulative oil profiles for the solution with the highest FOPT increase, i.e., the JNT2M1 solution.

Field and well oil production profiles for JNT2M1 solution

In figures 4.27 and 4.28 we plot the field and well cumulative oil production profiles for the JNT2M1 solution over the 5174 day production time frame. As before, dark red and gray curves represent solution and base case profiles, respectively. Dashed lines represent those FOPT and WOPT profiles that correspond to the original base case realization. Average FOPT and WOPT production profiles are plotted as thick lines.




Figure 4.27: Field oil production totals for multiple realization case using the JNT2M1 solution (red) and base case well configuration (gray) over a production time frame of 5174 days. Dashed lines correspond to the original base case realization. Thin lines correspond to FOPT curves for individual realizations, while thick lines represent the average FOPT over all realizations.

Field oil production total plots. Figure 4.27 shows the field oil production total for each model realization running with the JNT2M1 and base case wells over a production time frame of 5174 days. We know from Table 4.8 that, after 5174 days of production, the JNT2M1 average FOPT curve yields a 2.9% increase over average base case FOPT. In Figure 4.27 we see this curve has a somewhat greater increase compared to the base case mean FOPT curve in the time period between 1000 and 2500 days. From the figure it appears that the drop from this high-value period is related to production profiles from at least two different realizations. At about 1500 days these total oil production profiles veer off the general trend of the mean (see e.g., second red curve from the top), and significantly underperform with respect to the rest of the FOPT curves for the solution. For the longer field production horizon, low-performances on some realizations have a large effect on ultimate mean recovery. In the following we explore the performance of each of the wells over the various realizations.

Well oil production total plots. Figure 4.28 presents the four graphs that show the cumulative oil profiles for each of the JNT2M1 wells over all model realizations. These

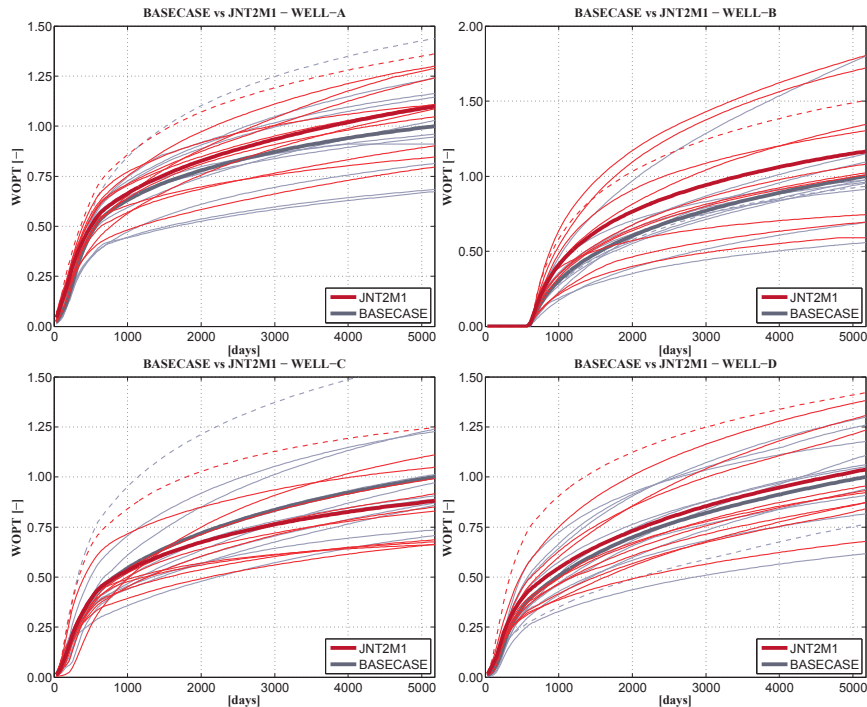


Figure 4.28: Well oil production totals for multiple realization case using the JNT2M1 solution (red) and base case wells (gray) over a production time frame of 5174 days. Dashed lines correspond to the original base case realization. For each well graph, thin lines correspond to WOPT curves for individual realizations, while thick lines represent the average WOPT over all realizations. Upper left and right graphs correspond to wells A and B, while production profiles for wells C and D are shown in the lower left and right graphs, respectively.

graphs finally demonstrate how solutions developed without taking into account the entire realization set can form low-performing WOPT curves that, from early on, and particularly when implemented over the larger production time frame, effectively reduce optimization gains, as seen both in Table 4.7, and Table 4.8. For the A, B and D wells, several high-performing WOPT curves, among them the ones corresponding to the original base case realization (dashed line), balance several low-performing curves to yield higher WOPT means than their respective base case wells. Still, among these wells, only the A well manages to increase its mean WOPT over the greater time horizon. The B well manages to keep its mean WOPT increase somewhat constant over the extended horizon, even though several of its individual realization curves start performing poorly after about 1000 and 1500 days of production.

The C well production curves shown in the lower left graph in Figure 4.28 confirm the negative contributions from this well seen in tables 4.7 and 4.8. In this graph we see the solution and base case mean curves separate already after about 1000 days of production. After 1200 days we see from Table 4.7 that the mean curve for the solution is 3.2% below the mean base case curve. In particular for this well, there are numerous curves that

have a close to flat rate of increase after about 800 days of production. Subsequently, after 5174 days of production, the mean FOPT for this solution is 12% less than the mean base case FOPT.

4.2.3 Hybrid solution tests on multiple realizations

In the beginning of this section we suggested a heuristics-based procedure could be devised to modify final well configurations and further improve results. The procedure would apply engineering heuristics (Koen, 2003) to create modified, or hybrid, solutions. The heuristics⁸ would be designed to adapt individual solutions according to given field development criteria, either independently, or by combining the various well trajectories from the solution set into new solutions. Establishing such a procedure would be the subject of future work. In this section we discuss some of the ideas behind the procedure, and provide the results from a test using a simple rule.

The purpose of the procedure would be to compensate for the different approximations made during optimization, as well as to further develop the solutions by specifically taking into account important considerations from field development work (ranging from pilot well drilling to geological uncertainty). As we have discussed in previous sections, to be able to efficiently implement our optimization methodology, we have both made several approximations to the field development setup, and placed some of its more general considerations outside of the scope of the optimization effort. For example, for implementation reasons, we have used a time horizon for reservoir model simulation that is substantially shorter than the one planned for the development of the reservoir, and we have not taken multiple realizations into consideration during our optimization. Throughout our testing, we have seen the effects of these design choices. In particular, we have seen how the gains from the optimization effort decline once simulations are run for the full time horizon, and also when, as we have shown in this section, the optimized well configurations are brought into the larger context of field development work, where it is crucial to consider geological uncertainty and test the solutions over multiple realizations.

The idea behind the proposed procedure is to find those areas of reservoir knowledge and field design targets and bounds that are directly related to the well placement problem, and to synthesize this information into a set of rules, or heuristics. Subsequently, these heuristics would use the solution data set from the optimization effort to develop new solutions. These new solutions would then have design characteristics that satisfied the field development concerns embedded in the heuristics. These set of heuristics would be able to address both specific or more general types of field development concerns. More importantly, the building blocks for the new solutions would be those final well configurations that, although developed using a narrower scope than the one drawn for the heuristic procedure, are still the product of an optimization process, i.e., they are the results of a systematic search.

As mentioned above, establishing such a procedure would be the subject of further work. Critically, this work would require not only a research-type of study, but a closer collaboration with the engineering team in charge of field development. Based on the accumulated knowledge and expertise of the field development team, a collaborative effort would involve discussions, as well as testing and confirmation of proper language and in-

tention behind the different heuristics. This process would also ensure that the rules were accurate representations of main field development concerns. At the same time, it is likely that this process will produce concrete specifications about the problem itself. Clearly, these specifications would be coded directly as constraints within the optimization framework. (In a sense, heuristics may be thought of as loose type of constraints that work on solution data. As such, they reshape and/or create new solutions a posteriori so that these solutions may now satisfy considerations that were left out of optimization scope, or were not possible to implement as constraints during the optimization in the first place.)

In this work we apply a simple rule on the solution data as an example of this type of procedure. The rule is stated as follows:

- take the best-performing solution; and,
 - replace low-performing wells with high-performing wells from other solutions; also, if necessary, use wells from base case configuration

The main concern behind this rule is to minimize the loss in optimization gains when solutions are run for the field production time frame and over the set of multiple model realizations. Obviously, much more advanced heuristics using several conditional statements can be established and applied to the solution set.

We apply the rule stated above on the JNT1M1 solution⁹ and obtain the three hybrid solutions shown in Table 4.9. The first hybrid solution is built as follows: J1M1H1 is equal to JNT1M1 except its D well is replaced by the base case D well. The two other solutions are built in similar manner, as demonstrated in Table 4.9.

Table 4.9: The table shows the well composition of the three hybrid solutions developed from applying a simple heuristic procedure on the JNT1M1 solution.

Hybrid solution	Composition
J1M1H1	→ JNT1M1 [A, B, C] + BASECASE [D]
J1M1H2	→ JNT1M1 [A, B] + BASECASE [C, D]
J1M1H3	→ JNT1M1 [A, B, C] + FXD2CUT [D]

The resulting well configurations for the hybrid solutions are shown in Figure 4.29, together with base case wells. For those solutions where base case wells have been introduced, there is an overlap between base case and solution wells in the figure. The saturation (HuPhiSo) maps are given after 5174 days of production. As for the original solutions, each of the hybrid solutions is tested over the multiple realization set for both sets of production time frames. We start our discussion with the 1200 day production horizon, and continue with the 5174 day production time frame starting from page 172.

Hybrid solutions: 1200 days production time frame

Table 4.10 shows the differences in final field and well cumulative oil production for each of the hybrid solutions compared to base case values after a 1200 day production time

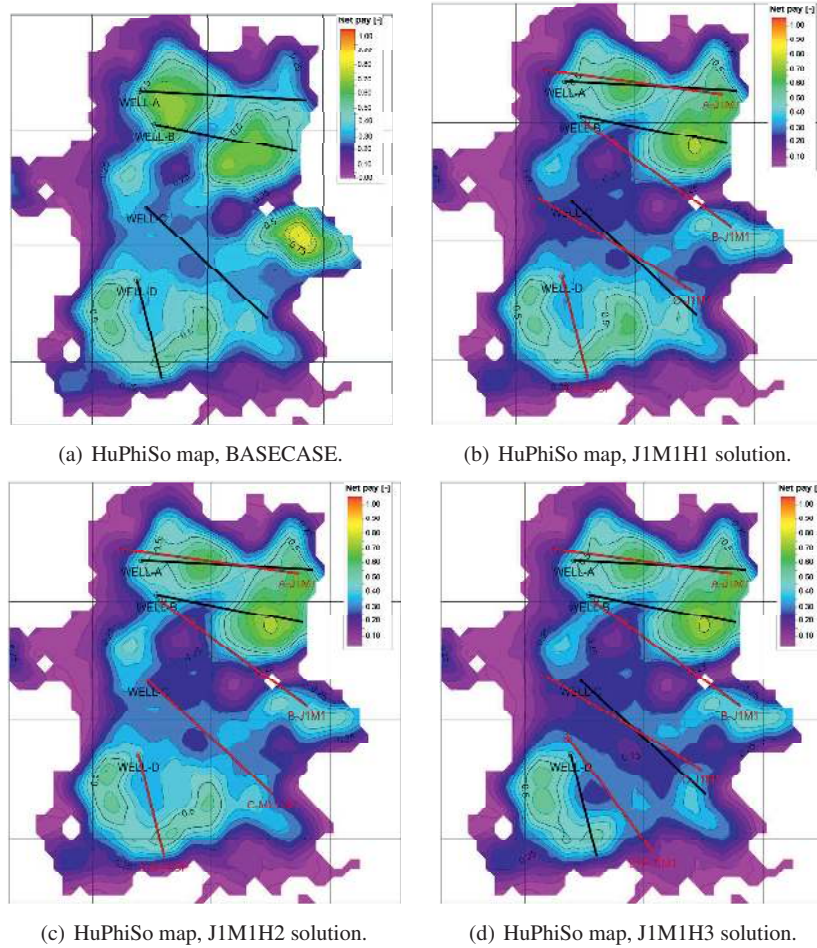


Figure 4.29: HuPhiSo oil saturation maps after 5174 days of production for the three hybrid solutions: J1M1H1, J1M1H2 and J1M1H3. Upper left map shows corresponding base case map. Scales have been normalized for confidentiality reasons.

frame. For additional reference, we have furthermore, from Table 4.8, reintroduced the increases corresponding to the parent JNT1M1 solution for this time frame. From this table we see that the J1M1H3 solution yields the highest increase in mean field oil production total ($\Delta\langle\text{FOPT}\rangle$) at 8.3% over base case. This represents a 2% increase over its regular counterpart, the JNT1M1 solution in Table 4.7.

The FOPT profiles corresponding to the J1M1H3 solution are shown in Figure 4.30. Similarly, the cumulative oil profiles for each of the J1M1H3 wells over all model realizations are presented in Figure 4.31. As before, solution profiles are plotted as dark red curves, while base case production plots are drawn in gray.

Table 4.10: Mean FOPT increase ($\Delta\langle\text{FOPT}\rangle$) and well mean WOPT increases ($\Delta\langle\text{WOPT}\rangle$) for all hybrid solution tested over the multiple realization set. Each of the increases has a corresponding standard deviation from the mean (σ). Production data is obtained from Eclipse simulations running original base case production strategy (\mathbf{x}_c^S) over a 1200 day production horizon. Values from the base case configuration as well as the parent JNT1M1 solution are set as reference.

Hybrid	ECLIPSE ₁₂₀₀ ($\mathbf{x}_{p,H}^*, \mathbf{x}_c^S$)									
	$\Delta\langle\text{FOPT}\rangle$		$\Delta\langle\text{WOPT}\rangle$							
			WL.-A		WL.-B		WL.-C		WL.-D	
	[%]	[σ]	[%]	[σ]	[%]	[σ]	[%]	[σ]	[%]	[σ]
BASECASE	0.0	0.119	0.0	0.215	0.0	0.284	0.0	0.298	0.0	0.255
JNT1M1	6.3	0.213	4.2	0.171	31.1	0.442	-3.0	0.231	6.6	0.315
J1M1H1	6.0	0.171	3.2	0.183	30.0	0.427	3.5	0.298	-1.1	0.264
J1M1H2	5.0	0.154	2.3	0.179	29.7	0.446	0.3	0.300	0.2	0.272
J1M1H3	8.3	0.198	2.0	0.187	31.6	0.444	-2.3	0.257	16.2	0.230




Figure 4.30: Field oil production totals for multiple realization case using the J1M1H3 solution (red) and base case well configuration (gray) over a production time frame of 1200 days. Dashed lines correspond to the original base case realization. Thin lines correspond to FOPT curves for individual realizations, while thick lines represent the average FOPT over all realizations.

