MULTILEVEL FIELD DEVELOPMENT OPTIMIZATION UNDER UNCERTAINTY USING A SEQUENCE OF UPSCALED MODELS

A DISSERTATION
SUBMITTED TO THE DEPARTMENT OF
ENERGY RESOURCES ENGINEERING
AND THE COMMITTEE ON GRADUATE STUDIES
OF STANFORD UNIVERSITY
IN PARTIAL FULFILLMENT OF THE REQUIREMENTS
FOR THE DEGREE OF
DOCTOR OF PHILOSOPHY

Elnur Aliyev
May 2015

I certify that I have read this dissertation and that, in my opinion, it is fully adequate in scope and quality as a dissertation for the degree of Doctor of Philosophy.

(Louis J. Durlofsky) Principal Adviser

I certify that I have read this dissertation and that, in my opinion, it is fully adequate in scope and quality as a dissertation for the degree of Doctor of Philosophy.

(Roland Horne)

I certify that I have read this dissertation and that, in my opinion, it is fully adequate in scope and quality as a dissertation for the degree of Doctor of Philosophy.

(Oleg Volkov)

Approved for the University Committee on Graduate Studies

Abstract

Determining the well locations and settings that maximize reservoir performance is a key issue in reservoir management. Computational optimization procedures are commonly applied for this purpose. In conventional optimization methods, flow simulation is performed over the basic reservoir model. If optimization under uncertainty is considered, multiple geological realizations, which quantify reservoir uncertainty, are typically used. This can result in great computational expense, as commonly used optimization methods may require thousands of function evaluations (each of which entails multiple flow simulations). Efficient surrogate models, which can be used to reduce the computational requirements, would thus be highly beneficial.

In this thesis, a multilevel optimization procedure, in which optimization is performed over a sequence of upscaled models, is developed for use in combined well placement and control problems. The multilevel framework is applicable for use with any type of optimization algorithm. In this work it is implemented within the context of a particle swarm optimization – mesh adaptive direct search (PSO-MADS) hybrid technique. An accurate global transmissibility upscaling procedure is applied to generate the coarse-model parameters required at each grid level. Distinct upscaled models are constructed using this approach for each candidate solution proposed by the optimizer at each grid level. We demonstrate that the coarse models are able to capture the basic ranking of the candidate well location and control scenarios, in terms of objective function value, relative to the ranking that would be computed using fine-scale simulations. This enables the optimization algorithm to appropriately select and discard candidate solutions.

The multilevel optimization framework is further extended for use in optimization

under geological uncertainty. Toward this goal, we introduce an accelerated multilevel optimization procedure, in which the full PSO-MADS algorithm is used only at the coarsest grid level. At subsequent grid levels, standalone MADS (a local optimization algorithm) is applied. An optimization with sample validation (OSV) procedure is also incorporated into the multilevel method. This approach enables us to optimize over a limited set of geological realizations. The combined use of accelerated multilevel optimization and OSV leads to substantial speedup compared to direct application of the standard multilevel optimization procedure.

Optimization results for two- and three-dimensional example cases, which involve both single and multiple geological realizations, are presented. The multilevel procedure for single realizations is shown to provide optimal solutions that are comparable, and in some cases better, than those from the conventional (single-level) approach. Computational speedups of about a factor of five to ten are achieved. For optimization under uncertainty, we use the accelerated multilevel procedure with sample validation to further reduce the computational cost of the optimization. Speedups of a factor of 10–20 are achieved by the accelerated multilevel approach relative to the conventional procedure for examples with ten realizations. An additional speedup factor of about two is observed through incorporation of OSV for cases involving 100 realizations. Our overall findings thus suggest that this framework may be quite useful for practical field development. We also investigate the application of the multilevel Monte Carlo approach for field development optimization under geological uncertainty as an alternative to the multilevel optimization technique. Although this approach provides speedup relative to the conventional (single-level) treatment, it is not as efficient as the multilevel procedure developed in this work.

Acknowledgments

I am thankful to my advisor, Prof. Louis J. Durlofsky, for the tremendous effort he invested in this dissertation and in me to become a researcher. From the first day on, he clearly defined the requirements and explained the level of commitment for a Ph.D. degree. He guided me through tough times, and never doubted my abilities. He was with me at every phase of the learning process, patiently watching over my progress. He set a great example for me as a researcher, teacher and advisor. I consider myself lucky to be one of his students.

I would also like to thank my defense committee, Prof. Roland Horne, Prof. Khalid Aziz, Prof. Yinyu Ye and Dr. Oleg Volkov, for their insightful comments on this dissertation. I am additionally grateful to the members of the Smart Fields and SUPRI-B research groups, especially Prof. Tapan Mukerji for his help during all the seminars and courses. I would also like to thank Obi Isebor (who provided the PSO–MADS optimizer used in this work), Mehrdad Shirangi (who developed the OSV procedure used here), Matthieu Rousset (who provided code for global upscaling in two dimensions), Rustem Zaydullin, Hai Vo and Denis Voskov for their help on my research. It has been a great pleasure to have them in my group and to take advantage of their ideas and suggestions. I have become closer to all my friends in the Smart Fields and SUPRI-B groups throughout all these years. Also, I am very thankful to the industrial affiliates of the Smart Fields Consortium at Stanford University for financial support during my studies.

I also would like to thank the BP Reservoir Management Team in Houston, Texas for giving me the opportunity to work with them during four summers. I want to thank Mike Litvak for being the most valuable mentor anybody could have. Also, I

would like to thank Bob Gochnour, Jerome Onwunalu, Jeff Parke, Patrick Angert, Jamie DeMahy, Scott Lane, Gherson Penuela, David Sparling and Karl Kramer. I am grateful for the meetings, their help, the fruitful discussions, and the tremendous insight they have provided me during my internships.

I would like to thank all my friends and colleagues in the ERE department. They have made life at Stanford a memorable experience. Especially, I would like to mention Ramil Ahmadov, Abdulla Kerimov, Selcuk Fidan, Matthieu Rousset, Mohammad Shahvali, Orhun Aydin, Markus Buchgraber, Asli Gundogar, Alex Galley, Vladislav Bukshtynov, Timur Garipov, Hangyu Li, Kyu Koh Yoo, Boxiao Li, Ognjen Grujic, Maytham Al Ismail, Khalid Alnoaimi and Ernar Sagatov for always being there for me. They helped me through both the good and bad times. I am blessed to have such good friends. Moreover, I want to thank my close friends outside Stanford, Nazrin, Orkhan, Afaq, Yasmin, Sabrina, Jafar, Murad, Tural, Vusal, Fuad, Akram, Toghrul, Jafar, Fariz, Elvin, Garib, Ramil, Ismail, Sultana and many others for giving me support and confidence all these years, and the countless fun moments. It has been a privilege to have this crowd as my life-long friends. I would also like to thank my undergraduate professors Serhat Akin, Ender Okandan, Mahmut Parlaktuna and Evren Ozbayoglu.

Finally, I would like to thank my parents Nuraddin Aliyev and Tamella Abilova and siblings Ziya Aliyev and Yegana Aliyeva for their moral support. They have always encouraged me to become a more successful person. My parents were the most influential people in my life. Without their support, I would have not achieved this. I would like to dedicate this thesis to them.

Contents

List of Tables

List of Figures

Chapter 1

Introduction

The determination of the well locations and controls that maximize a particular objective is of primary importance in oil field development and operation. The reservoir performance associated with a particular set of well locations and controls is evaluated using flow simulation. Many optimization procedures have been developed for this problem, and they typically require large numbers of simulations. Conventional approaches entail performing these simulations using the actual reservoir model (or models), which may be a high-resolution description.

Multiple geological realizations are typically used to characterize reservoir uncertainty. In optimization under uncertainty, flow simulation over multiple geological models must be performed for each function evaluation required during optimization. This leads to very large computational demands. In this case the objective function is evaluated by averaging over the individual model flow responses.

Different approaches can be used to reduce the computational requirements associated with optimization. The use of efficient but accurate surrogate models could lead to a substantial reduction in computational requirements. In addition, using an appropriate set of 'representative' realizations selected from the full set of geological models can further reduce the computational requirements.

In this work, we develop an optimization framework that greatly reduces computational requirements for field development optimization. The main component of this

framework is a multilevel optimization procedure that uses a sequence of upscaled models. For optimization under uncertainty, we also incorporate model selection and validation procedures into the multilevel optimization framework. This reduces the number of realizations that must be simulated during the optimization.

1.1 Literature review

Extensive literature exists on the general topics of field development optimization and well control optimization. Here, we first discuss derivative-free optimization algorithms, which are the types of algorithms used in this work. Then, joint field development and production optimization is described. Next, optimization under geological uncertainty and optimization with sample validation are discussed. Studies involving the use of proxies in optimization, and multigrid-based optimization, are then reviewed. Finally, upscaling procedures are discussed.

1.1.1 Optimization procedures for oil field problems

Oil field optimization problems can be categorized as well placement problems, well control problems, or joint optimization problems involving both well placement and control. In well placement problems, the locations of wells are optimized, while in well control problems, well settings such as flow rates and/or bottom-hole pressures (BHPs) are optimized. Joint optimization problems involve the simultaneous optimization of both sets of variables. Derivative-free optimization methods do not require gradients, and they can be used for all of these problem types. They are usually, however, less efficient than adjoint-gradient methods, though adjoint methods are invasive and require access to simulator source code. In this section we present studies that applied derivative-free methods for oil field optimization.

The most common approach used for well placement optimization is probably the genetic algorithm (GA) [26, 51], which is a stochastic evolutionary procedure. Guyaguler et al. [25] used GA to optimize well placement. Yeten et al. [61] optimized type, location and trajectory of nonconventional wells using GA. Litvak and Angert [39] applied GA to optimize field development in giant oil fields. Recently, Bouzarkouna et al.

[9] applied another procedure called covariance matrix adaptation evolution strategy (CMA-ES) for well placement problems. They reported that CMA-ES outperformed GA on the well placement optimizations considered. Onwunalu and Durlofsky [43] applied a different global optimization algorithm, particle swarm optimization (PSO), which is based on the social interactions of animal groups, for well location problems. They found that PSO provided better results than GA for a variety of well location optimization problems. It should be noted, however, that there are many variants of both GA and PSO, and certain variants may perform better for particular problems.

Global search algorithms can be hybridized with local derivative-free search methods to improve overall efficiency and performance. Yeten et al. [61] combined a local hill-climber algorithm with GA and demonstrated that the hybrid algorithm outperformed standalone GA. Guyaguler and Horne [24] hybridized GA with a polytope method and applied it to well placement problems. Recently, Isebor et al. [29, 31, 32] presented a hybrid algorithm that is a combination of (global) PSO [20] and (local) mesh adaptive direct search, or MADS [5], and demonstrated that the hybrid procedure outperformed the standalone PSO and MADS algorithms. In our optimizations here, we will utilize this PSO–MADS algorithm.

Gradient-based approaches are commonly applied for well control optimization problems. Wang et al. [57] compared various optimization algorithms including steepest ascent and simultaneous perturbation stochastic approximation (SPSA) for production optimization in a closed-loop reservoir management framework. They showed that the steepest ascent method is the most efficient among all of the algorithms considered. Sarma et al. [47] and Brouwer and Jansen [10] applied adjoint-gradient-based optimization to waterflooding problems. The adjoint procedure is more efficient because it uses gradients that are constructed efficiently from the underlying simulator. Echeverría et al. [23] compared several derivative-free optimization methods, including Hooke-Jeeves, general pattern search (GPS), and GA to gradient-based algorithms for well control optimization problems. They concluded that gradient-based sequential quadratic programming, GPS and a hybrid method combining GA with an efficient local search method, were the most effective.

The optimization of well placement and well controls can be addressed in a sequential or in a joint (simultaneous) manner. In the sequential approach, well locations are optimized with a specified set of well controls or a particular well control strategy. Then, well controls are optimized in a separate optimization problem (with well locations fixed). In the joint approach, well locations and controls are considered together. The multilevel optimization procedure developed in this thesis can be used with both sequential and joint approaches. Because the problems addressed in this work involve joint optimization, we now discuss recent work in this area.

Bellout et al. [8] implemented a nested joint optimization approach involving direct search (for well placement) combined with adjoint-gradient-based optimization (for well controls). They demonstrated that this approach provided improved objective function values compared to a sequential procedure. Similar findings were reported by Li and Jafarpour [38], who developed a method in which they alternated between the two optimization problems. They used the well-distance constrained SPSA algorithm for well placement and a gradient-based optimization for well control. Humphries et al. [28] used a hybrid optimization algorithm, which is a combination of a stochastic procedure (PSO) and a direct search (GPS), for joint optimization problems. They came to somewhat different conclusions, as they did not observe consistently better solutions using a joint optimization approach. This may be due to the specific algorithmic treatments employed in their work, including the use of heuristics for well control during well placement optimization.