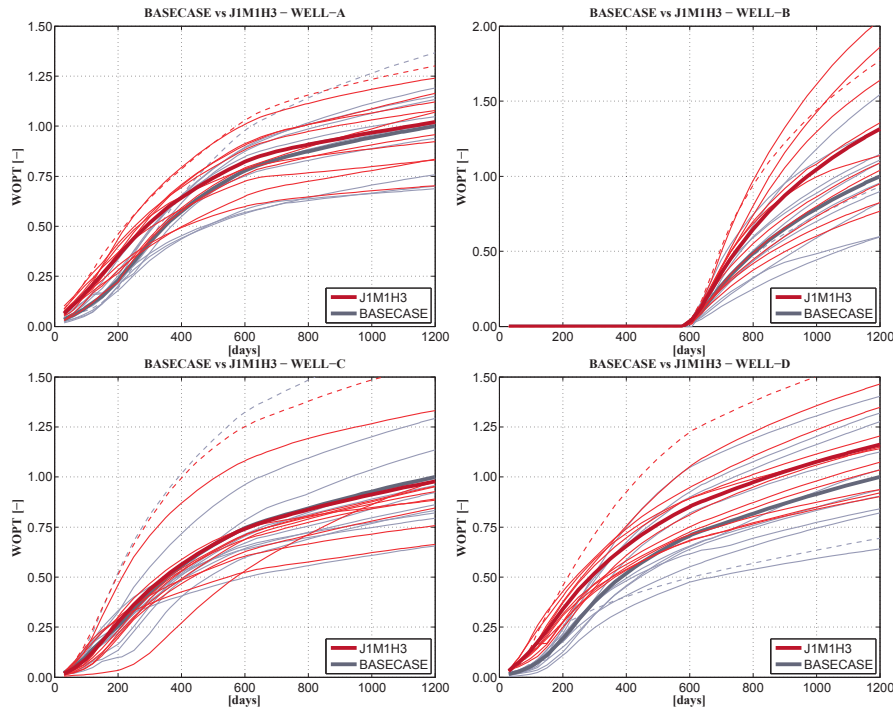


Figure 4.31: Well oil production totals for multiple realization case using the J1M1H3 solution (red) and base case wells (gray) over a production time frame of 1200 days. Dashed lines correspond to the original base case realization. For each well graph, thin lines correspond to WOPT curves for individual realizations, while thick lines represent the average WOPT over all realizations. Upper left and right graphs correspond to wells A and B, while production profiles for wells C and D are shown in the lower left and right graphs, respectively.

Field oil production total plots. In Figure 4.30 we see the hybrid solution producing a mean FOPT curve with significantly better performance than the base case mean FOPT over the entire 1200 day production time frame. Also, visually, the mean FOPT curve corresponding to the hybrid solution yields a clear improvement compared to field oil production curve shown for the JNT2M1 solution shown in Figure 4.25 (page 161).

Well oil production total plots. Figure 4.31 shows the oil production total for each well in the JN1M1H3 solution over all model realizations. Recall that the J1M1H3 solution consist of a D well that has been introduced from the FXD2CUT solution. Comparing these curves with the corresponding plots for the JNT2M1 solution shown in Figure 4.26, we notice a moderate decrease for the A and a slight improvement for the C well (these two wells are essentially the same for the JNT1M1 and JNT2M1 solutions). From tables 4.7 and 4.7 we have that the WOPT gain drops from 4.2% to 2.0% for the A well, while the improvement for the C well is below 1%.

More importantly, we observe that the WOPT increase for the new D well, introduced

from the FXD2CUT solution, is substantially larger than the increase when using the original D well (see Table 4.7). In fact, the increase from the FXD2CUT D well, which produced a 16.1% WOPT increase in the FXD2CUT solution (see Table 4.7), produces practically the same increase in the hybrid solution, namely 16.2% (see Table 4.10). Together with an improved mean WOPT performance for the J1M1H3 C well (not replaced in this solution) at -2.3% (from Table 4.10) instead of at -3.0 (see Table 4.7), we observe that this pairing of C and D wells yield an improved solution for the multiple realization case.

Hybrid solutions: 5174 days production time frame

Table 4.11: Mean FOPT increase ($\Delta\langle\text{FOPT}\rangle$) and well mean WOPT increases ($\Delta\langle\text{WOPT}\rangle$) for all hybrid solutions tested over the multiple realization set. Each of the increases has a corresponding standard deviation from the mean (σ). Production data is obtained from Eclipse simulations running original base case production strategy (\mathbf{x}_c^S) over a 5174 day production horizon. Values from the base case configuration as well as the parent JNT1M1 solution are set as reference.

Hybrid	ECLIPSE ₅₁₇₄ ($\mathbf{x}_{p,H}^*, \mathbf{x}_c^S$)									
	$\Delta\langle\text{FOPT}\rangle$				$\Delta\langle\text{WOPT}\rangle$					
			WL.-A		WL.-B		WL.-C		WL.-D	
	[%]	$[\sigma]$	[%]	$[\sigma]$	[%]	$[\sigma]$	[%]	$[\sigma]$	[%]	$[\sigma]$
BASECASE	0.0	0.110	0.0	0.235	0.0	0.313	0.0	0.258	0.0	0.213
JNT1M1	2.5	0.158	7.1	0.178	17.2	0.398	-11.5	0.199	3.7	0.253
J1M1H1	3.9	0.131	5.3	0.202	15.9	0.389	-1.0	0.273	-1.1	0.211
J1M1H2	3.3	0.122	5.4	0.198	16.0	0.392	-3.8	0.262	-0.2	0.218
J1M1H3	3.6	0.152	4.9	0.213	17.3	0.402	-9.4	0.205	7.3	0.198

Table 4.11 presents the FOPT and WOPT increases for the hybrid solutions against base case for the 5174 day production horizon. Also, for additional reference, the increases corresponding to the parent JNT1M1 solution are reintroduced from Table 4.7. We notice from Table 4.11 that the best-performing hybrid solution for the 5174 day production time frame is now the J1M1H1 solution. This solution has an increase in mean FOPT of 3.9% over mean base case FOPT. However, at this time frame, the results from the two other hybrid solutions also show similar increases of 3.3% and 3.9% for the J1M1H2 and J1M1H3 solutions, respectively. We also note that, for this time frame, the field and well-wise standard deviations for all the hybrid solutions are not much greater, and often close to, the respective base case standard deviations.




Figure 4.32: Field oil production totals for multiple realization case using the J1M1H1 solution (red) and base case well configuration (gray) over a production time frame of 5174 days. Dashed lines correspond to the original base case realization. Thin lines correspond to FOPT curves for individual realizations, while thick lines represent the average FOPT over all realizations.

Field oil production total plots. Figures 4.32 and 4.33 show the corresponding FOPT and WOPT profiles for the J1M1H1 solution when running over a 5174 day production time frame. In Figure 4.32 we see the hybrid J1M1H1 solution produces a somewhat improved mean FOPT curve. In particular, we notice a single-realization FOPT curve (second red line from top) that has a significant decline starting at about 700 days and continuing until the end of the production horizon.

Well oil production total plots. Figure 4.33 presents the four graphs for the cumulative oil profiles for each J1M1H1 well. These graphs can be compared to those in Figure 4.28, which present the same type of profiles, but for the JNT2M1 solution. The greatest differences between these two figures are in the graphs for the C and D wells. Crucially, both the solution C and D wells now have mean WOPTs that closely match those of their counterpart base case wells. In contrast, for the JNT2M1 solution, the D well has a slight increase over base case mean, while the C well has a substantially lower mean WOPT than the base case mean, presumably because of the relative closeness between the C and D wells in that solution. Thus, replacing the solution D well with the FXD2CUT D well has yielded a well configuration where the actual improvement over base case comes from the A and B wells.

Summary of solution tests on multiple realizations

In this section we have tested all our solutions on a set of 11 realizations both for the 1200 day production time frame used in the optimization procedure, and for the 5174 day

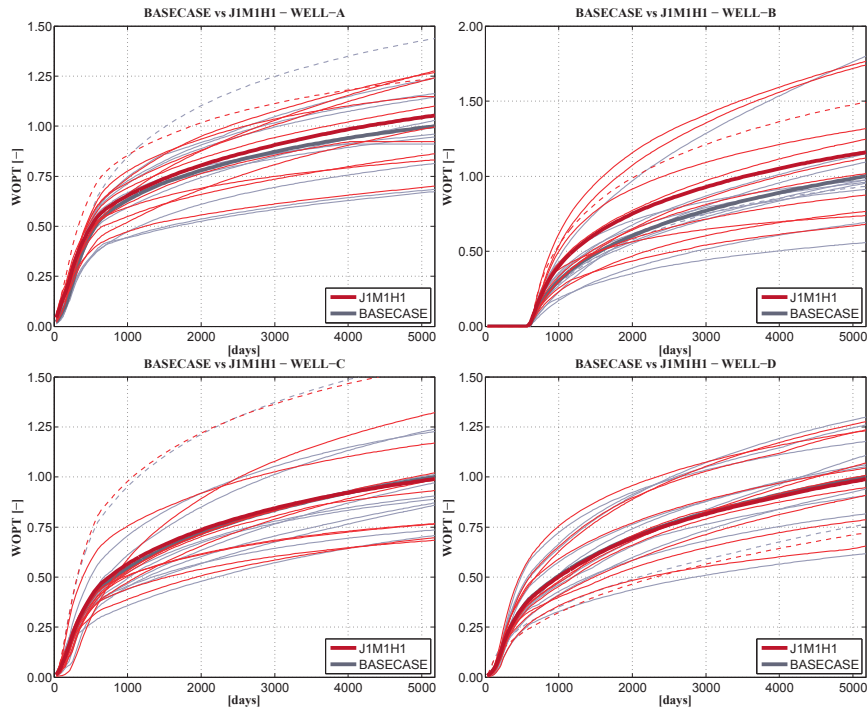


Figure 4.33: Well oil production totals for multiple realization case using the J1M1H1 solution (red) and base case wells (gray) over a production time frame of 5174 days. Dashed lines correspond to the original base case realization. For each well graph, thin lines correspond to WOPT curves for individual realizations, while thick lines represent the average WOPT over all realizations. Upper left and right graphs correspond to wells A and B, while production profiles for wells C and D are shown in the lower left and right graphs, respectively.

production horizon used in the original field development case. Taken together, these results represent the difficulty of applying solutions to a case with a greater production time frame than the time horizon used during optimization, and also of testing these solutions on realizations that were not included in the optimization procedure. We have also tested a simple heuristic procedure to treat solution data. This procedure was suggested as a step to integrate optimization results within the larger exploration and learning process that is the work of field development planning. Below, we summarize the main results and points of discussion presented in this section.

- As expected, solution results developed using only a single realization decrease significantly once the solutions are implemented over multiple realizations.
- Crucially, though single-realization solutions may perform poorly for some realizations, they may still yield substantial increases in mean field production, if applied within the optimization time frame.
- Within the 1200 day production time frame, some solutions perform poorly for most

realizations, but at least two solutions (JNT2M1 and JNT2OPT) retain substantial increases in mean field oil production total of almost 7% over base case.

- Test results confirm that, when applied for the larger production horizon, poor solution performances over several realizations effectively drain optimization gains developed for the shorter time frame.
- The bulk of solutions, obtained using various approximation and a limited optimization scope, may be treated by heuristic rules built based on other considerations not included in the optimization procedure. Hybrid solutions containing specifically selected parts from the various solutions may be developed to answer to a broader set of considerations from the field development work process.

4.3 Final topics on field case application

In this final section, we briefly discuss topics related to the solution testing effort conducted in this chapter, and the various limitations due to application design. We also discuss possible further developments of the Martin Linge field case, and how these developments may be treated by the optimization framework presented in this thesis.

We start this section by discussing the applicability of solutions to the field case, and continue with a brief description of a prospective production scenario for the Martin Linge oil reservoir involving inflow control valves. At the end of this section, we offer a summary and final comment of the pilot study conducted in this and the previous chapter.

4.3.1 Applicability of solutions and limits of application design

Developmental stage of application. One goal of the current work is that the results from the optimization procedure can possibly provide useful information to the field development work process of the operator. However, the application of the methodology is still at a very early stage of development, and it requires substantial maturation before its use can be properly considered within the work process of the operator. At the current stage of development, there are several issues that seriously limit the use of the procedure and its results within a field development and operations environment. Some of these issues are related to problem definition and application design in general, while others are mainly related to scope setting and model approximations that have been necessary to get the pilot application off the ground. Issues related to scope, and then approximations, are discussed next. At the end of the section we treat issues regarding problem definition and how to facilitate the translation of expert industry knowledge for use in research applications.

Testing issues. In the following, we will discuss how the core optimization results presented in Chapter 3 have been restricted by the scope of the optimization procedure, and how this has limited their applicability within the work process of the operator. We refer to the testing of obtained solutions for parameters that were not included in the original scope of the optimization procedure as testing issues. This type of testing is part of the application design, and provides us with important feedback on scope definition,

performance and applicability of solutions. Importantly, it demonstrates the limitations of the results obtained, and helps clarify which areas in application development that need further improvement. As we have seen, two of the main testing issues have been simulation over the field case production time frame of 5174 days, and the testing of the solutions over multiple model realizations. Extensive results from this testing have been provided in this chapter, (in particular Section 4.1: Test results from solution cases (page 107 and Section 4.2: Field model tests on multiple realizations) (page 158). Here we briefly discuss their implications.

Testing over original time horizon. Optimization has been performed using reservoir simulations running for only 1200 days instead of the larger field case production time frame of 5174 days. (This approximation has already been commented in Section 3.4.1, and further discussed in Section 4.1.) Crucially, test results from this chapter show that the gains obtained from the implementation of the solutions drop by about 50% once the solutions are run using the larger production time frame. That the solutions can only be considered optimal within a fraction of the production horizon intended for the development of the reservoir is a serious disadvantage for their further use. At best, the alternative well configurations can, at least for the first years of production, serve as estimates of recovery patterns not previously thought of. Production from individual well trajectories may also serve as standalone recovery studies from regions of the reservoir not previously examined. At worst, the reduced-time approximation invalidates many well trajectories that may have seemed promising for further study. Clearly, for future applications, improvements should be made to have the procedure consider the entire field case time horizon during optimization.

Testing over multiple realizations. A key factor not included in the scope of the procedure has been geological uncertainty. Still, each solution obtained by the procedure has been tested over a set of field case model realizations provided by Total E&P Norge AS. The solutions have been run over the set of realizations, and expected oil production totals and variation data have been computed, using both the 1200 day and the 5174 day production time frames. The results from these tests show severe drops in optimization gains from all solutions. For the best-performing solution, i.e., the solution with the highest field oil cumulative in the single-realization case, the optimization gain is reduced from almost 26% to about 7% for the 1200 day production time frame, and to less than 3% (from 13%) for the longer horizon. For the remaining solutions, the loss in gains is even more drastic. Some solutions yield mean recoveries that are indistinguishable from the mean recovery obtained when using the initial well configuration. Finally, from these data, it seems clear that none of the solutions can robustly handle the spread in structural uncertainty represented by the multiple realizations. In general, not taking into account the structural uncertainty of the reservoir within the optimization procedure, puts all the obtained solutions at a substantial disadvantage since obtaining a significant degree of certainty around any type of solution is critical from a field development point of view. It is therefore important that future applications aiming at contributing information to the field development work process, consider geological uncertainty during the optimization procedure, e.g., by optimizing using the expected value of the performance measure as cost function. Implementing an optimization effort using both a long-term production horizon, and computing for multiple model realizations at the same time, is a signifi-

cant undertaking. To balance computational demand, also decisions regarding application design have been made. Issues related to approximations within the procedure itself are discussed next.

Limitations due to application design. This work has shortened the simulation time frame and excluded model uncertainty to reduce computational demand, even though, as already mentioned, the applicability of our optimization solutions has been greatly reduced due to these decisions. Still, even after these limitations in scope, computational cost for this application work remained high, and further approximations, now with respect to design configuration, were seen as necessary to reduce overall cost. The main adjustment in this regard was replacing the field case model with a work model approximation. At this point, however, we attempted to counter the expected loss in applicability by establishing a relatively accurate approximation of the field case model (see Section 3.1: Targets and strategy for application development, page 42).

Future applications also need to explore alternative design configurations for the application process, to make it more effective, and thus to achieve a better balance between runtime and the development of useful information within an industry context. (The reader may compare this point to the description of our second strategy component, Section 3.1.1, page 45.) In particular, different options for simulation should be explored further, e.g., regarding the extended use of surrogate models, or whether, for some limited application scopes, we should consider using the field case model directly. Furthermore, if future work entails the continued use of a work model, further effort should be put on requiring fewer approximations for the model transfer. Finally, future applications need to ensure a better alignment of solutions with the business needs of the Industry Partner (thus increasing applicability, as the term has been discussed in this section), possibly by building a more refined understanding of the problem facing the operator (more on this topic below).

Limitations related to application design have been discussed here lastly because, based on our experience, once the technology components are sufficiently developed (e.g., the algorithms that make up the parts of the optimization framework), the challenge then of how to best define the application scope is best handled at the level of collaboration work. In our case, we believe further improvement in the use of collaboration processes is important to, one, achieve a better overview and understanding of the tasks performed within field development operations, and two, to obtain a better translation of expert knowledge and problem considerations from the industry operations side to the realm of applied research. Clearly, the topics discussed here are closely related to our emphasis on problem definition and knowledge translation as crucial tasks within the type of application effort performed in this thesis. The reader may refer back to Section 3.6: Collaboration work (page 105) where we treat both problem definition and knowledge translation as the most consequential tasks within the collaboration work performed during our application effort. Also in that section, we proposed an iterative test procedure to improve the translation process.

4.3.2 Other more advanced production scenarios

In general, the optimization methodologies already presented in Chapter 3 may be extended to deal with more advanced production scenarios than the ones treated so far. For the development of the Martin Linge oil reservoir a possible production scenario is to equip the horizontal wells with inflow control valves (ICVs; also referred to as devices, ICDs). In this configuration, the wells will be rigged with several flow valves placed along the well bore. The objective for this configuration is to encourage a steady displacement of oil towards the well bore, and avoid or delay water breakthrough, by managing the settings of the valves. A reasonable formulation of an optimization problem for this production scenario would represent the setting of each valve as a control variable. Furthermore it would be interesting to treat the problem with a formulation that penalized water production, i.e., use NPV as cost function, and possibly use a discount factor to prioritize early production. In total, compared to the control problem currently treated in this work, a more interesting production scenario for control optimization could certainly be construed when operating wells using ICVs. Finally, we might further speculate that a more sophisticated production case like this one would have a greater influence on the master search for well placement, which, at least conceptually, lends support for a joint treatment.

4.3.3 Final comparison to current base case configuration

Throughout this work we have consistently used the base case model and well configuration provided to us by Total E&P Norge AS at the beginning of the study. In the meantime, the base case configuration has undergone further work by the field development team, in addition to receiving further adjustments based on suggestions from the drilling contractor. To create a coordination point between the results presented in this study, and the ongoing field development work, in Figure 4.34 we plot the updated base case well configuration together with the original base case well configuration used in this study. (It should be noted that the updated base case well configuration is the result of an independent work process that has not been influenced by the results in this work.) Using this reference graph, we can compare any of the well placement solutions developed in this study, via the original base case configuration, to the updated base case wells, if necessary. In Figure 4.34 we plot the J1M1H1 well placement solution as subject for further comparison.

4.3.4 Final comment on pilot study

This application effort has been a pilot study to test the developed optimization approach on a field case. In the preceding sections we have performed extensive tests of the obtained solutions to facilitate the consideration of these results by the development geoscience team at Total E&P Norge AS. Several factors affect the results from the pilot study, and should be considered in future developments of the optimization framework. By design, the application framework performs optimizations with limited scope, e.g., reduced production time frame, to make the solution process more effective. Moreover, a fundamental

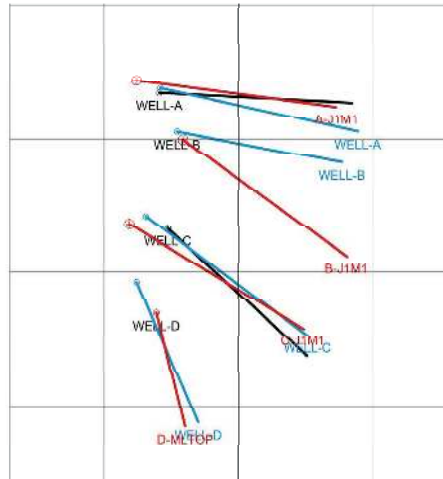


Figure 4.34: Original base case wells used in this work (black lines) are drawn together with updated base case wells (blue lines). Notice base case B wells overlap with blue line on top. Wells for hybrid JIM1H1 solution are also drawn (red lines).

limitation of the application framework design is that geological uncertainty is not considered during the optimization process. In our case, this causes a significant part of the benefit from the solutions to be lost when assessing the robustness of the various solutions against structural uncertainty. Additionally, several approximations were included in the creation of the work model used for optimization. The design of the optimization framework, combined with the reduced scope and the various approximations, impose severe restrictions on the use of the set of solutions achieved by the optimization process. Work is currently underway to further improve the optimization framework and core methodology to deal with these limitations, and to increase the applicability of the optimization process as a tool to assist well placement design in field development operations.

Chapter 5

Chapter Summaries and Further Work

This final chapter both summarizes the results from this thesis, and presents research topics for further work. It consists of two sections. The first section summarizes the results from each chapter individually. For each chapter summary, this section starts with some concluding remarks, then presents an outline of the main results, and finally provides a general critique about the work conducted in the respective chapter. Finally, the second section covers various topics of current and future research based on the work in this thesis.

5.1 Chapter summaries

This section offers itemized overviews of the main results from each of the preceding chapters. (Naturally, only concluding remarks, and no result outline, is provided for the introductory Chapter 1.) The various result outlines are prefaced by some concluding remarks that explain the overall purpose of the work. A by-product of this treatment is to place the different types of work performed in this thesis within a larger perspective of solving for integrated problems. Each chapter summary ends with a brief, general critique.

Summary of Chapter 1

We start this thesis with a brief overview of the field development work process. The point of this overview is to show that the development of a petroleum field involves a large number of different decisions, and that most of these decisions are challenging because they usually involve complex, interdependent systems. Crucially, these systems commonly require the allocation of substantial computational resources and engineering effort. We then outline the general aim of this work, which is to develop and test optimization methodology to support petroleum engineering work tasks performed in field development operations. In particular, the tasks treated in this thesis are directly related to the location and operation of wells for the production of hydrocarbon assets.

Specifically, we introduce the work in this thesis as targeting two crucial tasks for the development of a petroleum asset; namely, the overall well placement configuration for the production of the hydrocarbons, and the well control settings that determine how the

wells in the configuration are to be operated over the production time frame of the asset. In the introduction of this work, the issue of how, within a coupled optimization effort, the well control decision influences the search for optimal well placement configuration, is highlighted as a central and recurring topic throughout this thesis. Focusing on this topic, we describe how the overall work in this thesis revolves around testing for the significance, and cost, of jointly optimizing for well placement and controls. Furthermore, in this introductory chapter we emphasize the importance of comparing the joint approach against other approaches that perform the optimization sequentially, and that either simplify, or not take into account, the clear interdependency between the well placement and well control problems. Finally, we finish Chapter 1 by providing an overview of the various contributions from the study of this topic, taken both from method development (Chapter 2) and field case application (Chapter 3 and 4).

General critique. A general note regarding Chapter 1 is that the background work describing the various topics treated in this thesis, e.g., well placement and control optimization, optimization algorithms, etc., could clearly have been included within the structure of this chapter. Indeed, it would have been fitting to introduce this contextual text in the first chapter before we presented the overall description of the work in this thesis. However, we believe the very general nature of the discussion in Chapter 1, and the standalone and distinctive characters of chapters 2 and 3 (dealing with method development and field case application, respectively) make it suitable to rather introduce independent literature reviews at the beginning of each of these chapters. Combined, these separate literature reviews, found in sections 2.1 and 3.2, constitute the necessary thesis background, while best complementing their respective chapters, and overall thesis structure.

Summary of Chapter 2

The core of the methodology for the overall work in this thesis is developed in Chapter 2. In this chapter, methodology constructs, referred to as joint and sequential approaches, are defined as different ways to solve for the coupled well placement and control problem. From this chapter on, and throughout this thesis, the joint and the sequential fixed and reactive approaches are treated as opposing constructs, and their merit measured and discussed comparatively. The approaches are tested and contrasted through progressively more complex cases (though, the Chapter 2 cases are still relatively simple compared to the field case application in Chapter 3). The simplest case in Chapter 2 consists of a five well problem where the optimal location of an injector needs to be determined within a heterogeneous permeability field with four producers fixed at each corner of the reservoir. For this particular case, the low dimensionality of the well placement part of the problem allows the feasible space of injector well locations to be explored exhaustively. Three exhaustive searches are launched while using either gradient-based optimization, or fixed or reactive strategies, to deal with the well control part of the problem. The second fundamental case from this chapter optimizes the location and control scheduling of three producers and two injectors. This water–flooding configuration increases the diversity of production scenarios, and renders the search more challenging than the arrangement with only one injector. Results obtained from the comparison of the different approaches for the two cases, as well as some important observations from the work in

Chapter 2, are summarized below.

Control-optimized objective function. A key aspect of the joint scheme is that the objective function at the well placement outer loop is an optimized value of the cost function considered in the inner optimization of the well controls. Consequently, the joint scheme results in the solution of the outer optimization satisfying optimality conditions not only for the well placement problem but also for the well control part.

Smoothing of the optimization surface. For the first case, results for the three exhaustive explorations corresponding to each control strategy (fixed, reactive and optimized) yield a much rougher surface for the inner loop function associated with the fixed control strategy than the surfaces obtained with the other strategies. These results demonstrate that an efficient control strategy within the inner loop, e.g., gradient-based optimization, can significantly improve the performance of less promising locations, and consequently achieve a smoothing of the outer loop optimization surface with respect to the well placement variables.

Exhaustive search solutions. From the well locations with the highest net present value obtained from the three exhaustive explorations, we observe the joint scheme clearly out-performing the sequential methodologies, even after the additional gradient-based control optimization step. The joint approach yields an increase of 4.2% and 5.9% in NPV compared to the sequential fixed and reactive approaches, respectively. For this case, final oil saturations show less bypassed oil for the joint solution than the best solutions from the sequential approaches.

Reduction due to final control optimization step. Using the various approaches, searches for optimal injection well locations where launched from 12 different initial points. Results show that the additional control optimization to some extent reduces the differences in results obtained by the sequential and joint approaches. Before the control optimization step, the average optimized NPV by the joint approach is 24.1% and 9.8% larger than the average optimized NPVs from the fixed and reactive approaches, respectively. After the additional optimization, these percentages decrease to 10.3% and 6.1%.

Results from the more complex second case. Results when optimizing the location and controls of two injectors and three producers, show that, in terms of NPV, the sequential fixed and reactive strategies clearly under-perform the joint approach. The average (maximum) NPV over all of the runs obtained with the joint approach is 18.2% (14.4%) and 20.6% (7.3%) higher than with the sequential fixed and reactive schemes, respectively.

Smaller standard deviation of NPV results. Results from the second case show that the joint approach yields smaller standard deviation σ of the NPV than the sequential methodologies. This result is consistent with the smoothing of the well placement optimization landscape observed for the joint strategy in the first case.

Cost of joint versus sequential approaches. Despite the clear gain from the joint approach, we recognize its cost is significant compared to the sequential alternatives. Results from the second case show that the average number of simulations required by the joint approach is about one order of magnitude higher than that needed by the sequential methodologies.

Parallelization of well placement search. The increase in computational cost is somewhat mitigated by the implementation in parallel of some of the pattern search algorithms used for the well placement part of the approach. For example, the computation of two pattern search algorithms, parallelized using eight and 20 computing cores, yielded speedup factors of 4.1 and 6.4, respectively.

General critique. In total, the results from this chapter show that jointly optimizing for well placement and controls can yield significant improvements over sequential alternatives. However, these results are based on relatively simple cases treating only vertical wells. In general, additional work is needed on more complex cases to further validate these results. Furthermore, the cost of performing a joint type of approach is prohibitively high. This cost constitutes a serious obstacle towards any practical application of the joint approach, and needs to be further treated by additional research.

Summary of Chapter 3

Chapter 3 focuses on extending and applying the work developed in the previous chapter on to a real field case. The extension of the methodology to treat a more complex field case poses various technical challenges to the implementation developed in Chapter 2. Resolving these challenges, and the actual optimization of well placement and controls while using a real field case, are the main contributions from Chapter 3. However, since dealing with a real field case and treating with the IO Center Industry Partner operating this asset, before these challenges are dealt with, we have emphasized the need for there to be an overall frame of understanding regarding both the research development and the field case application effort performed in this thesis. For this reason, at the beginning of Chapter 3, we have focused on structuring our thinking into a consistent framework that uses the general concept of strategy within technology management to both clarify targets, and chart a clear course of action for the application work. Using this general management structure, the work in this and subsequent chapters consists of creating sufficient alignment between the stated application targets, and the various work processes (i.e., tactics) set up to resolve the defined challenges.