Isebor et al. [31, 32] proposed a PSO–MADS procedure that optimizes well locations and controls simultaneously. Using this approach, joint optimization was found to consistently outperform sequential optimization for combined field development and well control problems [31]. As noted above, the method used in this work is this PSO–MADS procedure.

1.1.2 Optimization under geological uncertainty

Reservoir models that are used in field development optimization contain many uncertain parameters due to uncertainty in the subsurface reservoir geology (reservoir structure, faults, permeability, porosity), fluid properties, etc. Multiple geological

realizations are created to characterize subsurface uncertainty. For a particular well configuration in optimization, flow simulation must then be performed over multiple geological realizations, which leads to substantial computational expense.

A number of researchers presented studies on well control optimization under uncertainty. Aitokhuehi and Durlofsky [2], van Essen [54], Sarma et al. [48], Su and Oliver [52] and Wang et al. [57] optimized well controls with multiple geological models. They used gradient-based algorithms in their optimizations.

Several researchers investigated well placement optimization under geological uncertainty. Guyaguler and Horne [24] applied decision tree tools and a utility theory framework to solve well placement optimization under geological uncertainty to maximize expected NPV. Cameron and Durlofsky [11] considered well placement and control optimization under geological uncertainty for carbon storage problems. They generated multiple geological realizations to represent uncertainty in aquifer models and then used optimization to minimize the risk of leakage.

Optimization under geological uncertainty is computationally expensive. Although uncertainty may be better represented with a large number of realizations, computational requirements increase as the number of realizations increases. There have been some studies on improving the efficiency of optimization under uncertainty. Artus et al. [4] optimized monobore and dual-lateral well locations under geological uncertainty. They applied a cluster-analysis-based proxy model to estimate the cumulative distribution function (CDF) of the objective function. Artus et al. [4] reported that with this approach, by using about 10% or 20% of the realizations, they were able to achieve comparable results to those obtained using all realizations. Wang et al. [58] applied a retrospective optimization (RO) framework that solves a sequence of optimization subproblems using an increasing number of realizations. They showed that, by using the RO procedure with cluster-based sampling, the computational expense of well placement under uncertainty was reduced by about an order of magnitude compared with optimizing using the full set of realizations at all iterations.

Shirangi and Durlofsky [50] introduced an alternative approach to reduce the computational demands of optimization under geological uncertainty. They presented a

systematic optimization with sample validation (OSV) procedure to reduce the number of geological models used in optimization. This procedure shares some similarities with the retrospective optimization (RO) procedure presented by Wang et al. [58]. As in RO, OSV divides the optimization problem into subproblems with increasing numbers of realizations. However, in OSV, the performance of the optimization in a subproblem is assessed with a specific validation criterion. The objective function value is calculated at the beginning and end of the optimization using all realizations. Then, the relative improvement (RI), which is the ratio of objective function value improvement with all realizations to the improvement with only the representative realizations, is calculated. If RI does not meet the defined criterion, then the optimization is repeated with a larger number of representative realizations.

There are some differences in the model selection strategies used in the OSV and RO approaches. Wang et al. [58] used random and cluster sampling in RO to select the representative models to be used in optimization. They constructed clusters based on cumulative oil production, original oil in place (OOIP) and the location of the water-oil contact. Shirangi and Durlofsky [50], by contrast, used the CDF of the objective function values to select the representative geological models. To generate this CDF, all realizations are run using the current best well configuration. Models are then selected such that the representative set of models provides a CDF in essential agreement with the CDF based on the full set of models. In this study, we incorporate the OSV procedure into our multilevel optimization framework.

Another approach that can be used to reduce the computational cost of optimization under uncertainty is the multilevel Monte Carlo (MLMC) method. The MLMC method enables the efficient computation of the expected value of a reservoir simulation output (e.g., the objective function) over multiple realizations. MLMC is based on a so-called telescopic sum of different numbers of realizations at different coarsening levels. Müller et al. [40] combined MLMC with streamline simulation to assess uncertainty in problems involving two-phase flow in random heterogeneous porous media. They showed that results using MLMC led to an order of magnitude speedup and were similar to results using all realizations. In this work, we will use MLMC in field development optimization and compare it with optimization using a sequence of

upscaled models.

1.1.3 Treatment of nonlinear constraints

Bound, linear and/or nonlinear constraints are typically required in optimization problems. Bound constraints appear when there is a specified range for the optimization variables. Well rate constraints (with BHPs as the control variables) and water cut constraints are examples of nonlinear constraints. Including these constraints in field development optimization often renders the optimization more difficult.

Isebor et al. [32] applied a filter method to handle nonlinear constraints. Use of a filter method is somewhat similar to performing biobjective optimization, where the first objective is minimizing or maximizing the objective function, and the second objective is minimizing the aggregate constraint violation. Feasibility is typically achieved after some number of optimization iterations.

Isebor et al. [32] applied the filter method for field development optimization with a single realization. However, this approach may not be the most appropriate for optimization under uncertainty. In this case, a feasible solution that honors all constraints in all realizations might not even exist. Even if such a solution exists, it may be overly conservative. For this reason, we do not use filter methods in this work.

Penalty methods are also widely used in constrained optimization problems. Guyaguler et al. [25] applied a penalty method to treat constraints in the optimization of well locations and water pumping rates in a Gulf of Mexico field. Echeverría et al. [23] used a penalty method for production optimization with constraints. Both of these studies considered problems with only a single realization.

Penalty methods have been used in mining applications to treat constraints in optimization with multiple realizations. Dimitrakopoulos [17] formulated an objective function, using a stochastic programming formulation with a penalty, to optimize mining operations under geological uncertainty. This formulation penalizes realizations that violate the constraints. In this work, this type of formulation is used to handle nonlinear constraints in field development optimization under uncertainty.

1.1.4 Proxy-based optimization

The general problem of optimizing with expensive function evaluations has also received significant attention in other application areas. Shan and Wang [49] provide a survey of strategies for addressing high-dimensional optimization problems with computationally-expensive black-box functions. Bandler et al. [7] reviewed the application of surrogate models in engineering design optimization. They showed that, through model refinement, the efficiency and robustness of the optimization scheme can be significantly improved. Echeverría [21] and Echeverría and Hemker [22] used surrogates with multiple levels of accuracy to solve optimization problems. They iteratively corrected an existing surrogate to improve the quality of the surrogate model during the optimization process. This correction is local in nature. They applied their surrogate-based optimization algorithm for optimization problems in the fields of magnetics, electronics and photonics, and showed that computation time can be significantly reduced.

Surrogate or proxy models have been used for a variety of oil field optimization problems. Yeten et al. [61] and Guyaguler and Horne [24] used artificial neural networks and kriging as statistical proxies in well placement optimization. Doren et al. [18], Cardoso and Durlofsky [13] and He and Durlofsky [27] applied reduced-order modeling procedures based on proper orthogonal decomposition (the latter two studies also incorporated trajectory piecewise linearization) for well control optimization. Reduced-physics models have also been used for optimization. Examples include the use of streamline procedures for well control optimization in waterfloods [59] and the use of simplified simulation models for optimizing horizontal wells and hydraulic fractures in shale gas production [60].

Upscaled models have additionally been applied within optimization frameworks. Abukhamsin [1] performed well placement optimization using coarse-scale models and found that the optimal locations differed from those using fine models. Krogstad et al. [34] recently applied global transmissibility upscaling, as is used here, to generate coarse models for gradient-based well control optimization. A high level of solution accuracy and clear computational benefits were reported. Neither of these studies, however, used a sequence of upscaled models, nor did they consider the joint

optimization of well locations and controls, as is accomplished here.

Lewis and Nash [36] and Nash and Lewis [41] used a multigrid optimization method to accelerate nonlinear programming algorithms. According to Lewis and Nash [36], although multigrid optimization is computationally efficient for problems governed by elliptic partial differential equations (PDEs), it can also be applied for systems described by other types of equations. Multigrid optimization entails the use of a sequence of grids, in common with our approach. These authors showed that results close to those from traditional optimization approaches can be obtained more efficiently with multigrid methods.

1.1.5 Upscaling techniques

Reservoir models with a large number of grid blocks are computationally expensive to run. If such models are used for optimization, which may require many thousands of simulation runs, elapsed times may be excessive. A variety of upscaling methods can be used to reduce the computation time required for simulation. Durlofsky and Chen [19] describe many of the existing upscaling methods in detail.

As discussed by Durlofsky and Chen [19], upscaling methods can be classified in terms of the types of properties that are upscaled. In single-phase upscaling, permeability or transmissibility is upscaled, and in two-phase upscaling, relative permeability is additionally upscaled. Although the upscaling of both single-phase and two-phase properties provides better accuracy than the use of single-phase upscaling alone, two-phase upscaling requires substantial additional computation. Also, single-phase upscaling is in some cases more robust. For these reasons, in this study we use only single-phase upscaling techniques.

Single-phase upscaling methods can be further classified in terms of the region over which the coarse-scale properties are computed. These methods can be divided into four categories: local, extended local, local-global and global upscaling, based on the region used in the computations [19, 37]. Local upscaling methods use only the fine-scale cell information for the target coarse block. Extended local methods

include neighboring fine-scale cells in the upscaled property computation. Local-global methods use global coarse-scale simulations to estimate boundary conditions for local and extended local upscaling computations. Global upscaling methods are generally considered to provide the most accurate coarse models among all these upscaling techniques, though they are typically more expensive. In this study we use a global upscaling method since the time required for upscaling is relatively small compared to the other computations performed in the optimization.

The coarse-scale property computed in single-phase upscaling can be either absolute permeability or interface transmissibility. Coarse-scale transmissibility includes both grid-block permeability and geometry effects. In fact, transmissibility upscaling often provides more accurate coarse results than permeability upscaling. This has been shown in [46], for a local upscaling method, and in [15] for a local-global upscaling procedure.

Based on the discussion above, in this work we apply a global transmissibility upscaling procedure. Specifically, we use a method of the type described by Chen et al. [16] and Zhang et al. [62]. In [16], boundary conditions were specified along portions of the reservoir boundary, while in [62], flow was driven by wells. In our work, as in [37], transmissibility and well index values for the coarse model are calculated from the global velocity and pressure fields, for a specific set of well locations and controls.

1.2 Scope of work

Finding the optimum well locations and controls is a key problem in reservoir management. Field development optimization procedures are, by nature, computationally expensive. Incorporating geological uncertainty into the optimization substantially increases computational requirements. The main focus of this dissertation is to improve the efficiency of field development optimization algorithms, for both single and multiple realization optimization problems.

Toward this goal, we introduce a multilevel optimization procedure that uses a sequence of upscaled models. We extend the framework to handle optimization under geological uncertainty. A stochastic programming formulation with an appropriate

penalty function is implemented to handle nonlinear constraints in optimization. We improve the efficiency of the procedure by only using a local search method after some number of optimization iterations. Finally, we incorporate the sample validation (OSV) procedure to further reduce the computational requirements for optimization under uncertainty.

The main objectives of this research are:

- to develop a new multilevel optimization framework. We introduce an efficient optimization approach that allows us to replace most of the expensive fine-scale simulation runs with less expensive runs that use upscaled models. Any optimization procedure could be used as the core optimizer.

- to extend the multilevel optimization framework to handle optimization under geological uncertainty with constraints. The objective function includes a penalty term, which is nonzero for realizations that violate nonlinear constraints. As the optimization proceeds and simulation runs become more expensive, we introduce a treatment where we replace PSO–MADS with MADS, a local optimization algorithm. This approach, called the accelerated multilevel procedure, further improves the efficiency of the multilevel optimization framework.

- to incorporate the OSV procedure into the multilevel optimization framework. The OSV method reduces the number of realizations used in the optimization process and thus acts to further reduce computational requirements.

- to apply MLMC to field development optimization under uncertainty. This approach can then be compared to our multilevel optimization procedure.

1.3 Dissertation outline

In this dissertation, we introduce a new and efficient procedure for the joint optimization of well location and control. In Chapter 2 we pose the field development and

well control problem as a mixed-integer nonlinear programming (MINLP) problem, as described by Isebor [29]. The formulation is extended to treat optimization under geological uncertainty. Then, we briefly describe the PSO–MADS hybrid procedure used in this work. We also discuss our penalty-based approach for treating constraint violations in multiple-realization problems. The iterative global upscaling method and multilevel optimization procedure are then described. The incorporation of a sample validation procedure into the multilevel optimization framework is discussed. Finally, we describe the use of a MLMC method for optimization under geological uncertainty. We note that part of the work presented in Chapter 2 has appeared in [3].

In Chapter 3 we apply the multilevel optimization procedure to field development and well control problems. We consider two- and three-dimensional examples involving a single geological realization. Multilevel optimization results are compared to results from the conventional approach where only the fine-scale model is used. Most of the examples in Chapter 3 have been presented in [3].