The first section of Chapter 3 starts with definitions of targets and strategy for the research development and field case application effort. Within the wider perspective of technology management, Chapter 3 proceeds by identifying challenges (both technical and collaboration-based), and by defining work processes to deal with the more complex real field case and overall organizational work. These work processes are organized into a work process loop set up to guide the entire application effort; from work model validation and problem definition, to optimization effort and solution testing. Finally, within this work process loop, we describe the technical framework for the optimization effort

that includes all the algorithms, as well as the technical extensions developed to accomplish the field case application work. Below we summarize the results achieved by the application of this optimization framework.

Mean FOPT increase from joint optimization runs. A total of nine different joint and sequential optimization runs have been performed using different configurations for the well-length constraint. Comparing the result of each sequential run to its corresponding joint counterpart, we observe the joint solutions yield higher FOPT increases for each solution pair. From our optimization effort using an approximated work model, the FOPT values for the joint solutions result in a mean increase in FOPT of close to 33% over the FOPT from the base case configuration.

Mean FOPT increase from sequential optimization runs. In this optimization case, FOPT values corresponding to the sequential runs result from an optimization of controls performed after a well placement configuration has been found using fixed controls. Averaging these FOPT values, and comparing them to the base case FOPT, we obtain a mean increase in FOPT of almost 26% for the sequential optimization runs.

Comparison joint vs. sequential runs. As mentioned before, a main axis for analysis is the comparison between the joint and the sequential approaches. For this particular problem case, the solutions obtained using the joint approach yield, on average, a 7% higher FOPT increase compared to the solutions obtained using the sequential approach. However, as for the second case in Chapter 2, the cost of the joint approach is significant compared to the sequential alternative. For this case we have that the mean total number of reservoir simulations required by the joint runs is almost 7 times higher than the mean number of reservoir simulations required by the sequential runs (i.e., we have averages of 6973 and 1025 total number of reservoir simulations for the joint and sequential approaches, respectively).

Influence of control optimization routine for each approach. For the work model case in Chapter 3, we have that the mean FOPT increases from the control routine are modest when optimizing using either the initial well configuration, or any of the final well configurations obtained from the sequential runs. However, based on function evolution data for the joint runs, we have that the control routine, once embedded within the well placement search, yields, on average, considerable higher increases when applied over whole ranges of different well placement trial solutions. This result agrees with our previous observation from Chapter 2 results regarding the possible smoothing of the outer loop optimization surface. Correspondingly, the results from this work model case support the notion that a smoothing of the outer loop optimization surface with respect to the well placement variables may be achieved, when the performance of less promising locations during the well placement search is improved (sometimes significantly) by the use of an efficient control strategy, in this case gradient-based optimization, within the inner loop.

Results associated with control optimization routine. Function evolution data shows that when using the joint approach, each well placement trial solution obtains a

general improvement in objective function value due to the inner loop control procedure of, on average, 7.5%. In contrast, only a mean increase of slightly more than 1% is observed when performing the control optimization part using either the initial well configuration, or any of the final well configurations obtained from the sequential runs.

Main observation from cost function evolution graphs. One comparison made of cost function evolutions curves is between the curves corresponding to the joint and sequential solutions, in terms of number of cost function calls for each of the two approaches. In this comparison, we notice each joint solution displays a better-performing cost function evolution curve, in terms of quicker progression and higher final FOPT, than its corresponding sequential cost function curve.

Transfer of solutions to real field case. The final sections of Chapter 3 focus on transferring the solutions obtained using the work model on to the original field case model. Due to the various approximations in the work model, we confirm a general decrease in cost function values once solutions are transferred over to the Eclipse field case model.

Joint vs. sequential comparison across solution transfer. Notably, we see that the joint runs still out-perform their sequential counterparts even after the various solutions are transferred to the field case model. Transfers implementing the complete solutions, and only the well placement part together with the original field case well schedule settings, are tested. For the first type of transfer, the joint solutions yield a mean increase in FOPT of slightly more than 25%. In comparison, the corresponding sequential solutions yield, on average, an increase in FOPT of somewhat less than 20%. Results for the second transfer have similar differences.

General critique. In summary, dealing with a real field case in the application of our methodology has posed significant difficulties with respect to implementation, but it has also allowed us to deal with at least two important challenges that further develop our overall research. The first challenge covers the application of the methodology on to more complex cases. Through this field case application we have shown that the joint approach yields significant improvements over sequential alternatives, also when optimizing horizontal well trajectories with various well placement design constraints. The second challenge has to do with the real field case application itself, in that the vast complexity of field development operations has required substantial thinking regarding organization as well as around research problem definition, planning and development. In a whole range of fundamental ways, from problem definition to consideration of results, the resolution of the second challenge has contributed significantly to the treatment of the first challenge, and vice versa.

However, in terms of final application and overall results, we see from the treatment in this chapter that the joint approach is only viable if we impose substantial approximations to the work model used for optimization. As mentioned in our general critique of the results in Chapter 2, more work is required to significantly reduce the computational

cost of the embedded routine while retaining the overall gains achieved by the joint approach. In total, we have that the joint approach is not yet efficient enough for practical implementations involving large scale problems.

Summary of Chapter 4

In Chapter 4 we test the well placement solutions developed in Chapter 3 on the original base case Eclipse model for the Martin Linge oil reservoir. As such, this chapter completes the second target specified for the application work in this thesis. Recall that a main element of the second target was to further shape (in this case by rigorous testing) the output of the optimization effort into a product that could potentially contribute to the overall field development work process of the industry operator. To this end, the entire Chapter 4 is devoted to testing and analyzing all solutions obtained by the optimization process with respect to various considerations important within the perspective of field development operations. In particular, the various solutions are additionally tested for two realistic field development considerations that were not included in the optimization effort, i.e., the original larger production time frame and a multiple realizations field case scenario. Finally, we have put significant emphasis on treating the results from the solution tests in such a way that the various representations and analyses are similar to pieces of information commonly used by the engineering team. The overall purpose of this treatment is to make the obtained information readily available for further evaluation within the work process of the industry partner. All input configurations involving the various solutions, as well as all output data associated with the different simulation test cases for this chapter, have therefore been re-created, run, and plotted using industry-standard tools and/or formats.

Chapter 4 starts by introducing the eight solution cases that have resulted from implementing the various solutions within the field case model. Resulting production profiles for the different wells from each configuration are given, in addition to saturation maps at the main two test time horizons (i.e., 1200 and 5174 days). Chapter 4 proceeds by studying the various solution cases collectively, using tables presenting differences in field and well oil production totals against base case values. These changes in production are also plotted against the sum and individual well lengths associated with each of the solution well configurations. Finally, all solution cases are tested over a multiple realization case scenario. At the end of this chapter, we also apply some simple heuristic rules as a way to modify and/or adapt the obtained solutions to the current engineering considerations within the field development work process of the operator (that were either approximated, or not included, in the problem formulation). The main results from this chapter are summarized below.

Overall analysis of individual solution well configurations.

- A number of solutions well configurations arrive at A well locations that have close to, or practically the same, positioning as the base case A well, but with a much shorter well bore. Still, these shorter solution wells have production profiles that closely match those of the base case A well.

- Those configurations that yield the greatest increases in field oil production total for both the 1200 day and the 5174 day time frames have B wells that aggressively target the eastern lobe of the Martin Linge oil reservoir.
- In practically all solutions, the oil production rates obtained from solution D wells are significantly larger than the oil production rate from the base case D well. This result is tightly connected with the fact that solution D well trajectories are in the majority of cases significantly longer and rotated counter-clockwise in a south-east alignment compared to the original base case D well trajectory.

Comparative analysis of solution field oil production totals.

- For the 1200 day production time frame, the JNT2M1 well placement solution yields the greatest increase in field oil production total with a 25.5% increase in FOPT compared to the base case configuration. For the same time frame, solutions JNT2OPT, JNT1M1 and FXD2OPT are close behind with corresponding increases in FOPT of 24.9%, 23.6% and 23.5%, respectively.
- For the 5174 day production time frame, both the JNT2M1 and the JNT2OPT solutions yield the highest increases in field oil production total at 13% compared to base case. A close second is the JNT2OPT2 solution with an increase in FOPT of 12.8%.

FOPT increases vs. changes in well length.

- While having total drainage lengths (the sum of all well lengths), more than 4% shorter than the total drainage length for the base case configuration, the JNT2CUT and the FXD2CUT solutions yield roughly a 10% increase in FOPT for the 5174 day production time frame.
- Having about the same total drainage length as the base case, the FXD2OPT2 solution yields FOPT increases of more than 20% and 10% for the 1200 day and the 5174 day production time frames, respectively.
- While having a 4% increase in total well drainage length, the JNT2OPT2, JNT2OPT and FXD2OPT solutions yield FOPT increases of about 25% and 12% for the 1200 day and 5174 day production horizons, respectively.
- For the 5174 day time frame, solution B wells with about 5% to almost 20% shorter well bore lengths than the base case B well yield WOPT increases ranging from 20% to 30%.

Solution tests on multiple realizations.

- Within the 1200 day production time frame, some solutions perform poorly for most realizations, but at least two solutions (JNT2M1 and JNT2OPT) retain increases in mean field oil production total of close to 7% over the base case configuration.

- Test results confirm that, when applied for the larger production horizon, more than half of the solutions yield mean FOPT results lower than those achieved by the base case well configuration. A final note is that the JNT2M1 solution manages to retain a roughly 3% increase in mean FOPT, while several solution wells show interesting production increases in their individual performances.
- Hybrid solutions containing specifically selected parts from the various individual solutions were developed to answer to the time frame and multiple realization considerations. The best solutions from this simple heuristic technique yield increases in mean field oil production total of about 8% for the 1200 day production time frame, and of roughly 4% for the 5174 day production time frame, both over corresponding base case mean FOPTs, respectively.

General critique. The general critique of the work in this chapter mainly follows the one made of Chapter 3. At a fundamental level, the real field case application of the overall methodology developed in this thesis is seriously hampered by the various approximations actually necessary to perform the optimization work. Substantial improvements, in both the performance and the usability of the application (the latter mostly referring to how well knowledge and problem specifications are transferred across inter-disciplinary boundaries), are required for the methodology and its results to receive proper recognition and general acceptance for use within a field development work process. In terms of performance, the main obstacle at the current level of method development and application capability revolves around the high cost of the approaches. Corollaries of this issue are the time frame reduction that was necessary to implement the approaches, and, of course, the various limitations imposed by the very expensive cost function calls, e.g., the inability to optimize using multiple realizations, and the considerable time it takes from a solution is launched until it can be presented to the field development engineering team for further treatment. Some of these issues are the subject for further work, which is discussed below.

5.2 Topics for further work

In this section we outline several topics that are currently being explored, or that will be treated in further work. To facilitate the discussion, we make a rough differentiation between topics dealing with the cost of implementing the approaches, and developmental topics that try to enhance the work in new, interesting ways. Clearly, both types of topics deal with advancing the work, but discuss the overall development from different perspectives. Topics regarding cost are mostly focused on developing techniques to tackle the high cost of the approaches. Some of the consequences related to the high cost of implementation have already been mentioned in the general critique of Chapter 4, where cost was identified as the main obstacle towards more extensive use of the methodology, e.g., when dealing with real field applications. The cost topics discussed below are: dealing with integrated problems; increased parallelization and taking further advantage of the modularity of the optimization framework; and a more advanced implementation of pattern search methods, and/or the introduction of other algorithms for well placement optimization. Developmental topics, on the other hand, are focused on extending the work

in the thesis in several interesting directions. These topics discuss the introduction of new cases and emphasize adding further complexity to the problems treated by the optimization framework. Topics in this context are: well control optimization using inflow control valves (ICVs); the development of surrogate models for the embedded control routine, and in general, as reservoir model approximations; and finally, dealing with reservoir model uncertainty. Notice that (somewhat) detailed descriptions of several of these topics have already been given at various places throughout this thesis. For this reason, and given the outline perspective of this section, we provide only brief comment on the various topics, and rather refer to those sections in this thesis where these topics are treated in greater length (where appropriate, we will provide references to relevant literature). We start our discussion with topics regarding cost.

5.2.1 Cost-reducing topics for further use of optimization framework

In this work, while developing the framework to solve for the core problem of joint well placement and control optimization, it has been natural to generalize the framework as a solution structure that, in a more robust and expansive version, can possibly also be used to solve for larger and/or other types of integrated problems. Along these lines, the general thought is that further development can eventually lead to a multi-purpose procedure that may be applied to solve problems with similar integrated structure, e.g., problems involving the combination of well placement and/or controls with the design of facilities and pipe network and/or the routing of well streams (as discussed in Section 1.1.2, page 3). Obviously, this is a high ambition, given that for the size of these problems, the emphasis on using the joint approach makes any solution effort exceedingly costly. For the basic version of the procedure presented in this work, when confronted with the increased computational cost of dealing with a real field case, we proceeded with the straightforward effort of extending the parallelization of the pattern search methods (see Section 3.2.1, page 56). The continuation of this effort as a cost-reducing measure is therefore an important topic for further work. The general focus for this work should be on the implementation of other algorithms with improved capability for distributed computing. Furthermore, in the search for better-performing algorithms, we should consider not only methods based on pattern search, but also explore the use of techniques that rely on global exploration, such as particle swarm optimization and/or genetic algorithms, that, as discussed in Section 3.2.1 (page 56), have also been used for the well placement part of the procedure. Notice, however, that the implementations of the non-linear well placement constraints used in this work may not be compatible with the stochastic nature of these algorithms, and that different implementations of these constraints may have to be developed if applying these global search techniques. As also mentioned in Section 3.2.1, it may be fruitful to explore more advanced implementations of pattern search techniques that perform more refined searches of the well placement space, e.g., methods that use cost function values from polling searches to approximate gradient directions. Finally on this topic, we should explore the field of bilevel optimization theory for insights that may improve the general performance of the proposed solution structure.

Moreover, future work should actively exploit the overall solution structure, particu-

larly in terms of the modular property of the optimization framework (see Section 3.2.1, page 53). A particular advantage of this property is that it can be used to accelerate overall performance by facilitating the introduction of surrogate models at different levels of the optimization framework. At the level on control optimization, a basic idea for further work is to estimate some of the optimization cost function values using a reactive strategy that may be “tuned” based on previously computed joint optimization results (see Section 2.5, page 40). Moreover, additional gains in performance may be achieved by further exploiting the inter-level relationship between the problems. Along these lines, work is currently underway that builds reduced-order models from sets of control solutions and uses these as surrogates within the control optimization routine (we refer to this enhancement as Joint+RCO; see Section 3.2.1, page 57, for a full description). Finally, a related topic that may be considered in future applications is the straightforward use of surrogate reservoir models as approximations to the full physics model during the whole, or parts, of the optimization procedure. Possibly, these surrogates can be upscaled variants of the original model (Nakashima and Durlofsky, 2010), or be developed by, e.g., using some of the advanced flow diagnostic tools developed by SINTEF (Lie et al., 2012).