In Chapter 4 we extend the multilevel optimization procedure to field development problems characterized by multiple realizations. Optimization is performed based on the expected value over these multiple realizations to find a robust solution. The performance of the multilevel optimization procedure is further improved by switching from PSO–MADS to MADS as the optimization proceeds. The sample validation procedure is applied to reduce the number of models used in the optimization. Results and timings are compared to those achieved using simpler procedures. We also apply MLMC for one of the examples to enable a comparison with our multilevel optimization procedure.

We conclude this dissertation with a summary and suggestions for future work in Chapter 5.

Chapter 2

Multilevel Optimization Procedure

In this chapter, we introduce the multilevel optimization procedure used to solve well location and control optimization problems. We first describe the formulation for jointly optimizing field development and well control, including the treatment of geological uncertainty. Next, the PSO-MADS hybrid optimization procedure is presented. A penalty-based method to handle nonlinear constraints in optimization under uncertainty is described. The iterative global upscaling procedure used in this work is then presented. Next, we introduce the multilevel optimization framework. We incorporate optimization with sample validation (OSV) into the framework to reduce the number of realizations used during optimization. Finally, we describe the use of MLMC in field development optimization.

2.1 Joint optimization of well location and control

In this section we present the optimization problems considered in this work. The PSO–MADS procedure and the treatment of constraints are also described.

2.1.1 Optimization problems

Following Isebor et al. [32], we pose the field development and well control optimization problem as a mixed integer nonlinear programming (MINLP) problem, which

can be stated as follows:

$$(P)\begin{cases}\max\limits_{\mathbf{u}\in U,\mathbf{v}\in V,\mathbf{z}\in\mathbb{Z}} & J\left(\mathbf{u},\mathbf{v},\mathbf{z}\right), \\ \text{subject to} & \mathbf{c}\left(\mathbf{u},\mathbf{v},\mathbf{z}\right)\leq\mathbf{0},\end{cases} \tag{2.1}$$

where J is the objective function we seek to optimize and $\mathbf{c}\in\mathbb{R}^m$ represents the nonlinear constraints. The bounded sets $V=\{\mathbf{V}\in\mathbb{Z}^{n_1};\ \mathbf{v}_l\leq\mathbf{v}\leq\mathbf{v}_u\}$ and $U=\{\mathbf{u}\in\mathbb{R}^{n_2};\ \mathbf{u}_l\leq\mathbf{u}\leq\mathbf{u}_u\}$ define the allowable values for the well placement variables $\mathbf{v}$ and well control variables $\mathbf{u}$. In this work, we use bottom-hole pressure (BHP) as the control variables, though rates could also be used. The vector $\mathbf{z}\in\mathbb{Z}$ denotes discrete categorical variables, which could designate, for example, whether a well is an injector or a producer. Here n_1 and n_2 indicate the number of optimization variables for well placement and well control, respectively.

For problems involving geological uncertainty, we optimize the expected reservoir performance by averaging over multiple geological realizations. In this case Eq. 2.1 can be generalized to

$$\left(\hat{P}\right)\begin{cases}\max\limits_{\mathbf{u}\in U,\mathbf{v}\in V,\mathbf{z}\in\mathbb{Z}} & E[J]=\dfrac{1}{N_{real}}\sum\limits_{s=1}^{N_{real}}J_s\left(\mathbf{u},\mathbf{v},\mathbf{z}\right), \\ \text{subject to} & \mathbf{c}\left(\mathbf{u},\mathbf{v},\mathbf{z}\right)\leq\mathbf{0},\end{cases} \tag{2.2}$$

where N_{real} is the number of (in this case equally probable) geological realizations and $E[\cdot]$ denotes expected value. Note that other objective functions, such as a utility function, could be used instead of $E[J]$.

2.1.2 PSO–MADS algorithm

We now briefly describe the MADS and PSO methods, and then discuss how they are combined in the PSO–MADS hybrid algorithm used in this work.

Mesh adaptive direct search (MADS), developed by Audet and Dennis [5], is a gradient-free, local optimization technique that can be classified as a pattern search algorithm. MADS, which is supported by local convergence theory, involves 'polling'

on a stencil in search space. A MADS iteration entails evaluating the objective function at all points on a stencil emanating from a central point. The central point is the best point, in terms of objective function value, found thus far in the optimization. For a problem with n optimization variables, there are $2n$ stencil points in our implementation of MADS, so $2n$ function evaluations are required at each iteration.

At iteration $k+1$, the central point is shifted to the point that provided the best objective function value at iteration k (different criteria may be used to select the new central point if all solutions are infeasible in terms of nonlinear constraints). The MADS stencil, in contrast to that in generalized pattern search, is not oriented in the coordinate directions of the search space, but rather in random directions that change with iteration. This leads to arbitrarily close poll directions and faster convergence in some cases. If no improvement is achieved at a particular iteration, the stencil size is reduced. MADS stopping criteria involve reaching a minimum stencil size or a maximum number of function evaluations. The algorithm naturally parallelizes because all of the function evaluations (reservoir simulations) can be performed simultaneously. For more details on MADS, see [5], [6], [35] and [31].

Particle swarm optimization (PSO) is a stochastic global search algorithm that was originally developed by Eberhart and Kennedy [20] and first applied for well location optimization (in oil field problems) by Onwunalu and Durlofsky [43]. PSO is based on the social behaviors of swarms of animals and, like a genetic algorithm, involves a set of candidate solutions at each iteration. A particular PSO candidate solution is called a particle and the set of solutions is referred to as the swarm (there are N_p particles in the swarm). Particle i moves through the search space according to the equation

$$\mathbf{x}_i^{k+1} = \mathbf{x}_i^k + \mathbf{v}_i^{k+1}\Delta t, \tag{2.3}$$

where $\mathbf{x}_i = (\mathbf{u}, \mathbf{v}, \mathbf{z})_i$ defines a well location and control scenario; i.e., the location of the particle in the search space, $\mathbf{v}_i$ is the particle velocity, k and $k+1$ indicate iteration level, and Δt is the 'time' increment, typically taken to be 1.

The velocity is comprised of three separate contributions — the so-called inertial, cognitive and social velocity components. The inertial term acts to maintain a degree

of continuity in particle motion by moving the particle in the same direction it was going in the previous iteration. The cognitive term moves the particle toward the best location (in terms of objective function and/or nonlinear constraint violation value) in the search space that it has encountered up to iteration k. The social term moves particle i toward the location of the best particle in its 'neighborhood.' In some PSO procedures, the neighborhood includes all N_p particles in the swarm, but in our implementation we use a random neighborhood topology in which particle i interacts with only a portion of the swarm. This portion changes, in a random fashion, over the course of the optimization. Stopping criteria for PSO involve a maximum number of iterations or function evaluations, or a minimum change in objective function value over one or more iterations. Like MADS, PSO naturally parallelizes because the flow simulations for all particles can be performed simultaneously.

PSO provides global search, though convergence, even to a local minimum, is not guaranteed. MADS, by contrast, provides essentially a local search (though some degree of nonlocality can be achieved by using a large initial stencil size), though it does lead to local convergence in many cases. By combining the two algorithms, the advantages of each approach can be exploited. This hybridization was accomplished by Isebor et al. [31, 32], who demonstrated that the PSO–MADS procedure outperformed both standalone MADS and standalone PSO. These references should be consulted for full details on the procedure. We note that an earlier hybridization involving PSO was developed by Vaz and Vicente [55].

The PSO–MADS hybrid algorithm is illustrated in Figure 2.1 for a minimization problem in two variables (i.e., a search space of dimension two). In the figure, the contour lines indicate objective function value. Note that there is a local minimum (in the lower left in Figure 2.1(a)) as well as a global minimum, which is indicated by the red star. The green points in Figure 2.1(a) depict PSO particles at iteration k. Figure 2.1(b) shows the particle positions at iteration $k+1$. If there is no improvement in the best particle for a specified number of subsequent PSO iterations, then MADS is applied, using the best particle (shown in red in Figure 2.1(b)) as the central point of the stencil. MADS iterations, depicted in Figure 2.1(c), proceed with a fixed stencil size until the solution stops improving, at which point we return to PSO. Other PSO

particles will be attracted to the best particle (through the social velocity term), which typically leads to additional improvement in the solution. The next time MADS is called, the stencil size is reduced. The algorithm proceeds until a minimum MADS stencil size is reached or until a maximum number of function evaluations have been performed.

In the examples presented in this work, we apply the PSO settings used in [31, 32]. Specifically, a random neighborhood topology is applied, and we set the coefficient of the inertial velocity term (ω) to 0.721, and the coefficients of the cognitive and social velocity terms (c_1 and c_2) to 1.193. The PSO swarm size (N_p) varies from case to case, but it is in the range $30 \leq N_p \leq 60$. The initial (maximum) MADS stencil size is 10 fine-scale grid blocks (well locations are always tracked on the finest scale), and the minimum stencil size is 1 fine-scale block. When we switch grid levels in the multilevel optimization procedure, the MADS stencil is reset to its initial size of 10 grid blocks. We note that the minimum MADS stencil size is usually not reached since the algorithm typically terminates once a (specified) maximum number of function evaluations are performed.

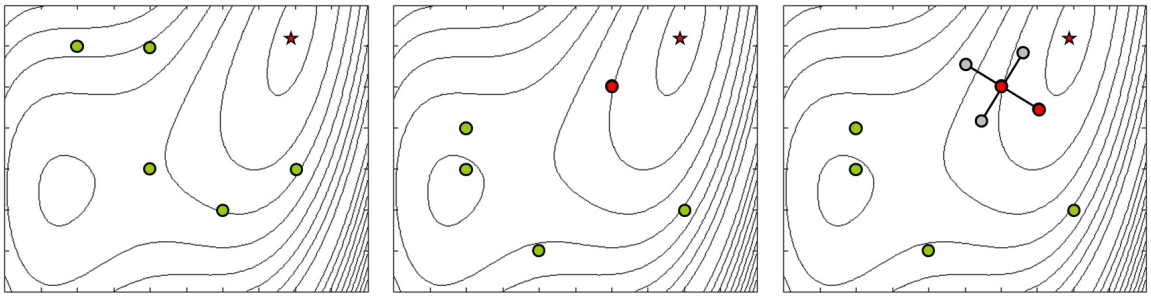

(a) PSO iteration k (b) PSO iteration $k+1$ (c) MADS iteration

Figure 2.1: Illustration of PSO–MADS iterations for a minimization problem in a two–dimensional search space, from Isebor et al. [32]

2.1.3 Nonlinear constraint treatment in optimization under uncertainty

Both bound and nonlinear constraints typically appear in field development and well control optimization problems. Bound constraints, which could include maximum or minimum allowable BHPs, are usually handled by the optimization algorithm, so little if any specialized treatment is required. Nonlinear constraints, by contrast, are more complicated. Examples of nonlinear (output) constraints include minimum oil production rate or maximum water cut from a well when BHP is the control variable. The constraints are nonlinear because they depend on simulation variables, and a nonlinear set of equations (i.e., the flow simulation equations) must be solved to determine these variables.

In Isebor et al. [31], a filter treatment for handling nonlinear constraints in PSO–MADS was developed. This involves viewing the problem in what is essentially a biobjective optimization fashion, and simultaneously minimizing the objective function along with an aggregate constraint violation function. This approach enables the algorithm to consider infeasible (in terms of the nonlinear constraints) solutions during the course of the optimization, even though the final solution is feasible.

When optimizing over multiple realizations (as in Eq. 2.2), however, requiring strict feasibility for all geological realizations might be too severe. A more appropriate procedure in this case may be to instead penalize realizations that do not satisfy the nonlinear constraints. Such an approach is described by Dimitrakopoulos [17], who used a stochastic programming formulation with a penalty for optimizing mining operations under geological uncertainty. Following this approach, we now express the optimization problem as

$$\left(\hat{P}\right)\begin{cases}\max\limits_{\mathbf{u}\in U,\mathbf{v}\in V,\mathbf{z}\in\mathbb{Z}} \; E[J^*]=\dfrac{1}{N_{real}}\displaystyle\sum_{s=1}^{N_{real}}\left(J_s\left(\mathbf{u},\mathbf{v},\mathbf{z}\right)-R_s\left(\mathbf{u},\mathbf{v},\mathbf{z}\right)\right),\end{cases} \tag{2.4}$$

where R_s is the penalty term. Our optimization procedure is quite general and can treat any appropriate objective function. In the examples presented in this paper, we

maximize undiscounted net present value (NPV). We express this objective function, with $\mathbf{x} = (\mathbf{u}, \mathbf{v}, \mathbf{z})$, as

$$J_s(\mathbf{x}) = P_{po}Q_{po}(\mathbf{x}) - P_{pw}Q_{pw}(\mathbf{x}) - P_{iw}Q_{iw}(\mathbf{x}) - C_{drill}(\mathbf{x}), \tag{2.5}$$

where P_{po} is the price of oil (\$/STB), P_{pw} and P_{iw} are the costs of produced and injected water (\$/STB), $Q_{po}(\mathbf{x})$ and $Q_{pw}(\mathbf{x})$ are the cumulative volumes (STB) of oil and water produced, $Q_{iw}(\mathbf{x})$ is cumulative water injected (STB), and $C_{drill}(\mathbf{x})$ is the drilling cost.

The specific form of R_s depends on the actual constraints in the problem. In contrast to penalty functions commonly used in optimization to drive solutions towards feasibility, this penalty is meant to reflect the actual economic cost of the constraint violation. For example, if we have a minimum oil production rate constraint ($q_{o,min}$) and a well does not meet this constraint, we shut the well in and additionally penalize the solution by defining R_s as follows:

$$R_s(\mathbf{x}) = P_{pen}Q_{unmet}(\mathbf{x}), \tag{2.6}$$

where P_{pen} is the 'penalty' price for oil (we would typically take $P_{pen} > P_{po}$) and the cumulative 'unmet' oil production (unmet because the well has been shut in) is given by $Q_{unmet} = q_{o,min} \times t_{shutin}$, where t_{shutin} is the total amount of time during which the well is shut in. If the well is shut in because of a different constraint violation (e.g., produced water rate exceeds the maximum), we still apply the penalty in this form. Treatments other than that in Eq. 2.6 could of course be used, and would in fact be required for cases that do not involve the specification of $q_{o,min}$.

2.2 Upscaling procedure

In this work, we apply a multilevel optimization procedure, where grids at several different levels of refinement are used over the course of the optimization run. The coarse-grid properties for each scenario are determined by applying a flow-based upscaling procedure. The problems considered here involve oil-water systems. In coarse-scale models for such systems, both single-phase (i.e., permeability and porosity) and

two-phase (relative permeability) properties can be upscaled. Chen and Durlofsky [14] and Durlofsky and Chen [19] discuss many of the upscaling methods available for problems of this type.

In the current implementation, we upscale only single-phase flow parameters and not the two-phase flow functions. This means we compute upscaled transmissibility, designated T^*, for each coarse block-to-block interface, and upscaled well index, WI^*, for each coarse block in which a well is completed. This type of coarse-scale model is in general less accurate than a model that additionally includes upscaled relative permeability functions [14]. However, this approach is more computationally efficient, since the calculation of upscaled relative permeability functions is time consuming, and it provides reasonable accuracy as the grid is refined. For very coarse grids our approach does incur some error but, as we will see, the ranking of solutions proposed by the optimization algorithm is largely maintained. In other words, candidate solutions that would be among the best (in terms of objective function value) when evaluated on the finest scale are among the best when evaluated on the coarse scale, and similarly for median and poor solutions.

A variety of single-phase parameter upscaling techniques exist, and here we apply what is essentially the most accurate method available. Specifically, we use a global transmissibility upscaling procedure of the type described by Chen et al. [16] and Zhang et al. [62]. With this method, we first solve the global single-phase pressure equation with flow driven by the actual wells:

$$\nabla \cdot (\mathbf{k} \nabla p) = q, \tag{2.7}$$

where p is pressure, $\mathbf{k}$ is the diagonal permeability tensor, and q denotes the well-driven source term. Wells are represented using the usual Peaceman [44] well index. Recall that at each iteration of either MADS or PSO, multiple solutions $\mathbf{x}_i$ must be evaluated, and each of these solutions corresponds to a different set of wells and controls. In our optimization framework, we solve Eq. 2.7 for each proposed well scenario $\mathbf{x}_i$ at each iteration of the optimizer. Thus, if the optimization involves 30 PSO particles, we solve the fine-scale pressure equation 30 times at each PSO iteration. The source term q also depends on the controls $\mathbf{u}_i$ embedded in $\mathbf{x}_i$, and

these controls can change in time. Because Eq. 2.7 is a steady-state equation, we average the time-varying behavior in $\mathbf{u}_i$ to arrive at the source term q.

We describe the upscaling procedure with reference to the schematic in Figure 2.2(a), which shows a 15×15 fine grid (lighter lines) and a 5×5 coarse grid (heavier lines). Wells are depicted by the $\times$'s. A region corresponding to two coarse blocks is highlighted in Figure 2.2(a) and extracted in Figure 2.2(b). Given the fine-scale pressure solution corresponding to well scenario $\mathbf{x}_i$, we compute T_j^* by relating the integrated (summed) flow rate through coarse interface j to the (estimated) coarse-scale pressure difference. This results in the following expression for T_j^*:

$$T_j^* = \frac{\sum_l f_l}{\langle p \rangle_{j-} - \langle p \rangle_{j+}}, \tag{2.8}$$

where f_l designates the flow rate through fine-scale interface l, which lies on coarse interface j, and $\langle p \rangle_{j-}$ and $\langle p \rangle_{j+}$ are the bulk-volume averages of the fine-scale pressure over coarse blocks $j-$ and $j+$. We note that, in the case of very small $\sum_l f_l$ and/or $(\langle p \rangle_{j-} - \langle p \rangle_{j+})$, it is possible that the resulting T_j^* will be anomalous (i.e., negative or extremely large). In such cases we replace the T_j^* from Eq. 2.8 with a value computed from the geometric averages of the fine-scale permeabilities in blocks $j-$ and $j+$.

For a coarse block j containing a well, WI_j^* is computed in an analogous manner:

$$WI_j^* = \frac{\sum_l f_l^w}{\langle p \rangle_j - p^w}. \tag{2.9}$$

Here the sum is over the l fine-scale well blocks lying in coarse well block j (the sum is needed for, e.g., vertical wells in three-dimensional models), f_l^w is the flow rate into or out of the well in fine-scale block l, $\langle p \rangle_j$ is the bulk-volume average of the fine-scale pressure over coarse block j, and p^w is the wellbore pressure (averaged over coarse block j if necessary). It is important to note that wells do not need to be in the center of the well block except at the finest scale (see Figure 2.2(a)). The effect of a well being off-center in a coarse block is captured by WI^* and the well-block T^*. This is a very useful feature of this upscaling method, as it allows us to consistently track the well location on the finest grid regardless of the refinement level in use in the optimization algorithm.

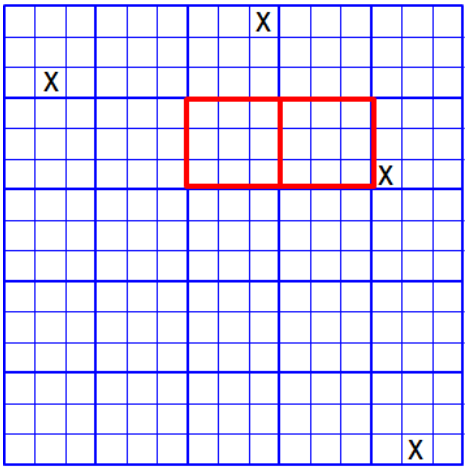

(a) 15×15 fine grid (lighter lines) and 5×5 coarse grid (heavier lines). Wells designated by $\times$'s

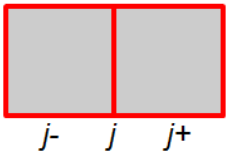

(b) Two coarse-block region for T_j^* computation

Figure 2.2: Schematic showing fine and coarse grids and transmissibility upscaling

The use of T^* and WI^* computed using Eqs. 2.8 and 2.9 generally provides coarse-scale models that can closely replicate the integrated single-phase flow behavior of the fine-scale model. As such, these coarse models are very useful in our workflow. Chen et al. [16] observed, however, that improved accuracy could be achieved in many cases by iterating on the coarse-model properties. These iterations involve first solving the pressure equation for the coarse-scale model with the T^* and WI^* computed from Eqs. 2.8 and 2.9. This provides pressure in every coarse-scale block. By replacing $\langle p \rangle_{j-}$ and $\langle p \rangle_{j+}$ in Eq. 2.8 with the actual coarse-scale pressures p^c_{j-} and p^c_{j+}, we obtain an updated estimate for T^* (and similarly for WI^*). Rather than use the new T^* and WI^* directly, a damping procedure is applied. In this work we apply 10 iterations of this procedure. This is still inexpensive because these are all coarse-scale computations. Refer to Chen et al. [16] for full details on this iteration procedure.

Figure 2.3 displays results for a channelized model using the iterative transmissibility upscaling procedure described above. The figure shows maps of y-direction transmissibility for various coarse grids (Figures 2.3(a)-(d)) and for the finest grid (Figure 2.3(e)). This model contains three production and two injection wells (shown as red and blue circles, respectively). We reiterate that the upscaled properties depend

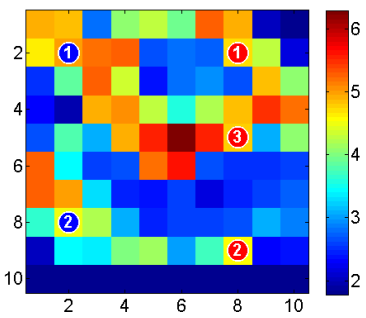

(a) Level 1 (10×10 grid)

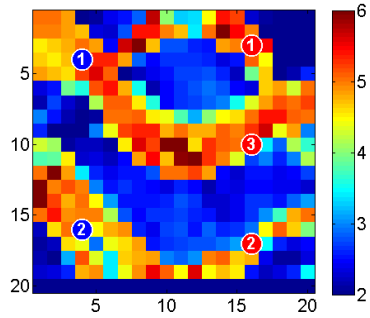

(b) Level 2 (20×20 grid)

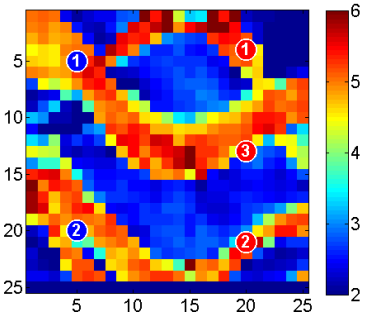

(c) Level 3 (25×25 grid)

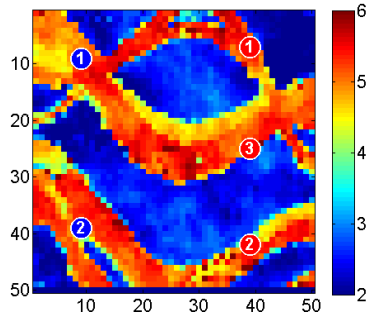

(d) Level 4 (50×50 grid)

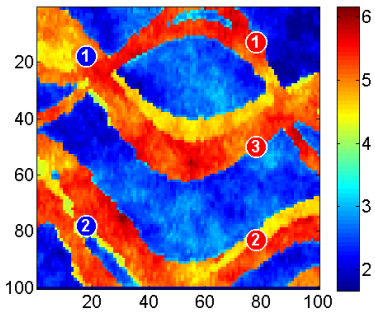

(e) Level 5 (100×100 grid)

Figure 2.3: Log_{10} transmissibility in y-direction at different coarsening levels. Production and injection wells shown as red and blue circles

on these well locations and their BHPs.

In this work, we extended an existing two-dimensional iterative transmissibility upscaling procedure to upscale three-dimensional models. Figure 2.4 shows a map of the y-direction transmissibility for a fine-scale $30 \times 30 \times 6$ model. This three-dimensional permeability model is used in Case 2 in Chapter 3, and the simulation parameters are presented in Table 3.2. The model contains four production wells and one injection well. All wells are perforated in all model layers. We applied the upscaling method to this three-dimensional model. The transmissibility map for one of the coarse-grid models, of dimensions $15 \times 15 \times 2$, is shown in Figure 2.5. Coarse layer 1 in this figure corresponds to layers 1–3 in the fine model, and coarse layer 2 to fine layers 4–6.

It is apparent that the channel resolution is fairly low for the $15 \times 15 \times 2$ coarse model, though even the coarsest grid provides results of high accuracy for single-phase flow quantities. This is demonstrated in Figure 2.6, where we show the steady state production rate for production well 1 (P1) at different coarsening levels. We see that the error in the single-phase flow rate, relative to the fine-scale solution, is less than 5% at all coarsening levels. Similar results are obtained for other wells.

Next, we use the upscaled models for two-phase flow simulations (using AD-GPRS). Oil and water production rates for P1 and P4 at different coarsening levels are plotted in Figures 2.7 and 2.8. We also plot the water injection rate for I1 in Figure 2.9. We see that the two-phase flow effects increase the error in the coarse-scale simulation results. However, as the number of grid cells in the upscaled models increases, the error in the two-phase simulation output reduces and the upscaled models provide results close to those from the fine model.