5.2.2 Topics for further development of optimization framework

As mentioned, the work in this thesis can be extended to treat larger, more complex integrated petroleum field problems. Crucially, this extended use is highly dependent on whether we are able to satisfactorily advance some of the cost-reducing topics described above. Disregarding this challenge for a moment, we notice there are other problems of interest that may be treated within the optimization framework. In particular, we refer to the introduction of more complex production cases (still within the framework of the integrated problem). For example, a more complex case that could be treated involves the optimization of controls when using inflow control valves in a production scenario with horizontal wells (see Section 4.3.2, page 178, for a highly relevant case example based on a Norwegian Continental Shelf development). For this type of work, it might be beneficial to complement the treatment of the more advanced production scenario with a study that specifically targets the optimization of controls subject to the type of non-linear production constraints likely to be applied to that case, thus matching increased case complexity for research purposes with realistic field considerations. In relation to this case, general research is currently underway that studies the characteristics of possible constraint formulations for this type of control optimization problems. (See also Section 3.4.2, page 84, for a brief discussion of different types of constraints for this problem.) For example, an interesting topic for research is whether (well and/or field) production constraints, e.g., water cut thresholds for shutting down wells, or water production limit rates, should be implemented at the solver level or within the reservoir simulator (as they often are, given the efficiency of such heuristic treatment). Furthermore, an important question in this regard is how each of these two approaches for constraint implementation affect the evolution of gradient-based optimization procedures. In particular, how is the accuracy of the adjoint gradient computation influenced by the heuristic treatment of production constraints within the forward reservoir simulation, and, if there is substantial degradation of the derivatives, how does the inconsistency between input controls and cost function

gradients affect the gradient-based technique driving the optimization. In closing, this overall topic is also coupled with the possible application of other gradient-based solvers for the specific purpose of decreasing the total number of calls to the reservoir simulator (see Section 3.2.1, page 57). Finally, as previously discussed, if further method developments are to be readily applied within field development operations, we need to take into account reservoir model uncertainty within the optimization process, e.g., by using the method by Wang et al. (2012), briefly discussed in Section 2.1, on page 16.

Summary of results

We end this thesis by providing a brief outline of the main results from this work. These results are presented according to the two conceptual parts that divide this thesis: method development and field case application. For a comprehensive summary of each chapter in this thesis, and of the results therein, the reader is referred to the previous chapter, Chapter 5.

Main results from part I: Method development

- Simple case results clearly show the joint scheme out-performing the sequential methodologies, even after the additional gradient-based control optimization-step performed by the sequential approaches. Exhaustive explorations of the well placement search space for a low-dimensional five-spot case yield optimal cost function values (NPV) that are 4.2% and 5.9% higher compared to the sequential fixed and reactive approaches, respectively.
- When optimizing the location and controls for a larger case including two injectors and three producers, we notice the performance of the sequential fixed and reactive strategies decrease compared to the joint approach. For this more complex case, we observe the average cost function of all of the runs obtained with the joint approach is 18.2% and 20.6% higher than with the sequential fixed and reactive schemes, respectively.
- For the establishment of the joint approach against sequential procedures, the clear trade-off between attaining higher objective function values and the increased computational cost this entails is an important consideration. For the five-well problem in the methodology part of this thesis, the average number of reservoir simulations required by the joint approach is about one order of magnitude higher than the computational resources needed by the sequential methodologies. However, a conscious design feature for the development of the methodology has been to mitigate the increased cost by implementing in parallel some of the pattern search algorithms used for the well placement part of the approach. For the five-well problem, the parallelization of these algorithms yields speedup factors of 4.1 and 6.4, respectively.

Main results from part II: Field case application

- Nine different joint and sequential optimization runs are launched using a work model approximation of the Martin Linge field case. A set of high-performance configurations, rather than a single best solution, allows the operator to choose the well configuration that best fits the broader well strategy of the development plan.
- Average cost function values (field oil production total; FOPT) for the joint solutions yield a mean increase in FOPT of close to 33% compared to the base case configuration. In contrast, the solution well configurations corresponding to the sequential runs result in a mean FOPT increase of almost 26%. For this optimization effort, we have that the joint runs require about 7 times more reservoir simulations than sequential runs.
- Cost function evolution data for each run show that the mean contribution from the embedded control routine is significantly higher compared to the average gain produced by the control optimization step for the sequential procedure. This result is consistent with similar observations made for the simpler cases, and points to a possible smoothing of the optimization surface with respect to the well placement variables due to the nested routine. Possibly, this smoothing may add some robustness to the overall well placement search conducted by the joint approach.
- The well configurations obtained using the work model are subsequently transferred to the field case model. Due to the multiple approximations necessary for the work model implementation, a general decrease in cost function values across all solutions is observed once these configurations are tested on the original reservoir model. Furthermore, the solutions are also tested for a larger production time frame (5174 days) than the one used during optimization (1200 days), and a field case scenario involving a set of 11 model realizations.
- For the 1200 day time frame, the JNT2M1 solution yields the greatest increase in FOPT with a 25.5% increase compared to the base case configuration. For the 5174 day time frame, both the JNT2M1 and the JNT2OPT solutions yield the highest increases in FOPT at 13% compared to base case. However, further test results show that, when applied for both for the larger 5174 day horizon and over the whole range of model realizations, the set of solutions lose their significant gains over the base case well configuration. Among other, this result underscores the importance of using more than only a single realization during the optimization procedure.
- A key result from the individual tests of well placement solutions is that those well placement configurations that yield the greatest FOPT increases have B wells that specifically target the eastern lobe of the Martin Linge oil reservoir. Furthermore, for the 5174 day time frame, solution B wells with about 5% to almost 20% shorter well bore lengths than the base case B well yield WOPT (well oil production total) increases ranging from 20% to 30%.
- For most solutions, the oil production rates obtained from solution D wells are significantly larger than the oil production rate from the base case D well. This

increase is mostly attributed to a substantially longer well bore and a significant rotation of this well compared to the base case D well trajectory. However, based on updated field studies, the lengths and realignments of these solution wells may be too severe with respect to current field development constraints. These wells should therefore be reconsidered within the routine, e.g., by further refining the overall boundary constraints for this particular well during optimization.

Bibliography

- Audet, C., Dennis Jr., J. E., March 2000. Pattern Search Algorithms for Mixed Variable Programming. *SIAM Journal on Optimization* 11 (3), 573–594.
- Audet, C., Dennis Jr., J. E., 2002. Analysis of Generalized Pattern Searches. *SIAM Journal on Optimization* 13 (3), 889–903.
- Audet, C., Dennis Jr., J. E., 2006. Mesh Adaptive Direct Search Algorithms for Constrained Optimization. *SIAM Journal on Optimization* 17 (1), 188–217.
- Aziz, K., Settari, A., 1979. *Petroleum Reservoir Simulation*. Applied Science Publishers, New York, NY.
- Bangerth, W., Klie, H., Wheeler, M., Stoffa, P., Sen, M., 2006. On Optimization Algorithms for the Reservoir Oil Well Placement Problem. *Computational Geosciences* 10 (3), 303–319.
- Barton, P. I., 1992. *The Modelling and Simulation of Combined Discrete-Continuous Processes*. Ph.D. thesis, Imperial College of Science, Technology and Medicine, London, UK.
- Beckner, B. L., Song, X., 1995. Field Development Planning Using Simulated Annealing - Optimal Economic Well Scheduling and Placement. In: *SPE Annual Technical Conference and Exhibition*. Dallas, TX, SPE 30650.
- Bellout, M., Echeverría Ciaurri, D., Durlofsky, L. J., Foss, B., Kleppe, J., 2012. Joint Optimization of Oil Well Placement and Controls. *Computational Geosciences* 16, 1061–1079.
- Boutaud de la Combe, J.-L., Kvinnsland, S., Marmier, R., Morante-Gout, J., 2012. Hild Extended Well Test: A Key Acquisition for the Field Development Decision. In: *SPE Annual Technical Conference and Exhibition*. San Antonio, TX, SPE 159275.
- Bratvold, R. B., Bickel, J. E., Kullawan, K., 2014. Value Creation with Multi-Criteria Decision Making in Geosteering Operations. In: *SPE Hydrocarbon Economics and Evaluation Symposium*. Houston, TX, SPE 169849.

-
- Brouwer, D. R., Jansen, J. D., 2004. Dynamic Optimization of Waterflooding With Smart Wells Using Optimal Control Theory. *SPE Journal* 9 (4), 391–402, SPE 78278.
- Cao, H., Jun. 2002. Development of Techniques for General Purpose Simulators. Ph.D. thesis, Dept. of Petroleum Engineering, Stanford University, Stanford, CA.
- Cardoso, M. A., Durlofsky, L. J., Feb. 2010. Linearized reduced-order models for subsurface flow simulation. *J. Comput. Phys.* 229 (3), 681–700.
- Centilmen, A., Ertekin, T., Grader, A., 1999. Applications of Neural Networks in Multiwell Field Development. Houston, TX, SPE 56433.
- Chang, S., Moyer, M., 2010. Lustre Best Practices: How Does Lustre I/O Work? Retrieved February Thursday 13, 2014, from http://www.nas.nasa.gov/hecc/support/kb/Lustre-Best-Practices_226.html.
- Christie, M. A., Blunt, M. J., 2001. Tenth SPE Comparative Solution Project: A Comparison of Upscaling Techniques. *SPE Reservoir Evaluation & Engineering* 4, 308–317, SPE 66599.
- Clerc, M., 2006. Particle Swarm Optimization. Wiley-ISTE, London.
- Conn, A. R., Scheinberg, K., Vicente, L. N., 2009. Introduction to Derivative-Free Optimization. MPS-SIAM Series on Optimization. Society for Industrial and Applied Mathematics and the Mathematical Programming Society, Philadelphia, PA.
- Custodio, A. L., Vicente, L. N., 2007. Using Sampling and Simplex Derivatives in Pattern Search Methods. *SIAM Journal on Optimization* 18 (2), 537–555.
- Dake, L. P., 1978. Fundamentals of Reservoir Engineering. Elsevier, Amsterdam.
- Dempe, S., 2002. Foundations of Bilevel Programming. Kluwer Academic Publishers, New York, NY.
- Doren, J., Markovinović, R., Jansen, J.-D., 2006. Reduced-order optimal control of water flooding using proper orthogonal decomposition. *Computational Geosciences* 10, 137–158.
- Douillard, A., Way, S., Arnaud, J., Adamsen, M., Mikkelsen, G., 2009. Derisking Hild Field Development Through Wide-azimuth OBC and Multi-azimuth Streamer Depth Imaging. In: 71st EAGE Conference & Exhibition. Amsterdam, The Netherlands.
- Eberhart, R. C., Kennedy, J., Shi, Y., 2001. Swarm Intelligence. Morgan Kaufmann, Waltham, MA.
- Echeverría Ciaurri, D., Conn, A., Mello, U. T., Onwunalu, J. E., 2012. Integrating Mathematical Optimization and Decision Making in Intelligent Fields. In: Proceedings of the SPE Intelligent Energy International. Utrecht, The Netherlands, SPE 149780.

-
- Echeverría Ciaurri, D., Isebor, O. J., Durlofsky, L. J., 2011a. Application of Derivative-Free Methodologies for Generally Constrained Oil Production Optimization Problems. *International Journal of Mathematical Modelling and Numerical Optimisation* 2 (2), 134–161.
- Echeverría Ciaurri, D., Mukerji, T., Durlofsky, L. J., 2011b. Derivative-Free Optimization for Oil Field Operations. In: Yang, X.-S., Koziel, S. (Eds.), *Computational Optimization and Applications in Engineering and Industry*. No. 359 in *Studies in Computational Intelligence*. Springer, Berlin/Heidelberg.
- Ertekin, T., Abou-Kassem, J. H., King, G. R., 2001. *Basic Applied Reservoir Simulation*. SPE Textbook Series. SPE, Richardson, TX.
- Fletcher, R., Leyffer, S., Toint, P., 2006. A Brief History of Filter Methods. Tech. Rep. ANL/MCS/JA-58300, Argonne National Laboratory.
- Forouzanfar, F., Li, G., Reynolds, A. C., 2010. A Two-Stage Well Placement Optimization Method Based on Adjoint Gradient. In: *Proceedings of the SPE Annual Technical Conference and Exhibition*. Florence, Italy, SPE 135304.
- Foss, B., Gunnerud, V., Dueñas Díez, M., 2009. Lagrangian Decomposition of Oil-Production Optimization Applied to the Troll West Oil Rim. *SPE Journal*, 646–652SPE 138891.
- Gill, P. E., Murray, W., Michael, Saunders, M. A., 1997. SNOPT: An SQP Algorithm For Large-Scale Constrained Optimization. *SIAM Journal on Optimization* 12, 979–1006.
- Gill, P. E., Murray, W., Saunders, M. A., Jan 2005. SNOPT: An SQP Algorithm for Large-Scale Constrained Optimization. *SIAM Review* 47 (1), 99–131.
- Gill, P. E., Murray, W., Saunders, M. A., Apr. 2007. User's Guide for SNOPT Version 7: Software for Large Scale Nonlinear Programming. Stanford University, Stanford, CA, and University of California, San Diego, CA.
- Goldberg, D. E., 1989. *Genetic Algorithms in Search, Optimization and Machine Learning*. Addison-Wesley Professional, Boston, MA.
- Griffin, J. D., Kolda, T. G., 2007. Nonlinearly-Constrained Optimization using Asynchronous Parallel Generating Set Search. Tech. Rep. SAND2007-3257, Sandia National Laboratories, Livermore, CA.
- Griffin, J. D., Kolda, T. G., Lewis, R. M., May 2008. Asynchronous Parallel Generating Set Search For Linearly-Constrained Optimization. *SIAM Journal on Scientific Computing* 30 (4), 1892–1924.
- Güygüler, B., Horne, R. N., Rogers, L., Rosenzweig, J. J., 2000. Optimization of Well Placement in a Gulf of Mexico Waterflooding Project. In: *Proceedings of the SPE Annual Technical Conference and Exhibition*. Dallas, TX, SPE 63221.