2.3 Multilevel optimization framework

As noted in the Introduction, upscaled models have been used previously for optimizing well locations [1] and well control [34]. Related procedures include the multigrid-based optimization methods developed by Lewis and Nash [36, 41]. In these approaches, coarse-resolution problems were used to generate search directions

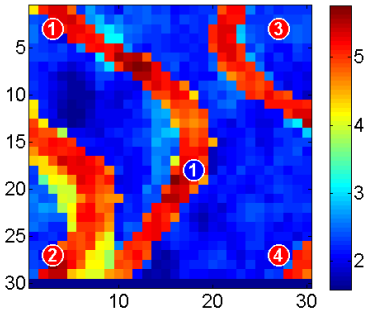

(a) Layer 1

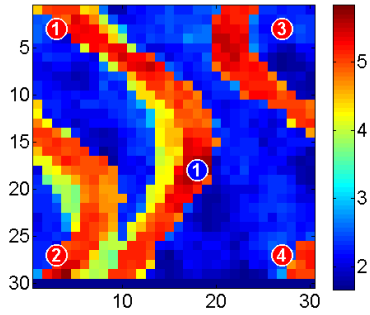

(b) Layer 2

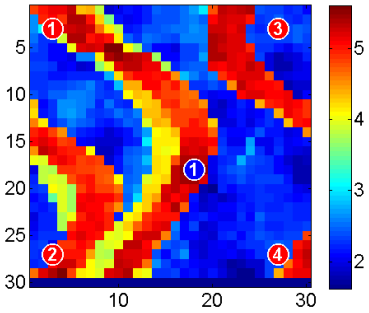

(c) Layer 3

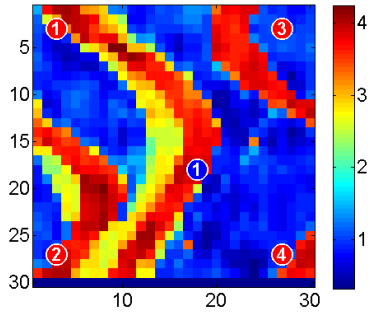

(d) Layer 4

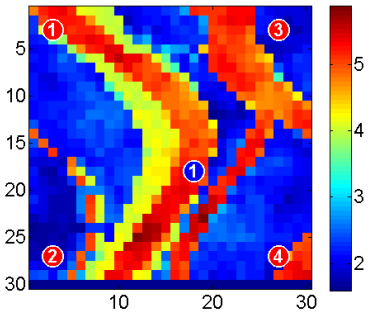

(e) Layer 5

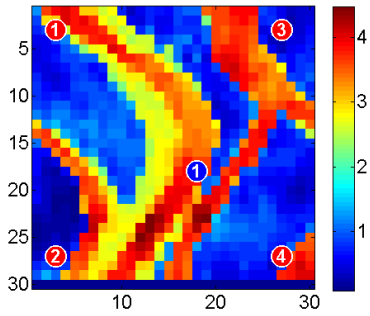

(f) Layer 6

Figure 2.4: Log_{10} transmissibility in y-direction for $30 \times 30 \times 6$ fine model. Production and injection wells shown as red and blue circles. Model from [29]

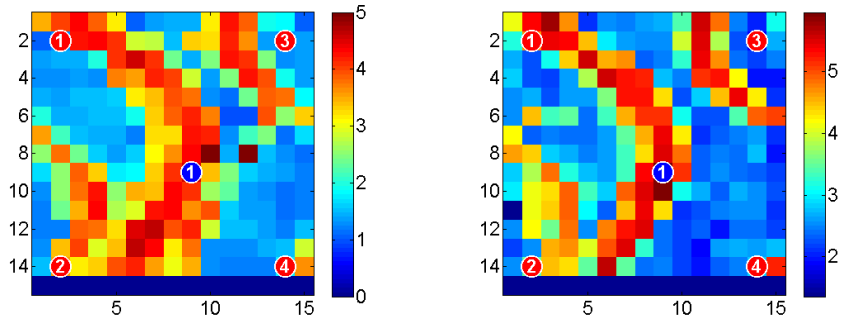

(a) Coarse layer 1 (b) Coarse layer 2

Figure 2.5: Log_{10} transmissibility in y-direction for $15 \times 15 \times 2$ coarse model generated from the model in Figure 2.4. Production and injection wells shown as red and blue circles.

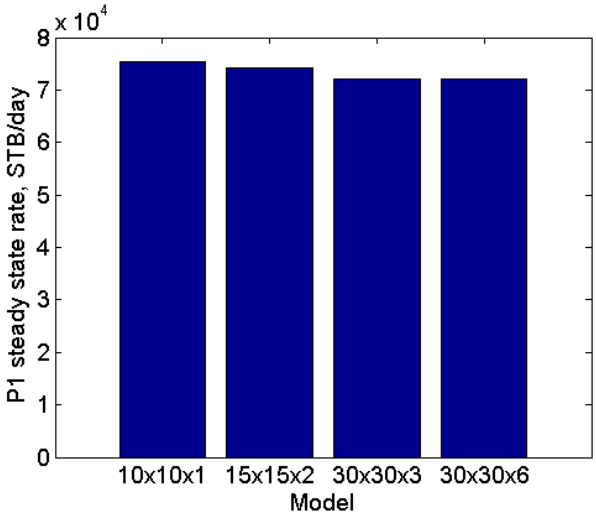

Figure 2.6: Upscaling results for single–phase flow (Case 2 in Chapter 3)

(for gradient-based optimization) for finer-resolution problems. Our approach differs from earlier methods in that we apply a multilevel approach, in which a set of subproblems are optimized sequentially. In addition, our approach is suitable for use in challenging field development problems, in which both well locations and well controls are optimized. Although our method is incorporated into the PSO–MADS framework, it could be used with a wide range of optimization procedures, including gradient-based, direct search and meta-heuristic algorithms.

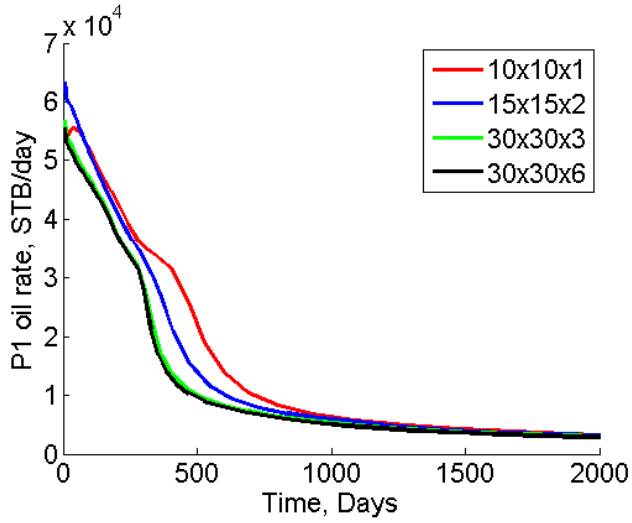

(a) P1 oil production rate

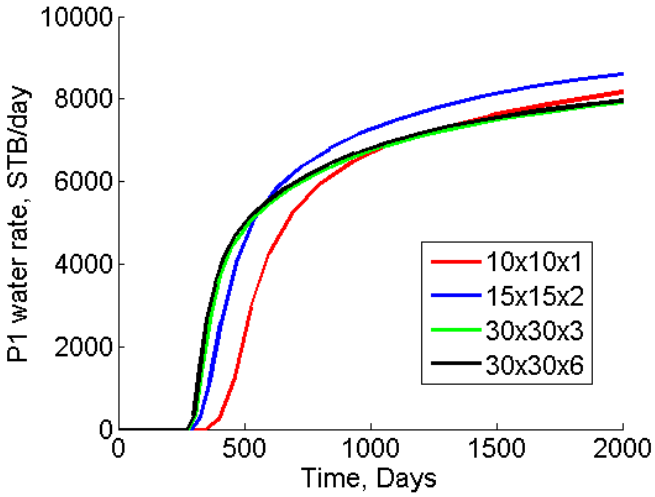

(b) P1 water production rate

Figure 2.7: Upscaling results for two-phase flow for well P1 (Case 2 in Chapter 3)

We now describe the multilevel procedure for optimization with a single geological model. The key component of the method is the global transmissibility upscaling technique presented in the previous section. For specificity, assume the finest-grid model contains 100×100 grid blocks (as in Figure 2.3(e)). We begin the optimization with a coarse model containing 10×10 grid blocks (as in Figure 2.3(a)). We refer to this coarse-grid optimization problem as P_1, and the grid as the Level 1 grid.

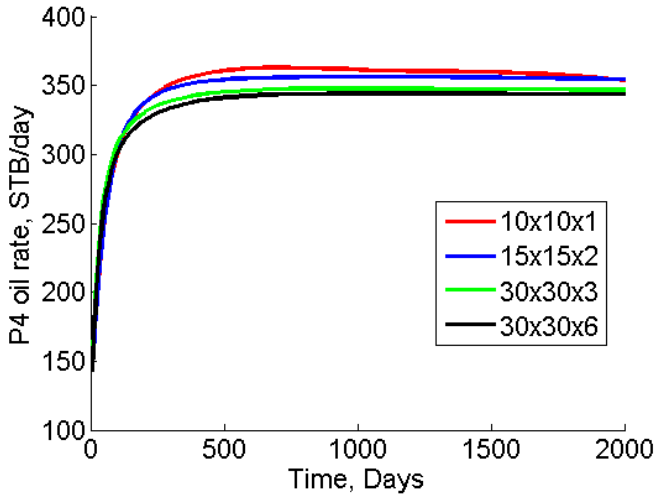

(a) P4 oil production rate

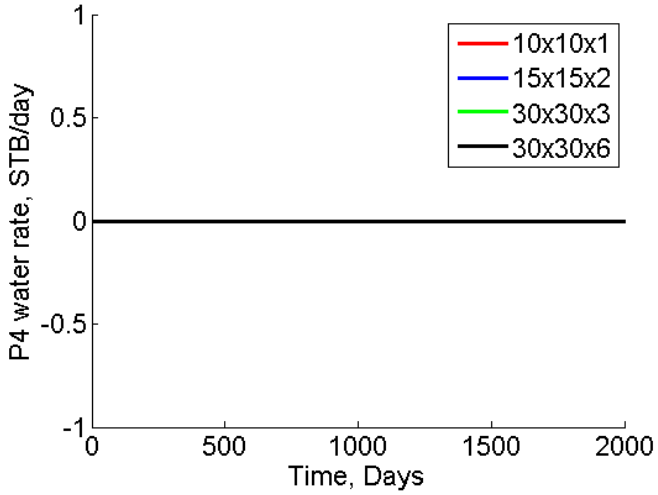

(b) P4 water production rate

Figure 2.8: Upscaling results for two-phase flow for well P4 (Case 2 in Chapter 3)

Recall that, at any iteration (involving either PSO or MADS), many different well scenarios are considered. For each scenario $\mathbf{x}_i$, we generate an upscaled Level 1 model. This is necessary since the upscaled properties depend on the well locations and BHP settings. Although these computations are more time consuming than simpler upscaling approaches, they can be much more accurate than alternate procedures, and they are still fast compared to solving the fine-scale two-phase flow problem.

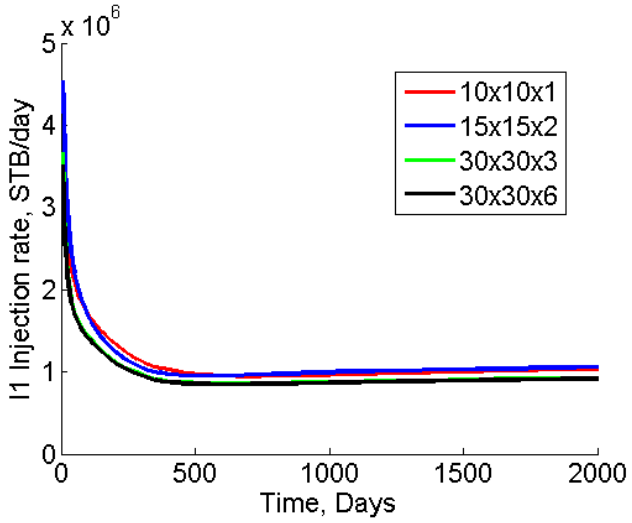

Figure 2.9: Upscaling results for well I1 (Case 2 in Chapter 3)

We refer to the result from the P_1 optimization problem as $\mathbf{x}_1^*$. This solution is then used as an initial guess for optimization problem P_2. In PSO, an initial swarm is required, and the other particles in the swarm are generated randomly. The Level 2 grid in this case contains 20×20 grid blocks (as in Figure 2.3(b), though the well locations would now correspond to $\mathbf{x}_1^*$). Note that, because well locations are always defined on the finest grid, they are not perturbed when we change grid levels. We proceed in this manner, using the solution to optimization problem P_{m-1}, designated $\mathbf{x}_{m-1}^*$, as the initial guess for problem P_m. The optimization at any level ends when we reach the stopping criteria. These are either a maximum number of function evaluations (this number can vary from level to level) or a minimum MADS stencil size. As noted earlier, the maximum number of function evaluations is typically reached before the minimum stencil size. When the optimization is started on the next level, the MADS stencil size is initialized to its maximum (10 fine-grid blocks in our examples).

In the two-dimensional multilevel optimization examples presented here, the fine-grid model is defined on a 100×100 grid. The grid levels included in the multilevel procedure are determined by applying the same integer upscaling ratio in both the $x-$ and $y-$directions; i.e., we consider 10×10, 20×20, 25×25 and 50×50 grids.

It should be possible, however, to skip some of these grid levels, or to use different stopping criteria at different levels. This should be a topic for future investigation.

Our multilevel approach shares some similarities with the retrospective optimization (RO) procedure presented by Wang et al. [58]. In that work, a sequence of subproblems was used for optimization under geological uncertainty. The number of geological realizations considered in the optimization increased from subproblem to subproblem. In both RO and the current procedure, early subproblems, which require substantial exploration by the optimizer and thus many function evaluations, entail much faster flow computations. Later subproblems involve more expensive function evaluations, but fewer iterations are required because the solution is 'closer' to the optimum (by optimum here, we mean the best solution found by the optimization algorithm).

To apply the multilevel procedure for optimization under geological uncertainty, we consider multiple (N_{real}) geological models, and now take the objective function to be the expected (average) response over the N_{real} realizations. Then, for each function evaluation, we compute $E[J^*] = (1/N_{real}) \sum_s (J_s - R_s)$ as in Eq. 2.4. This means that at each PSO iteration, $N_p \times N_{real}$ simulations are required, and at each MADS iteration, $2n \times N_{real}$ simulations must be performed. This can be very expensive computationally for large N_{real} (e.g., for $N_{real} \sim O(100)$), and in such cases a cluster-sampling procedure, such as that used by [58], could be applied to select a representative subset of realizations. Alternatively, a procedure such as a multilevel Monte Carlo method (described in Section 2.5), which entails performing a 'telescopic sum' involving different numbers of realizations simulated at different levels of refinement, could be considered [40]. The approach used most extensively in this work, optimization with sample validation (OSV), is described in the next section. With any of these treatments, the multilevel optimization procedure now requires that we construct coarse-scale models (i.e., perform flow-based upscaling) for each well location and control scenario $\mathbf{x}_i$, in each geological realization $\mathbf{m}_s$, evaluated in subproblem P_m. Although this corresponds to a large number of upscaling computations, it is still highly cost effective relative to performing fine-scale two-phase flow simulations.

The multilevel optimization framework described thus far uses PSO–MADS [31, 32] as the core optimizer at all grid levels used in the optimization. Using the PSO

component, which provides global search, at all levels of the optimization may, however, not be necessary. This is based on our observation that the best solution often changes relatively little after the first level, and many of these changes are local in character. We therefore introduce a modified (accelerated) multilevel field development optimization procedure that uses only the local optimizer (MADS) in the later levels. This will be shown to save a significant amount of computation time.

2.4 Optimization with sample validation

Because the number of realizations used in the optimization directly scales the number of simulation runs required, it is important to reduce the number of realizations to the extent possible. Toward this aim, we employ the optimization with sample validation (OSV) approach developed by Shirangi and Durlofsky [50], which we now describe. We assume that the subsurface geological uncertainty is characterized using a total of N_{real} realizations. We start the optimization with a small number (N_{rep}) of 'representative' realizations. 'Representative' realizations are selected from the full set of N_{real} realizations using the CDF approach that we will discuss later. At a particular grid level, we designate $\mathbf{x}^{l-1}$ and $\mathbf{x}^{l}$ as the initial and optimum solutions respectively. Then, the increase in the objective function value, with N_{rep} realizations, is $J^*\{\mathbf{x}^l, N_{rep}\} - J^*\{\mathbf{x}^{l-1}, N_{rep}\}$.

Our goal in the optimization is however to improve the objective function over all N_{real} realizations. We thus evaluate all N_{real} realizations for the $\mathbf{x}^{l-1}$ and $\mathbf{x}^{l}$ solutions to quantify overall improvement, which is given by $J^*\{\mathbf{x}^l, N_{real}\} - J^*\{\mathbf{x}^{l-1}, N_{real}\}$. We now compute the relative improvement (RI), which is the ratio of the improvement in objective function value with all N_{real} realizations to the improvement with N_{rep} representative realizations:

$$RI = \frac{J^*\{\mathbf{x}_m^l, N_{real}\} - J^*\{\mathbf{x}_m^{l-1}, N_{real}\}}{J^*\{\mathbf{x}_m^l, N_{rep}\} - J^*\{\mathbf{x}_m^{l-1}, N_{rep}\}}, \tag{2.10}$$

where the subscript m refers to the current coarsening level in the multilevel optimization framework.

We do not proceed to the next level in the optimization unless $RI \geq \theta$, where θ is a user-specified validation criterion (here we set $\theta = 0.5$). If $RI < \theta$, then the number of representative models N_{rep} is increased and the PSO-MADS procedure is repeated with the larger number of representative models. OSV is terminated if $RI \geq \theta$ is satisfied or if a maximum number of OSV steps is reached. See [50] for further details.

The number of OSV steps and the number of representative realizations at each step are specified by the user. Different approaches can be used to select the representative realizations out of the full set of N_{real} realizations. In this work, we base the selection on the cumulative distribution function (CDF) of the N_{real} model responses [50]. The CDF of the objective function is constructed based on $\mathbf{x}^{l-1}$, which is the current optimum scenario. Then, the two realizations that correspond to the P10 and P90 values (10th and 90th percentiles) on the CDF are selected. Next, the CDF between P10 and P90 is equally divided into $N_{rep} - 2$ regions and the median realization from each region is selected. This procedure is illustrated in Figure 2.10.

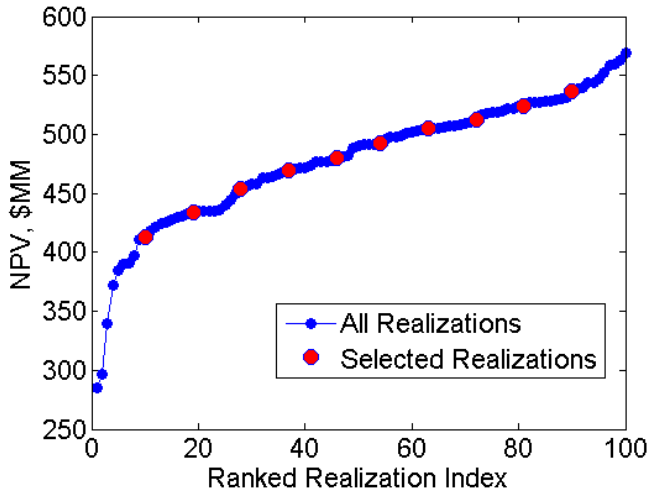

Figure 2.10: Ordered NPV plot for $N_{real} = 100$ realizations. The ten selected N_{rep} realizations are designated by the red points

A flowchart of the overall multilevel optimization procedure with OSV is shown in Figure 2.11. For each grid level in the multilevel optimization framework, we perform

the OSV procedure. We start OSV with a small number of representative models and then increase the number of realizations as required, until the RI criterion is satisfied or a maximum number of OSV steps is reached. Then, the model is refined (i.e., we proceed from P_m to P_{m+1}) and OSV is restarted with the finer model.

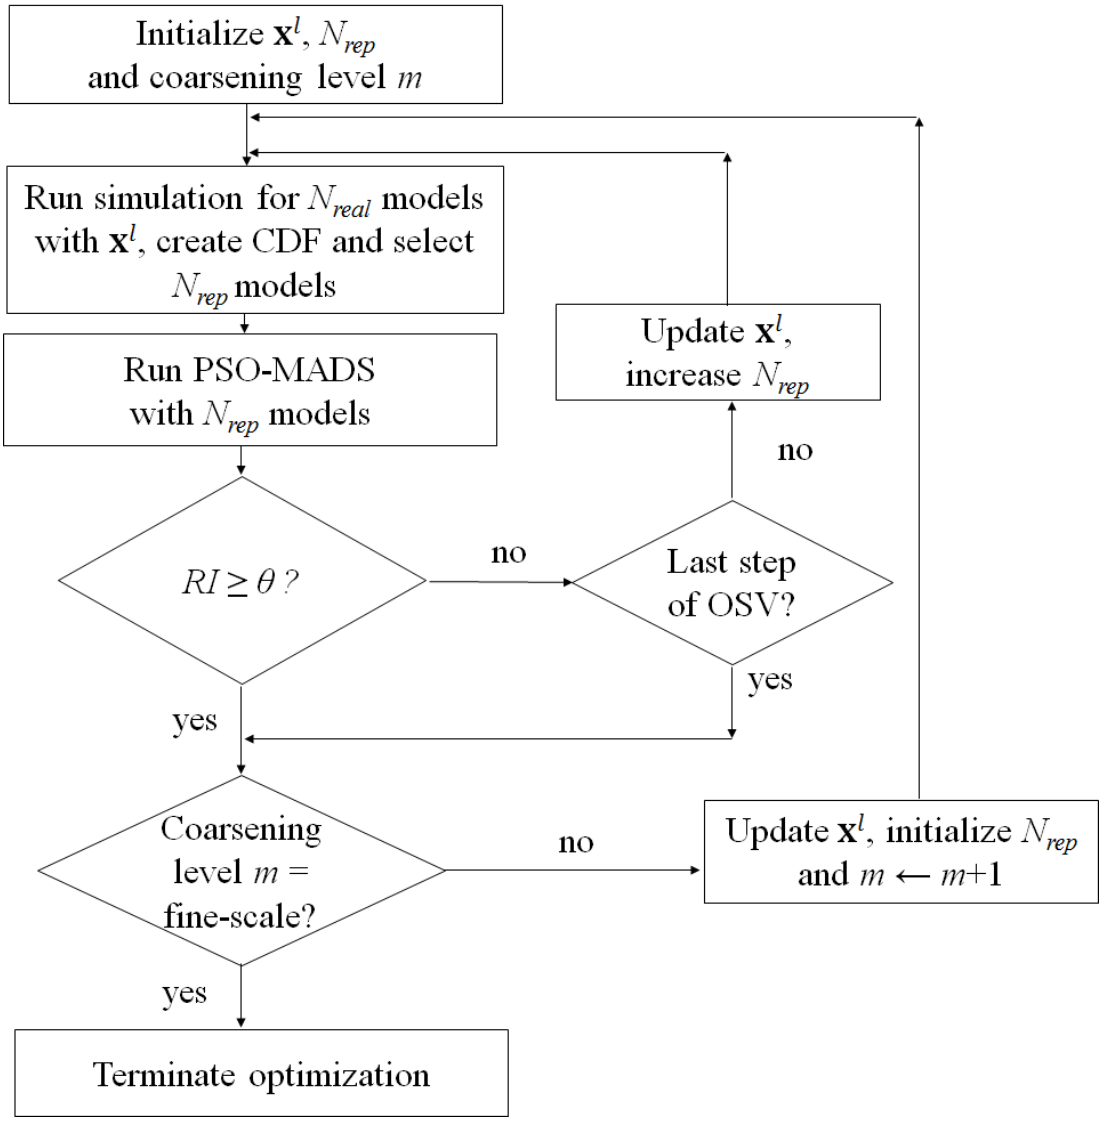

Figure 2.11: Flowchart of the multilevel optimization procedure with OSV

2.5 Optimization under uncertainty using MLMC

The multilevel Monte Carlo (MLMC) method enables an efficient estimate of the expected value of a reservoir simulation output over multiple realizations. MLMC is based on a telescopic sum with different numbers of realizations at different coarsening levels. Müller et al. [40] combined MLMC with two–phase streamline simulation to assess the impact of geological uncertainty in oil–water flow problems. They showed that MLMC provided accurate results with an order of magnitude speedup. We now describe the MLMC procedure and its application to field development optimization.

Assume for now that we wish to compute a fine-scale quantity using solutions on five grid levels. The fine-scale quantity, J_5^*, can be represented in terms of a telescopic sum as follows:

$$\begin{aligned} J_5^* = J_5^* + (J_4^* - J_4^*) + (J_3^* - J_3^*) + (J_2^* - J_2^*) + (J_1^* - J_1^*) = (J_5^* - J_4^*) + \\ (J_4^* - J_3^*) + (J_3^* - J_2^*) + (J_2^* - J_1^*) + (J_1^* - J_0^*) = \sum_{m=1}^{5} (J_m^* - J_{m-1}^*). \end{aligned} \tag{2.11}$$

Note that decreasing m corresponds to coarser grid levels, with J_1^* corresponding to the coarsest model. We set $J_0^* = 0$ in Eq. 2.11. The general form of the equation can be written as follows

$$J_M^* = \sum_{m=1}^{M} (J_m^* - J_{m-1}^*), \tag{2.12}$$

where J_M^* refers to a quantity on the finest grid level and M refers to the number of levels in the telescopic sum.