-
- Handels, M., Zandvliet, M., Brouwer, R., Jansen, J., 2007. Adjoint-based Well-placement Optimization under Production Constraints. SPE Reservoir Simulation Symposium SPE 105797.
- He, J., Sætrom, J., Durlofsky, L. J., Sep. 2011. Enhanced Linearized Reduced-order Models for Subsurface Flow Simulation. *J. Comput. Phys.* 230 (23), 8313–8341.
- Hooke, R., Jeeves, T. A., 1961. Direct Search Solution of Numerical and Statistical Problems. *Journal of the ACM* 8 (2), 212–229.
- Isaacs, W., 1999. *Dialogue and the Art of Thinking Together*. Doubleday/Random House, New York.
- Jansen, J. D., Brouwer, D. R., Naevdal, G., van Kruijsdiik, C. P. J. W., Jan. 2005. Closed-loop Reservoir Management. *First Break* 23, 43–48.
- Juell, A., Whitson, C. H., Hoda, M. F., 2010. Model-based Integration and Optimization – Gas-cycling Benchmark. *SPE Journal*, 646–657 SPE 121252.
- Koen, B. V., 2003. *Discussion of the Method: Conducting the Engineer’s Approach to Problem Solving*. Oxford University Press, New York, NY.
- Kolda, T. G., Lewis, R. M., Torczon, V., 2003. Optimization by Direct Search: New Perspectives on Some Classical and Modern Methods. *SIAM Review* 45 (3), 385–482.
- Kourounis, D., Durlofsky, L. J., Jansen, J. D., Aziz, K., 2014. Adjoint Formulation and Constraint Handling for Gradient-based Optimization of Compositional Reservoir Flow. *Computational Geosciences* 18 (2), 117–137.
- Læg Reid, T., June 2001. *Technology Strategy and Innovation Management in the Petroleum Industry*. M.Sc. Thesis, Alfred P. Sloan School of Management, Massachusetts Institute of Technology, MA.
- Lie, K., Krogstad, S., Ligaarden, I. S., Natvig, J. R., Nilsen, H. M., Skaflestad, B., 2012. Open-source MATLAB Implementation of Consistent Discretisations on Complex Grids. *Computational Geosciences* 16 (2), 297–322.
- Litvak, M. L., Hutchins, L. A., Skinner, R. C., Darlow, B. L., Wood, R. C., Kuest, L. J., 2002. Prudhoe Bay E-Field Production Optimization System Based on Integrated Reservoir and Facility Simulation. In: *SPE Annual Technical Conference and Exhibition*. San Antonio, TX, SPE 77643.
- Møyner, O., Krogstad, S., Lie, K., 2014. The application of flow diagnostics for reservoir management. *SPE Journal* Accepted June 2014.
- Nakashima, T., Durlofsky, L., 2010. Accurate Representation of Near-well Effects in Coarse-scale Models of Primary Oil Production 83 (3), 741–770.
- Nocedal, J., Wright, S. J., 2006. *Numerical Optimization*. Springer, 2nd Edition, New York, NY.

-
- NPD, 2013a. Norwegian Petroleum Directorate Factpages: Martin Linge field. Retrieved October Thursday 31, 2013, from <http://factpages.npd.no/factpages/>.
- NPD, 2013b. Norwegian Petroleum Directorate: Martin Linge fact sheet. Retrieved October Thursday 31, 2013, from <http://www.npd.no/en/Publications/Facts/Facts-2013/Chapter-11/Martin-Lange/>.
- NTNU/IO Center, 2013. IO Center Annual Report. Tech. rep., NTNU Center for Integrated Operations in the Petroleum Industry, <http://www.iocenter.no/>.
- NTNU/IO-Center, 2014. NTNU Center for Integrated Operations in the Petroleum Industry. <http://www.iocenter.no/>.
- Onwunalu, J. E., Durlofsky, L. J., 2010. Application of a Particle Swarm Optimization Algorithm for Determining Optimum Well Location and Type. *Computational Geosciences* 14 (1), 183–198.
- Onwunalu, J. E., Litvak, M. L., Durlofsky, L. J., Aziz, K., 2008. Application of Statistical Proxies to Speed Up Field Development Optimization Procedures. In: *Proceedings of the SPE Abu Dhabi International Petroleum Exhibition and Conference*. Abu Dhabi, UAE, SPE 117323.
- Peaceman, D. W., Jun. 1978. Interpretation of Well-Block Pressures in Numerical Reservoir Simulation. *SPE Journal* 18 (3), 183–194, SPE 6893.
- Pironneau, O., 1974. On optimum design in fluid mechanics. *Journal of Fluid Mechanics* 64, 97–110.
- Plantenga, T. D., October 2009. HOPSPACK 2.0 User Manual. Tech. Rep. SAND2009-6265, Sandia National Laboratories, Livermore, CA.
- Powell, M. J. D., 2009. The BOBYQA Algorithm for Bound Constrained Optimization Without Derivatives. Tech. Rep. NA2009/06, Department of Applied Mathematics and Theoretical Physics, Cambridge, UK.
- Rahmawati, S. D., Whitson, C. H., Foss, B. A., Kuntadi, A., 2010. Multi-Field Asset Integrated Optimization Benchmark. In: *SPE EUROPEC/EAGE Annual Conference and Exhibition*. Barcelona, Spain, SPE 130768.
- Sarma, P., Chen, W. H., 2008. Efficient Well Placement Optimization with Gradient-based Algorithm and Adjoint Models. In: *Proceedings of the SPE Intelligent Energy Conference and Exhibition*. Amsterdam, The Netherlands, SPE 112257.
- Sarma, P., Durlofsky, L. J., Aziz, K., Chen, W. H., 2006. Efficient Real-time Reservoir Management using Adjoint-based Optimal Control and Model Updating. *Computational Geosciences* 10, 3–36.
- Schein, E., 1992. *Organizational Culture and Leadership*, 2nd Edition. Jossey-Bass management series. Jossey-Bass, San Francisco, CA.

-
- Schlumberger, 2012a. ECLIPSE Reference Manual.
- Schlumberger, 2012b. ECLIPSE Technical Description.
- Schlumberger, 2012c. Petrel E&P Software Platform v.2012.7.
- Svanberg, K., 2002. A Class of Globally Convergent Optimization Methods Based on Conservative Convex Separable Approximations. *SIAM Journal on Optimization*, 555–573.
- Tchelepi, H., Aziz, K., 2012. AD-GPRS Manual. Department of Petroleum Engineering, Stanford University, CA.
- Thibaut, E., Leforgeais, B., 2012. Martin Linge Electric Power From Shore. In: Abu Dhabi International Petroleum Conference and Exhibition. Abu Dhabi, UAE, SPE 161960.
- Torczon, V., 1997. On the Convergence of Pattern Search Algorithms. *SIAM Journal on Optimization* 7 (1), 1–25.
- Total E&P Norge AS, December 2011. Plan for utbygging og drift av Hild; del 2: Konsekvensutredning.
- Total.com, 2013. Martin Linge Project (Norway). Retrieved October Thursday 31, 2013, from <http://total.com/en/energies-expertise/oil-gas/exploration-production/projects-achievements/others/martin-linde>.
- TU, 2013. Teknisk Ukeblad: Over 40 år etter første funn skal det produseres på Martin Linge. Retrieved June Thursday 19, 2013, from <http://www.tu.no/petroleum/2013/10/22/over-40-ar-etter-forste-funn-skal-det-produseres-pa-martin-linge>.
- Vlemmix, S., Joosten, G. J., Brouwer, R., Jansen, J.-D., 2009. Adjoint-Based Well Trajectory Optimization. In: EUROPEC/EAGE Conference and Exhibition. Amsterdam, The Netherlands, SPE 121891.
- Volkov, O., Kourounis, D., 2012. Optimization Module of AD-GPRS. Tech. rep., Department of Petroleum Engineering, Stanford University, CA.
- Voskov, D., Zhou, Y., 2012. Automatic Differentiation based General Purpose Research Simulator. Tech. rep., Department of Petroleum Engineering, Stanford University, CA.
- Wachter, A., Biegler, L. T., 2006. On the Implementation of an Interior-point Filter Line-search Algorithm for Large-scale Nonlinear Programming 106 (1), 25–57.
- Wang, C., Li, G., Reynolds, A. C., 2007. Optimal Well Placement for Production Optimization. In: Proceedings of the SPE Eastern Regional Meeting. Lexington, KY, SPE 111154.
- Wang, H., Echeverría Ciaurri, D., Durlofsky, L. J., Cominelli, A., 2012. Optimal Well Placement Under Uncertainty Using a Retrospective Optimization Framework. *SPE Journal* 17 (1), 112–121, SPE 141950.

-
- Wild, S. M., 2009. Derivative-free Optimization Algorithms For Computationally Expensive Functions. Ph.D. thesis, Cornell University, Ithaca, NY.
- Yeten, B., Durlofsky, L. J., Aziz, K., 2003. Optimization of Nonconventional Well Type, Location, and Trajectory. SPE Journal 8 (9), 200–210, SPE 86880.
- Zandvliet, M., 2008. Model-based Lifecycle Optimization of Well Locations and Production Settings in Petroleum Reservoirs. Ph.D. thesis, TU Delft, The Netherlands.
- Zandvliet, M., Handels, M., van Essen, G., Brouwer, R., Jansen, J.-D., Dec. 2008. Adjoint-based Well-placement Optimization under Production Constraints. SPE Journal 13 (4), 392–399, SPE 105797.
- Zhang, K., Li, G., Reynolds, A. C., Yao, J., Zhang, L., Sep. 2010. Optimal Well Placement using an Adjoint Gradient. Journal of Petroleum Science and Engineering 73 (34), 220–226.

Notes

Chapter 1

1. The actual search procedure may be deterministic or stochastic, but the principle of changing the current iterate with another one that yields an improved objective, regardless of how the new iterate was found, is deterministic by logic.

Chapter 3

1. The use of the concept of strategy, and related terminology, helps structure the application work described in this section because it encapsulates dimensions of planning and agency that are important in the development and application of research. A strategy can be defined as “a collection of plans and processes that integrate [...] vision, overall goals, policies and actions in a coherent way” (Lægreid, 2001). Obviously, the aim of our strategy within the context of this application work is much more limited (though important for our effort), namely to establish a purpose and direction for the work conducted in this thesis.

2. Lægreid (2001) is a very interesting reference because the author is a Ph.D. in Physics from NTNU who worked many years in the R&D department of Statoil AS, before conducting a M.Sc. degree in the Management of Technology at Alfred P. Sloan School of Management, MIT.

3. The work process structure developed earlier in Section 3.1 can be seen to describe all the work tasks related to the creation and use of the optimization framework and its results.

4. This geological situation yields a relatively large number of inactive reservoir grid cells throughout the reservoir, specially at the border of the main reservoir sands, where sand quality is lower. This impacts the projection algorithm that keeps reservoir trajectories within bounds, since the well toes or heels of these trajectories need to be projected onto active grid cells. A check of where the closest-lying grid block is located is therefore implemented before we perform the actual projection, see Section 3.4.3.

5. A much closer match in recovery volumes was possible (specially with respect to field gas production). However, this work model configuration was chosen for optimization because it was faster, and its numerical solution was more robust with respect to different well locations and control settings, than the other approximations.

-
6. In this regard, a reasonable problem definition can be seen as one that includes the main aspects of the decision to be made, and that offers the field development team the type of information required
 7. However, we might speculate that joint solutions are somehow more robust against certain well placement changes, in this case changes to well length. If true, this type of robustness would be beneficial once solutions are considered within the more general field development work process, where constraints not included in the optimization process, are likely to require changes to the solutions proposed by the optimization procedure. In any case, more work would be needed to further test for this type of robustness.
 8. On this topic, the reader may refer to the three distinctions made regarding information types introduced in Section 3.2.1, page 61, that discussed translation of field case understanding into workable constraint definitions.

Chapter 4

1. In general, the first priority when searching for reasonable well trajectories during the course of field development planning is to maximize the recovery of hydrocarbons in place. During this search, a common strategy is to target those reservoir areas which have a high probability of containing significant accumulations of hydrocarbons. These probabilities are estimated using careful analysis of seismic data and of petrophysical information from discovery and appraisal wells. In this process it is important to minimize uncertainty. A common guiding procedure to minimize uncertainty about where to place production wells is to locate the potential well trajectories close to existing reservoir entry points that discovery and appraisal wells have proven contain hydrocarbon bearing sands. Once promising locations have been found, they are extensively studied using reservoir simulation.
2. This decision is calibrated against the present uncertainty surrounding the depth and curvature of the top reservoir surface.
3. In most cases, well design strategy will try to keep drainage length to a minimum, because longer wells are more expensive to drill and operate, have a larger well-bore pressure loss and face a greater risk of running beyond reservoir sands. Other more distant concerns also have significant impact on the ultimate field development cost and final decision of where to place the base case wells. For instance, the location of the platform, its planned well slot capacity, and drilling considerations such as well path curvature, all influence to varying degrees the final design and cost of the base case well trajectories.
4. Besides serving as calibration points for the horizontal drains, these pilot wells will also provide information about shales both above and below the top structure of the reservoir.
5. We note that for this particular solution, the constraint projection (see Section 3.5) has extended the trajectory of the D-F1M1 well to a minimum length of 1200 meters, such that the toe appears just outside the accumulation border in the mapping. Since cells are inactive outside the border, this positioning has no practical implication for the production from this well, except that this well is effectively somewhat shorter than 1200 meters.
6. A longer version of this well which extended its heel to reach a length of 1525 meters was discarded during transfer to Eclipse (simulation results for this test case all correspond to a configuration using this modified well length). The reason for the modification was

that this well slightly violated the length constraint by being somewhat longer than 1500 meters. A shorter version of this well around the solution was chosen for testing. This also allows us to compare longer well bores from other solutions that produce from the eastern lobe against a short well bore that produces from the same area.

7. By the same token, it can be said that we have disregarded a significant time period of water production in the middle and at the tail end production. In particular, a more extensive control optimization scheme can be devised with an objective function that includes production from the entire horizon, and that penalizes water production, e.g., in a net present value formulation. The definition of this type of optimization problem is the subject of future work.

8. See (Koen, 2003): “Discussion of the method: conducting the engineer’s approach to problem solving” for an ample discussion of the application of heuristics within the field of engineering.

9. Notice that we use the JNT1M1 solution as base in our example hybrid procedure, and not the JNT2M1 solution, because only the former solution was ready at implementation time.

Appendix A

Remaining production profiles from validation work

Remaining production profiles from Section 3.3.3: "Field model transfer and validation", are presented here. Figure 6.1 shows well gas production rates (WGPR) and totals (WGPT). Well oil production rates (WOPR) and totals (WOPT) are shown in Figure 6.2 while well water production rates (WWPR) and totals (WWPT) are shown in Figure 6.3. Finally, Figure 6.4 shows well liquid production rates (WLPR) and totals (WLPT).

Approximation of well group controls (cont'd)

Discretization of Eclipse model solution.