In field development optimization under uncertainty, we need to compute the expected value of the objective function over multiple geological realizations of the fine-scale model. The telescopic sum is thus extended to compute $E[J_M^*]$, the expected value

of the fine-scale solutions over all N_{real} realizations. We can express $E[J_M^*]$ as:

$$E[J_M^*] = \sum_{m=1}^{M} \frac{1}{N_{real}} \sum_{i=1}^{N_{real}} (J_{i,m}^* - J_{i,m-1}^*). \tag{2.13}$$

Now, we introduce the MLMC approximation to reduce computation time. The key is to use different numbers of realizations at different coarsening levels. More realizations are used at coarser levels than at finer levels. For each level m we define N_m to be the number of models used at that level. The expected objective function value $E[J_M^*]$ is now approximated by the MLMC estimator $E'[J_M^*]$ as

$$E[J_M^*] \approx E'[J_M^*] = \sum_{m=1}^{M} \frac{1}{N_m} \sum_{i=1}^{N_m} (J_{i,m}^* - J_{i,m-1}^*). \tag{2.14}$$

Here $(J_{i,m}^* - J_{i,m-1}^*)$ denotes the difference between solutions at levels m and $m-1$ for realization i, which means they derive from the same permeability realization.

In [40], the N_m models were randomly sampled from the full set. In this work, we also employ the CDF approach described previously [50] to select models. This approach is compared to MLMC estimation using random sampling. Also, we investigate the effect of the number of grid levels (M) on MLMC performance.

Chapter 3

Single-Realization Optimization Results

In this chapter, we apply the multilevel optimization procedure described in Chapter 2 to optimization problems involving a single realization. The first two examples are for two- and three-dimensional models, with fixed numbers of injectors and producers, while the third example involves the optimization of well number and type. We also perform conventional optimization for each of these examples, using only the fine-scale model, to enable comparisons with the multilevel optimization results.

Any field development optimization algorithm can be used in the multilevel optimization framework. In this work, the PSO–MADS hybrid algorithm [29], described in Chapter 2, is used for all optimization runs. Key economic and operational parameters are presented in Table 3.1. The objective in all cases is to maximize undiscounted NPV, given in Eq. 2.5.

Flow simulations are performed using either Stanford's General Purpose Research Simulator (GPRS), described in [12] and [33], or the more recent Automatic Differentiation-based General Purpose Research Simulator (AD–GPRS), described in [63]. All cases involve oil-water flow. The relative permeability curves used in all runs are shown in Figure 3.1. Other simulation parameters are presented in Table 3.2.

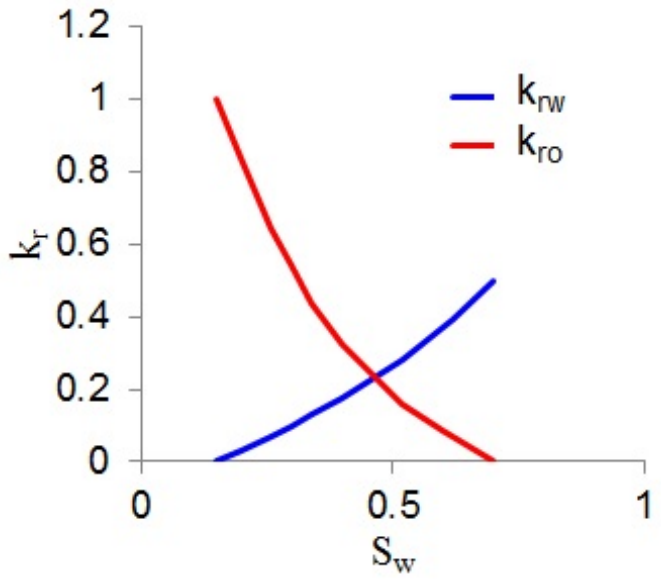

Figure 3.1: Relative permeability curves used for all simulations

Table 3.1: Optimization parameters used in the example cases

Parameter	Case 1	Case 2	Case 3
P_{po}, P_{pw} and P_{iw}, \$/STB	100, 5 and 5	100, 5 and 5	100, 10 and 5
Well drilling cost, \$ million per well	5	5	100
Injection BHP range, psi	6000 – 9000	6500 (fixed)	6000 – 9000
Production BHP range, psi	1500 – 4500	1500 (fixed)	1500 – 4500

3.1 Case 1: Two-dimensional channelized model

In this case we optimize the locations of three production wells and two injection wells, along with three BHP control values for each well. There are thus a total of $n = 25$ optimization variables (ten location variables and 15 control variables). The production time frame is 2190 days, and BHPs are determined at initial time, at 730 days, and at 1460 days. The reservoir model, shown in Figure 3.2, was generated by Isebor and Durlofsky [30] using SGeMS [45]. An initial guess for the well locations is also shown in Figure 3.2. For this case we specify $N_p = 50$ for PSO. Each MADS

Table 3.2: Simulation parameters used in the example cases

Parameter	Cases 1 and 3	Case 2
Grid size	100×100	30×30×6
Grid cell dimensions	32.8 ft×32.8 ft×32.8 ft	100 ft×100 ft×20 ft
Initial pressure p_i, at datum	2900 psi at 3280 ft	3500 psi at 8400 ft
Rock compressibility c_R, psi^{-1}	0.5 $\times 10^{-5}$	0.3 $\times 10^{-5}$
μ_o and μ_w, cp	3 and 1	1.20 and 0.31
ρ_o and ρ_w, lbm/ft^3	62.4 (both)	49.6 and 64.79
B_o and B_w, RB/STB	1.00	1.00

iteration also entails $2n = 50$ function evaluations (simulations). Given that we can access 50 computational cores, the optimization in this case is fully parallelized. Nonlinear constraints are not included in this example.

In this and subsequent cases, the multilevel optimization at each level stops when the maximum number of simulation runs (prescribed for that level) is reached. The specific numbers of simulation runs will be presented below. We run the single-level optimization for the same number of total simulation runs, which in this case is 13,300.

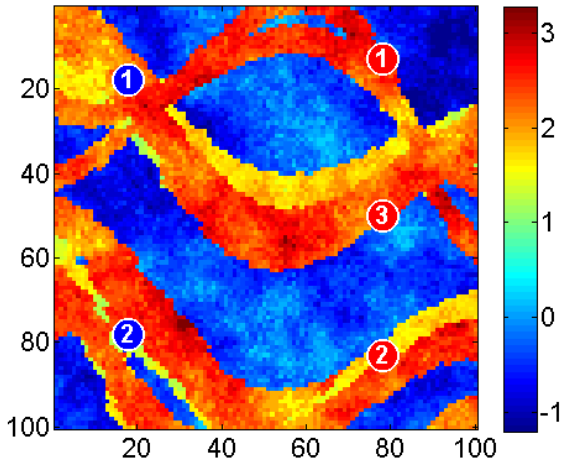

Figure 3.2: Log_{10} permeability (in md) for Case 1. Model from [30]

Before discussing optimization results, we first assess the performance of the upscaled models for this case. The well configuration considered in this assessment is shown in

Figure 3.2. Figure 3.3 displays the oil and water production rates for models at different levels of coarsening. The 10×10 models clearly overpredict oil production and underpredict water production by substantial margins, but they do capture general trends. As we refine, the coarse model results are seen to systematically approach the fine-scale (100×100) results for both oil and water production rates. This feature is very useful in the context of multilevel optimization.

Figure 3.4 displays the NPVs associated with 50 well scenarios (these are the 50 candidate scenarios considered by the optimizer after 6000 simulation runs) at various grid levels. The NPV computed for the fine model is plotted against the NPVs for the coarse models. If the coarse models were perfectly accurate, all points would fall on the 45° line. It is evident that, as the grid is refined, the points generally shift towards the 45° line. We also see that the coarse models typically overpredict NPV. This bias is consistent with the flow results shown in Figure 3.3, where the coarse models overpredict oil production while underpredicting water production. The degree of NPV overprediction in Figure 3.4 clearly decreases with increasing grid resolution.

An important feature of Figure 3.4 is the fact that the general ordering of the 50 models is reasonably well captured, even at the coarsest grid level. In other words, well scenarios that are among the best (in terms of NPV) for the finest grid are also identified as among the best for the 10×10 or 20×20 grids. This indicates that these models are indeed useful for optimization, where the relative performance of the various scenarios is of primary importance.

The timings for upscaling and the subsequent two-phase flow simulation for the various models are shown in Table 3.3. The number of simulation runs and elapsed time at the different grid levels are presented in Table 3.4. We see in Table 3.3 that the coarser models run much faster than the fine model. From Table 3.4 it is evident that although the elapsed times at the various grid levels are comparable (these timings assume 50 cores are always available; i.e., the optimization is fully parallelized), many more iterations are performed using coarse models than fine models. This results in substantial overall computational savings.

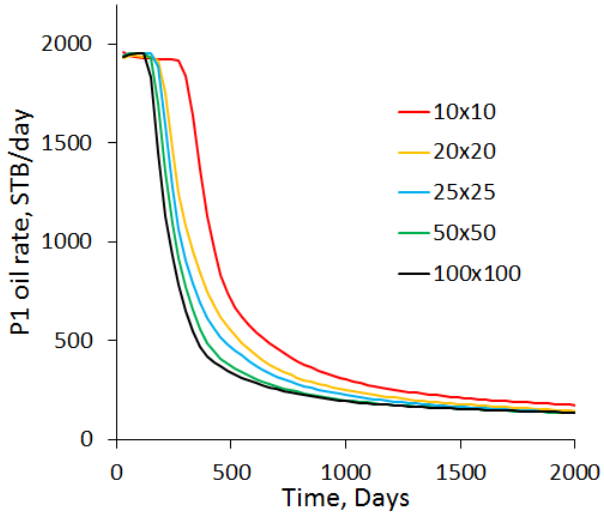

(a) P1 oil production rate

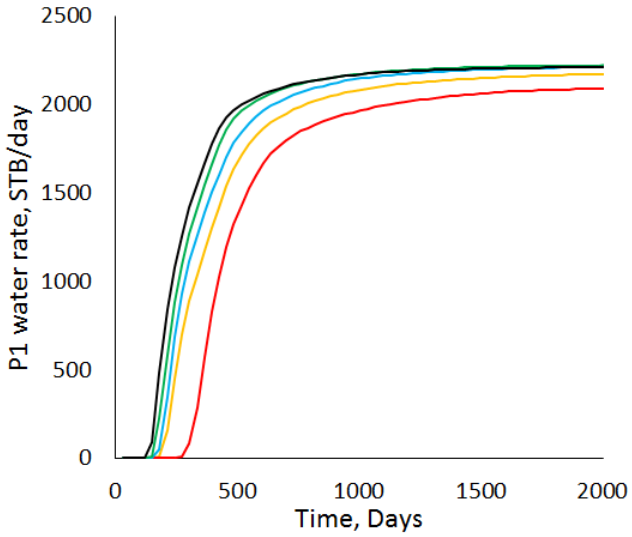

(b) P1 water production rate

Figure 3.3: Upscaling results for two-phase flow (Case 1)

Optimization results are reported in Table 3.5. Results are shown for both the conventional (single-level) optimization approach, with all computations performed on the 100×100 fine grid, and for the multilevel procedure. We perform three runs for each approach due to the stochastic nature of the PSO–MADS algorithm. In these runs, the single-level procedure uses 13,300 simulation runs, which is the same as is

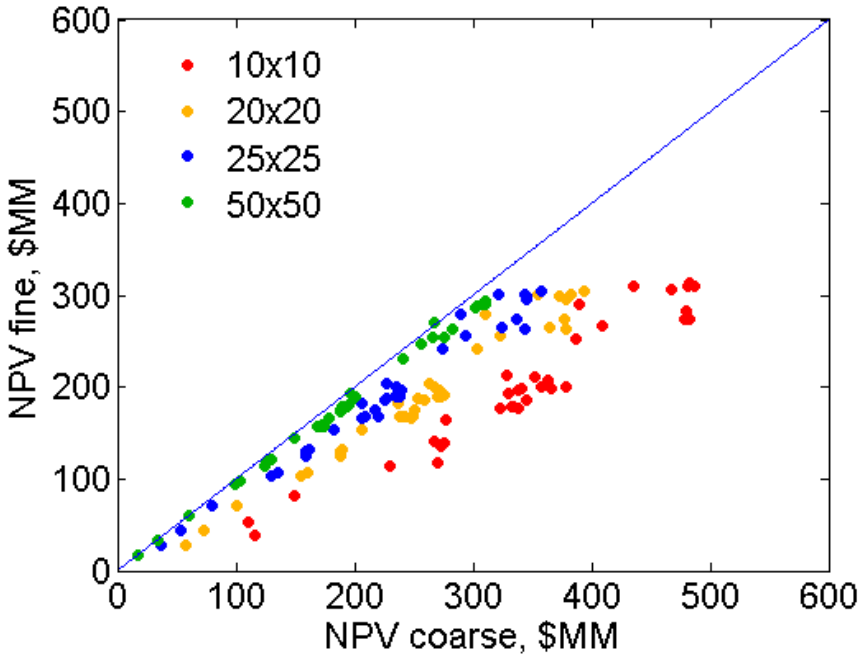

Figure 3.4: Comparison of NPVs evaluated at different grid levels for the 50 well scenarios after 6000 simulation runs (Case 1)

used in the multilevel runs. The best run for each method is shown in bold. The results from the two approaches are quite comparable. The conventional approach provides an average NPV that is about 0.2% higher than that from the multilevel procedure, but the best multilevel optimization result is about 0.2% higher than the best single-level result. As expected, the multilevel procedure is much faster than the conventional approach; here it requires only about one tenth of the elapsed time of the single-level method.