Resulting well bottom-hole pressures (WBHP) and well water production rates (WWPR) from Eclipse simulation case 3 with volume aquifer approximation (ECLbk_ovaq) are discretized over the entire production horizon. The accuracy of the discretization is checked by running an additional Eclipse simulation using only the discretized values, without gas lift the procedure nor well and field rate constraints (i.e. simulation case 3: ECLck_rvqw). In Table 3.4 we see we obtain reasonable comparison in cumulative volumes of gas, oil and water that are 0.994, 1.001 and 0.994, respectively, compared to base case volumes.

The entire production horizon is split into four different periods, and well bottom-hole pressures and water production rates are discretized within each period. Each period has a different discretization interval depending on the dynamic behavior of fluid production from the reservoir, e.g., high gas production early in the development requires a finer discretization than at later production times. The finer discretization at early production times is important to facilitate solver convergence in the AD-GPRS simulator. Figures 6.5 and 6.6 show the final discretization of WBHP and WWPR at each production period, respectively.

Application of discretized WBHP and WWPR

Figure 6.7 and 6.8 show how the original well control strategy is transformed into simpler form by the production quantity discretization. Figure 6.7 shows the control and constraint setup defined in the original field model. Figure 6.8 shows how the discretized controls enter as simulator inputs in Eclipse and AD-GPRS.

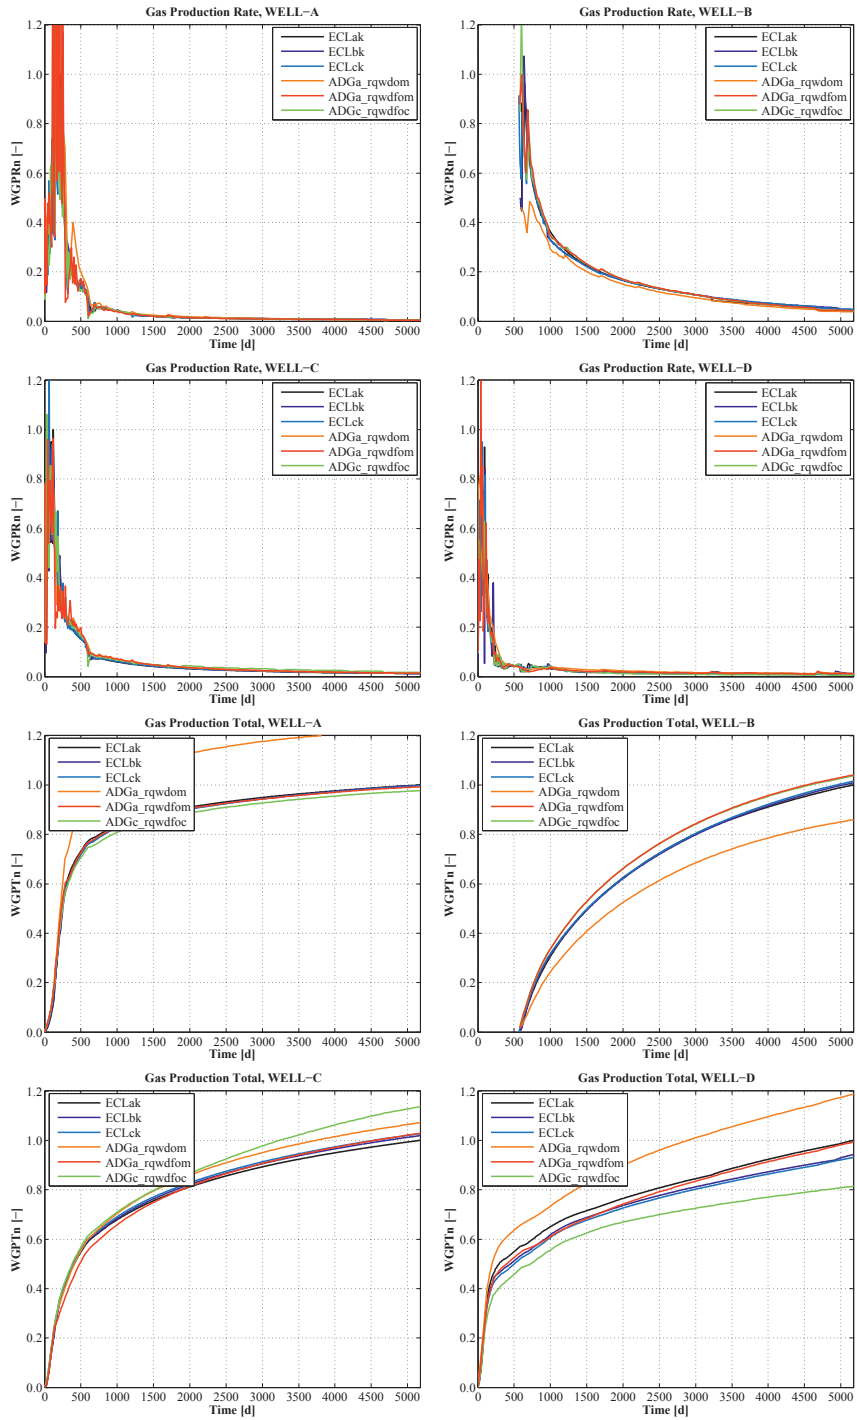


Figure 6.1: Well gas production rates and totals. Suffix “n” on production mnemonics signifies values have been normalized for confidentiality reasons.

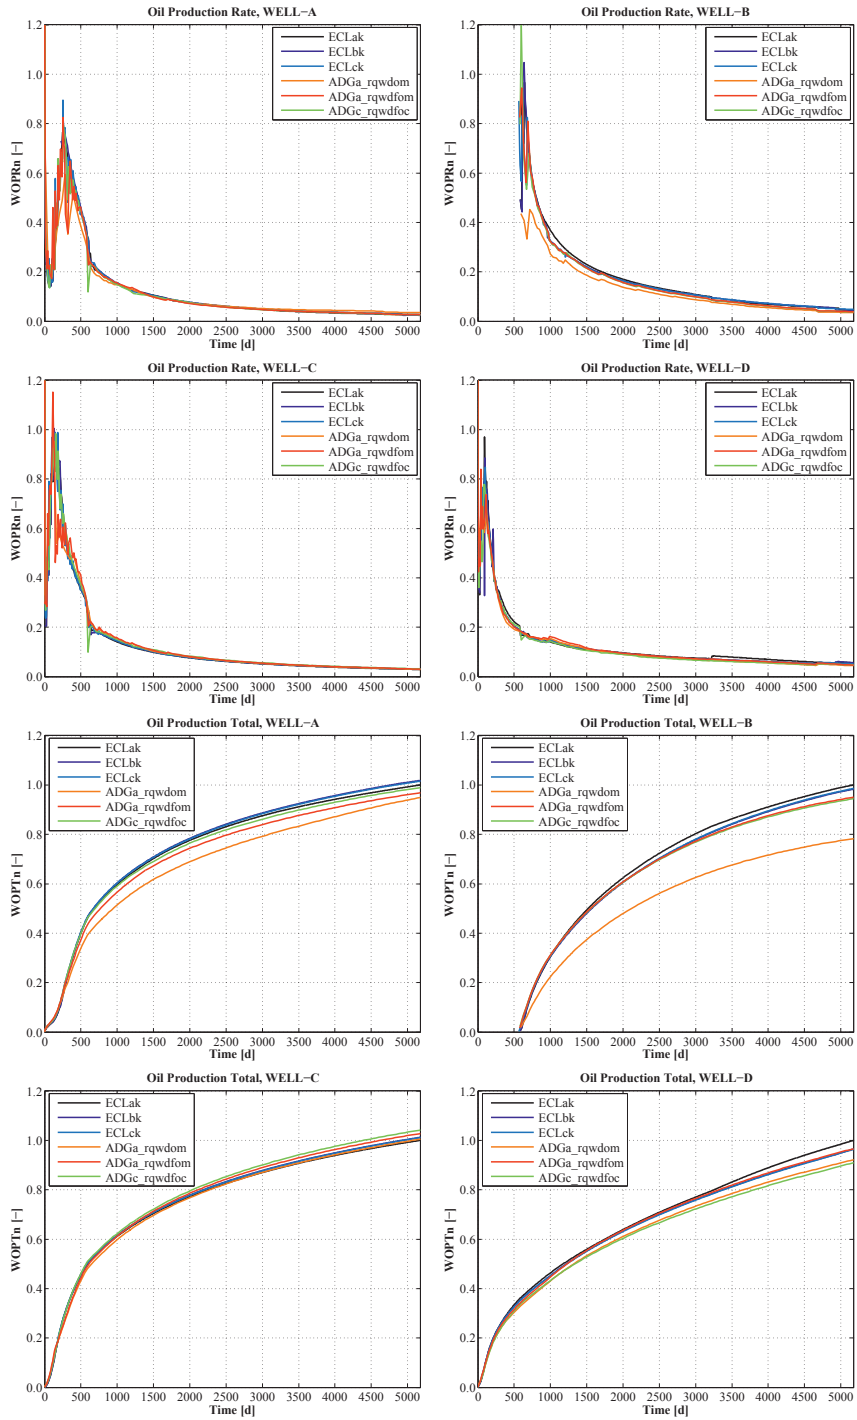


Figure 6.2: Well oil production rates and totals. Suffix “n” on production mnemonics signifies values have been normalized for confidentiality reasons.

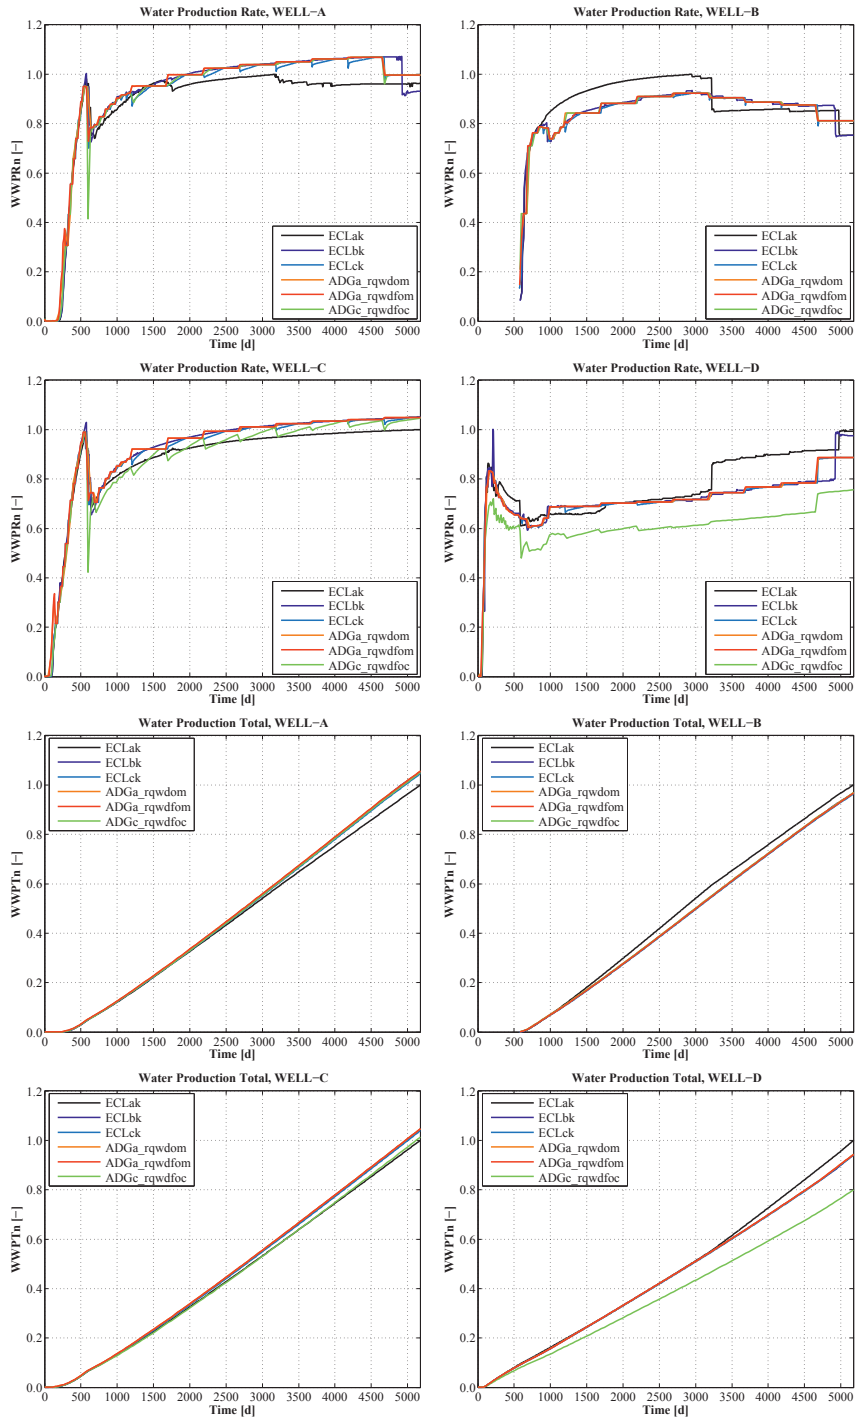


Figure 6.3: Well water production rates and totals. Suffix “n” on production mnemonics signifies values have been normalized for confidentiality reasons.

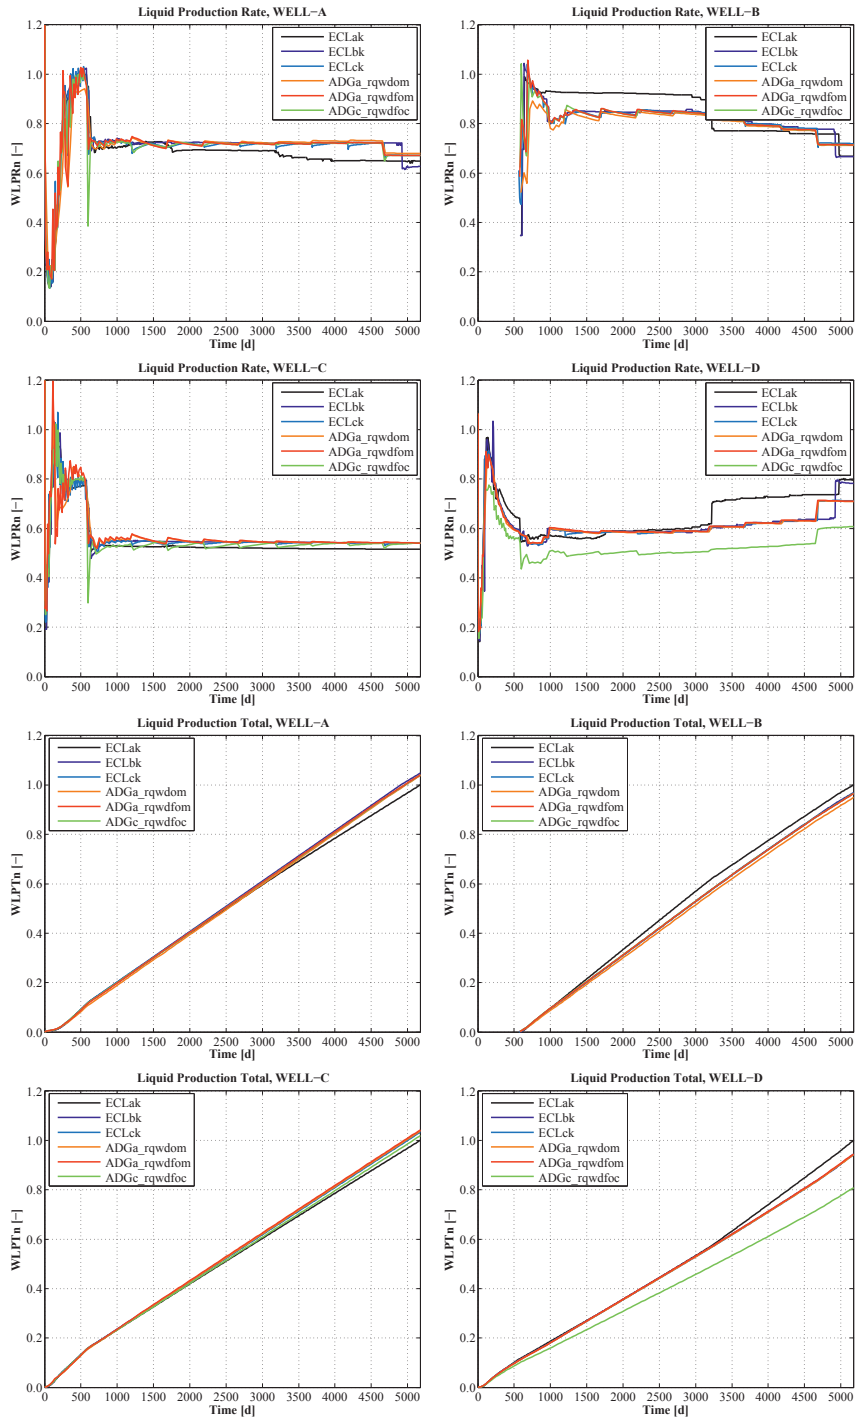


Figure 6.4: Well liquid production rates and totals. Suffix “n” on production mnemonics signifies values have been normalized for confidentiality reasons.

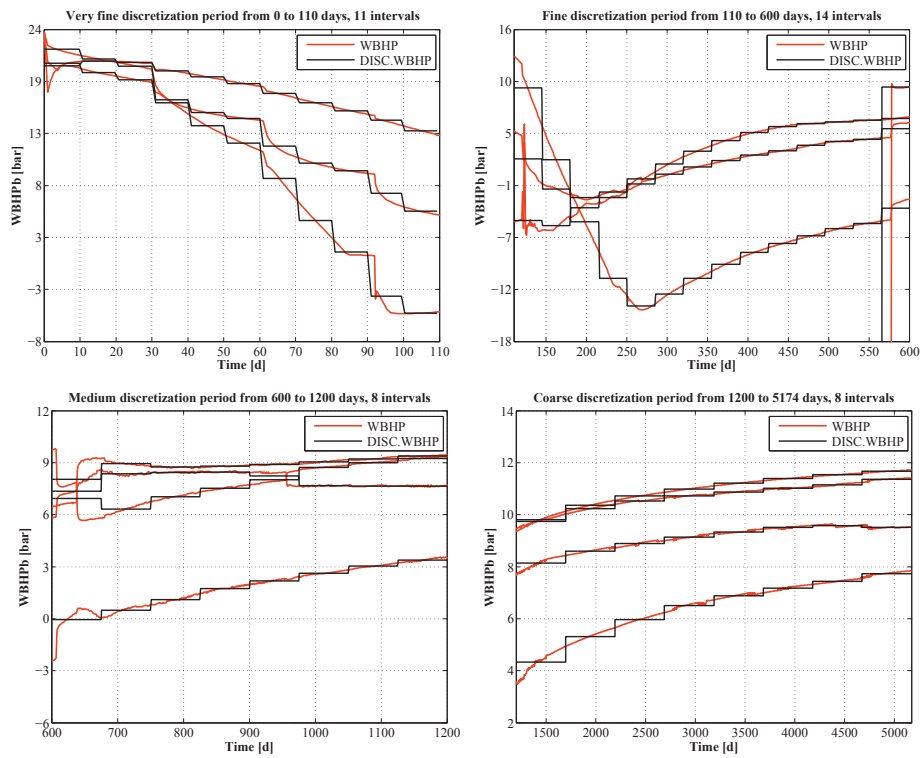


Figure 6.5: Discretized well bottom-hole pressures for the four production periods comprising the entire production time frame. Suffix “b” on pressure mnemonics signifies values have been scaled for confidentiality reasons.

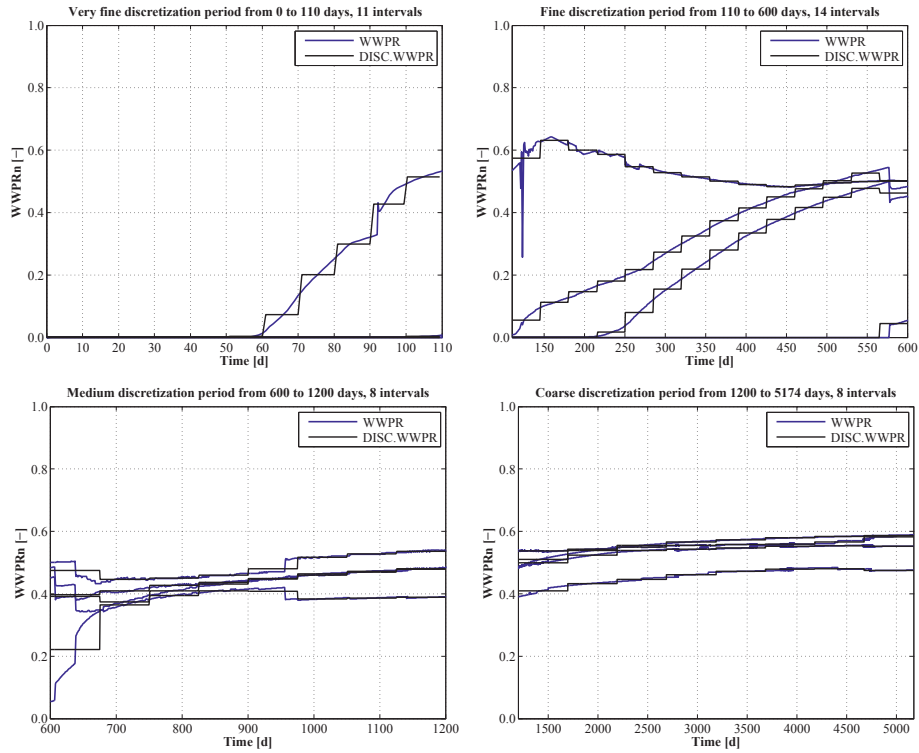


Figure 6.6: Discretized water production rates for the four production periods. Suffix “n” on production mnemonics signifies values have been normalized for confidentiality reasons.

ORIGINAL ECLIPSE WELL CONTROLS (EXAMPLE)										
WCONPROD										
NAME	STAT	CTRL	ORAT	WRAT	GRAT	LRAT	RESV	BHP	THP	
WL-X	OPEN	GRUP	1*	1*	1.5E6	5500	1*	70	30	
GCONPROD										
NAME	CTRL	ORAT	WRAT	GRAT	LRAT	OPRO				
OILPROD	ORAT	5500	1*	1.0E6	5500	RATE	1*	1*	1*	RATE RATE /
GASPROD	GRAT	10600	1*	2.0E6	10600	RATE	1*	1*	1*	RATE RATE /
2GROUPS	ORAT	10600	1*	2.1E6	10600	RATE	1*	1*	1*	RATE RATE /

Figure 6.7: Well and group controls and constraints for the two Eclipse simulation cases 1 and 2 using the original control setup, ECLak_orig and ECLbk_ovaq, respectively. Well control type and pressure constraint marked in bold.

MODIFIED ECLIPSE / AD-GPRS WELL CONTROLS									
WCONPROD									
NAME	STAT	CTRL	ORAT	WRAT	GRAT	LRAT	RESV	BHP	THP
WL-X	OPEN	BHP	1*	XXX	1*	1*	1*	XXX	1* /

Figure 6.8: Discretized well controls used in approximated cases 3 to 6, ECLck_rvqw, ADGa_rqwdo, ADGa_rqwdf, and ADGc_rqwdoc, respectively. Well control type and production/pressure constraints marked in bold.

Appendix B

Here we present two pieces of pseudo-code that detail the two general steps of the Joint+RCO approach (see Section 3.2.1). Different versions of this enhanced approach have been implemented for a case with two injectors and three producers, all vertical on a simple reservoir model. Final results are forthcoming, and further work will introduce the enhancement in future implementations of the optimization framework.

Computation of projection matrix Φ . The procedure for computing Φ is described in Algorithm 4. The POD procedure is applied separately for controls corresponding to injectors and producers, analogous to how pressure and saturation states are treated in Cardoso and Durlofsky (2010). Also Doren et al. (2006) segregate the pressure and saturation states “because they correspond to different physical processes and will consequently generate different dominant structures.” Here we superimpose this line of reasoning to control vector elements corresponding to injectors and producers, respectively. Therefore, in the training step of our implementation, we build separate snapshot and projection matrices for injector and production data.

Implementation case. Joint+RCO was implemented on a case similar to the one in Chapter 1. Similarly, the locations of the wells in this case are represented by areal (x, y) coordinates. The case contains a set of 5 wells (2 injectors, 3 producers) which yields a well position vector $\mathbf{x} \in \mathbb{R}^n$ with $n = 5 \times 2 = 10$. A $2n$ coordinate search around position vector \mathbf{x} yields a poll set of $2n = 20$ well position vectors, i.e., $P = \{\mathbf{x}^m | m = 1, \dots, 20\}$. Each \mathbf{x} is associated with a well control vector \mathbf{u} describing the time-dependent well controls for all wells in \mathbf{x} . For each well, the controls are defined by piecewise constant functions over $N_t = 20$ intervals. For this 5-well case, \mathbf{u} is a $5 \times N_t = 5 \times 20 = 100$ element vector. Snapshot data matrices containing injector and producer control elements are thus given as:

$$\mathbf{U}_{injs} = \{\mathbf{u}_{injs}^m | m = 1, \dots, 20\} \quad \text{where} \quad \mathbf{u}_{injs} = \begin{bmatrix} \mathbf{u}_{inj1} \\ \mathbf{u}_{inj2} \end{bmatrix}$$

$$\mathbf{U}_{prods} = \{\mathbf{u}_{prods}^m | m = 1, \dots, 20\} \quad \text{where} \quad \mathbf{u}_{prods} = \begin{bmatrix} \mathbf{u}_{prod1} \\ \mathbf{u}_{prod2} \\ \mathbf{u}_{prod3} \end{bmatrix}$$

For this case, optimal control injector and producer data matrices, \mathbf{U}_{injs} and \mathbf{U}_{prods} , have dimensions 20×40 and 20×60 , respectively.

Determining number of columns of Φ_{injs} and Φ_{prods} . We use a definition of fraction of total energy E_t (Cardoso and Durlofsky, 2010; He et al., 2011) to select the number of columns l in projection matrices Φ_{injs} and Φ_{prods} , i.e., l_{injs} and l_{prods} , respectively. The

Algorithm 4 : Polling procedure used for training – Full-order control optimization performed for all well placements vectors \mathbf{x} in poll set P_i . Projection matrices Φ_{injs} and Φ_{prods} are subsequently computed.

- 1: **for** all $\mathbf{x} \in P_i$ **do**
 - 2: Set initial controls \mathbf{u}_{init}
 - 3: Set lower and upper bounds for \mathbf{u} , i.e., \mathbf{u}_l and \mathbf{u}_u
 - 4: Run full-order control optimization with regular (relatively high) number of major iterations in SNOPT, e.g., 16,
 - 5: Save control solution \mathbf{u}^*
 - 6: Return $f^* = f(\mathbf{x}, \mathbf{u}^*)$
 - 7: **end for**
 - 8: Assemble data matrices \mathbf{U}_{injs} and \mathbf{U}_{prods}
 - 9: **for** \mathbf{U}_{injs} and \mathbf{U}_{prods} **do**
 - 10: Perform SVD(\mathbf{U})
 - 11: $\lambda_i = \sigma_i^2$
 - 12: Determine l
 - 13: Construct Φ
 - 14: **end for**
 - 15: Use Φ_{injs} and Φ_{prods} in subsequent projections
-

energy fraction is given as:

$$\frac{E_l}{E_t} = \frac{\sum_{i=1}^l \lambda_i}{\sum_{i=1}^k \lambda_i} = \frac{\sum_{i=1}^l \sigma_i^2}{\sum_{i=1}^k \lambda_i}$$

where k is the total number of columns in the respective projection matrix. λ_i and σ_i are corresponding eigenvalues and singular values, respectively. In this implementation we select a cut-off value of $\frac{E_l}{E_t} = .99$ to determine the sizes of both projection matrix Φ_{injs} and Φ_{prods} . We select the columns corresponding to the snapshot matrix eigenvalues that sum up to make this cut-off value. This value can be reduced to further decrease the number of optimization variables in reduced space, i.e., the dimension of \mathbf{z} . Also, independent cut-off values may be chosen for the construction of Φ_{injs} and Φ_{prods} .

Reduced-order control optimization The procedure for optimizing for controls in a low-dimensional space is described in Algorithm 5. In our current implementation, the reduced-order variable \mathbf{z} is unbounded during the control optimization in low-dimensional space. Instead, we enforce the lower and upper bounds from the full-order formulation (\mathbf{u}_l and \mathbf{u}_u , respectively) by applying projections of these bounds as linear constraints within the reduced-order optimization problem. Other configurations to ensure feasibility of reduced-order solution once transferred to high-dimensional space are being explored.

Algorithm 5 : Surrogate polling procedure – Poll set that uses reduced–order control optimization. This type of polling is performed until a new training polling procedure is called for (Algorithm 4).

- 1: **for** all $\mathbf{x} \in P_{i+1}$ **do**
 - 2: Set initial controls $\mathbf{z}_{init} = \Phi^T \mathbf{u}_{init}$
 - 3: Introduce linear inequality constraints $\Phi \mathbf{z} \geq \mathbf{u}_l$ and $\Phi \mathbf{z} \leq \mathbf{u}_u$ into reduced–order control problem formulation
 - 4: Run reduced–order control optimization using lower (minimal) setting of major iterations in SNOPT
 - 5: Return $f^* = f(\mathbf{x}, \mathbf{z}^*)$
 - 6: **end for**
-