Table 3.3: Upscaling and simulation times for different grid levels (Case 1)

Model size	10×10	20×20	25×25	50×50	100×100
Upscaling time per model, sec	3	4	4	5	–
Simulation time per model, sec	1	2	6	15	84
Total time per model, sec	4	6	10	20	84

The progress of the optimization for the conventional and multilevel approaches is

Table 3.4: Multilevel optimization computations (Case 1)

Model size	10×10	20×20	25×25	50×50	100×100
Total time per model, sec	4	6	10	20	84
Number of function evaluations	6000	4000	2000	1000	300
Elapsed time (fully parallelized), hours	0.13	0.13	0.11	0.11	0.14

Table 3.5: Optimization results for three runs. Best result shown in bold (Case 1)

	Run 1 ($MM)	Run 2 ($MM)	Run 3 ($MM)	Average NPV ($MM)	Time (hours)
Conventional opt.	322.5	**338.8**	329.5	330.3	6.2
Multilevel opt.	334.2	315.4	**339.5**	329.7	0.6

shown in Figures 3.5(a) and (b), where we plot NPV versus the number of simulation runs. The objective function value for the conventional approach improves very little after about 7000 simulation runs, though it does continue to increase (very slowly) until about 12,500 iterations. If we stopped the conventional optimization after 7000 simulation runs, we would obtain a speedup of five (instead of ten) using the multilevel approach. In Figure 3.5(b), the solid line shows the current NPV, at the grid level in use at that point in the optimization. The NPV for the current best $\mathbf{x}_i$, computed using the fine-scale model, is also shown (as the black points) when the grid level changes. There are clear jumps as we proceed from one grid level to the next, and these jumps are substantial for the coarser grid levels. There are also significant differences between the NPV computed for the fine-scale model and that computed for the coarse-scale model. However, the fine-scale NPV clearly increases with iteration even though the actual value is not very well approximated on coarse grids. This occurs because the coarse-grid models are indeed capable of identifying promising scenarios, which are then improved upon through additional PSO–MADS iterations.

The best well configurations found by each method are shown in Figure 3.6. The locations clearly differ between the two runs, but both scenarios show injection wells linked to production wells through channel sand. The optimized BHP controls for production well P1 and injection well I1 are shown in Figures 3.7(a) and (b). Results

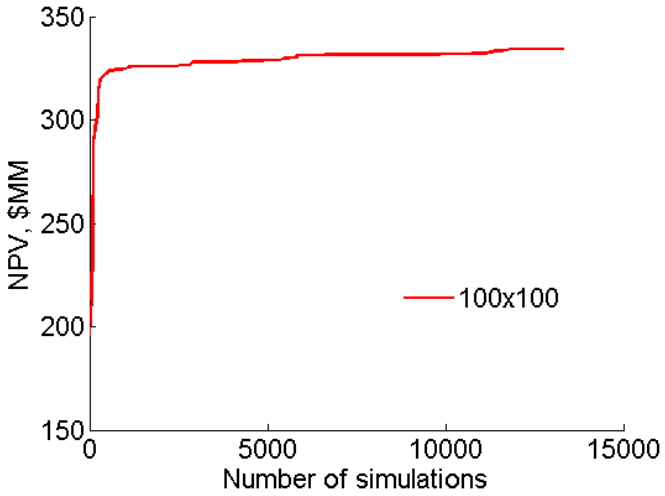

(a) Conventional approach

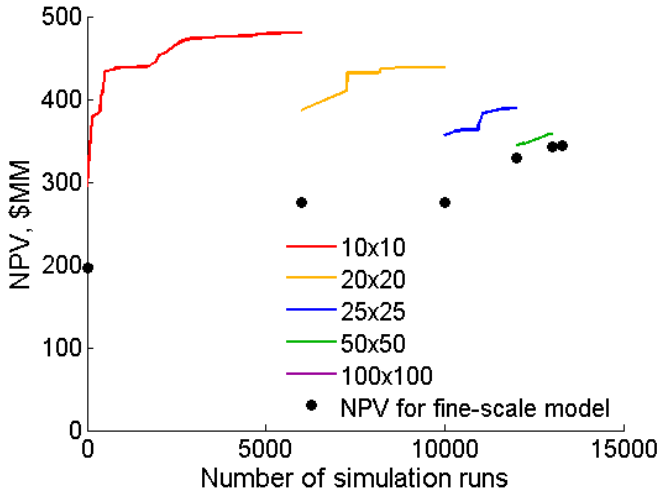

(b) Multilevel approach

Figure 3.5: Evolution of objective function (Case 1)

for the run leading to the highest NPV for each method are shown. The reference solution corresponds to the initial guess with well locations shown in Figure 3.2. The multilevel optimization introduces BHP variation in time, while the BHPs from the conventional approach are nearly constant. Final saturation maps for these optimized solutions are presented in Figures 3.8(a) and (b). The channels are well swept in both cases, as would be expected.

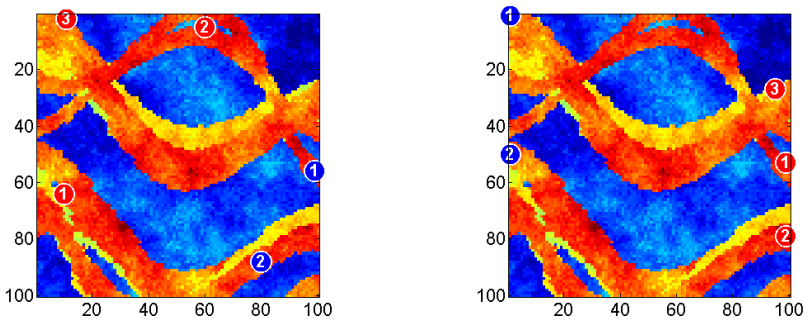

(a) Conventional approach (b) Multilevel approach

Figure 3.6: Best solutions found by the two methods (Case 1). Background shows $\log_{10} k$

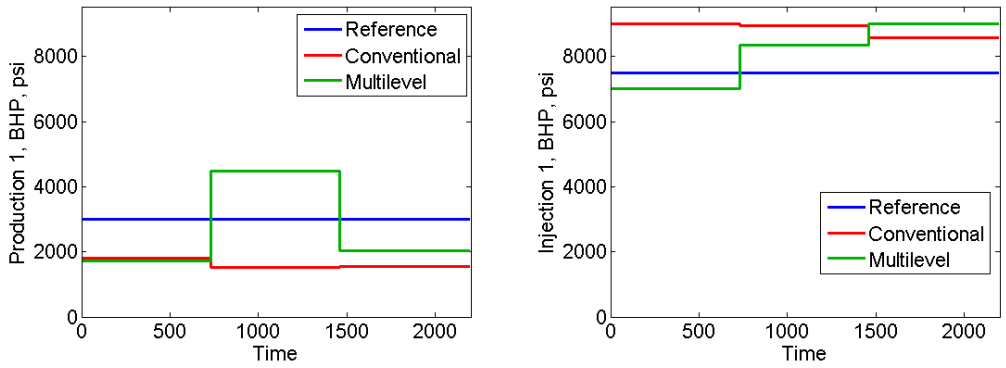

(a) Production well P1 (b) Injection well I1

Figure 3.7: Optimum BHPs found by the two methods (Case 1)

We plot the cumulative production profiles for these optimized solutions and for the reference (initial guess) scenario in Figures 3.9(a)-(c). From the production profiles, it is evident that cumulative oil production in both optimized solutions is about 1 MM-STB higher than in the reference scenario. Cumulative water injection and production are also less than in the reference scenario. The two optimized solutions provide similar cumulative production and injection profiles. This agreement is consistent with the very similar saturation maps in Figure 3.8.

The optimum well locations at the end of each level in the multilevel optimization are displayed in Figure 3.10. Note that the injection wells do not shift very much after

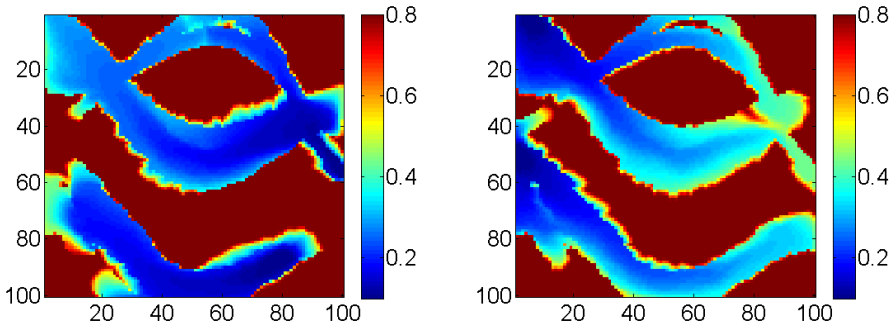

(a) Conventional approach (b) Multilevel approach

Figure 3.8: Comparison of the final oil saturation maps (red indicates oil and blue water), for the optimized solutions found by the two methods (Case 1)

the Level 1 optimization is completed, and the producers are very near their final locations after the Level 3 optimization. This illustrates the ability of the coarser models to identify near-optimal well locations.

3.2 Case 2: Three-dimensional channelized model

In this example we optimize well locations only; BHP values are prescribed (see Table 3.1). Four vertical production wells and one vertical injection well are considered. For each well, we optimize the areal location $(x,\ y)$ and the vertical completion interval $(z_1,\ z_2)$. This corresponds to a total of 20 optimization variables. For PSO we take $N_p = 30$, and for MADS we use $2n = 40$ poll points. For this case 40 computational nodes are required to fully parallelize the optimizations. The permeability field for this case, taken from Isebor [29], is shown in Figure 2.4 (in the previous chapter). Nonlinear constraints are not included in this example. The total simulation time in this case is 2000 days.

Figure 3.11 shows the NPVs for the 30 well scenarios after 3600 simulation runs computed at each grid level. As in Figure 3.4, we see that the accuracy of the models increases with grid refinement and that the general ordering is captured even using the coarsest $(10 \times 10 \times 1)$ models. The timings and number of simulation runs for

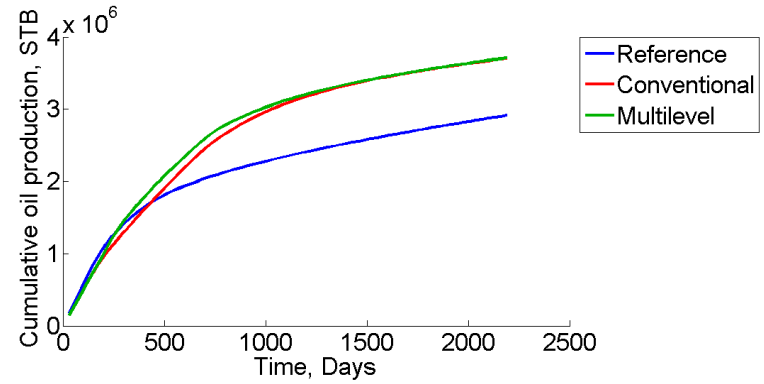

(a) Cumulative oil production

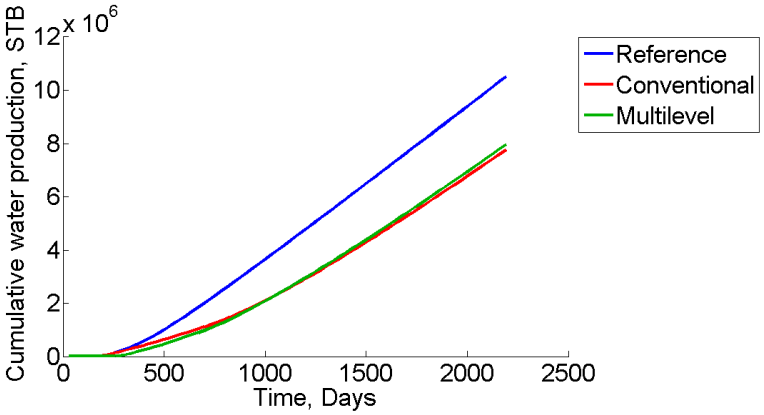

(b) Cumulative water production

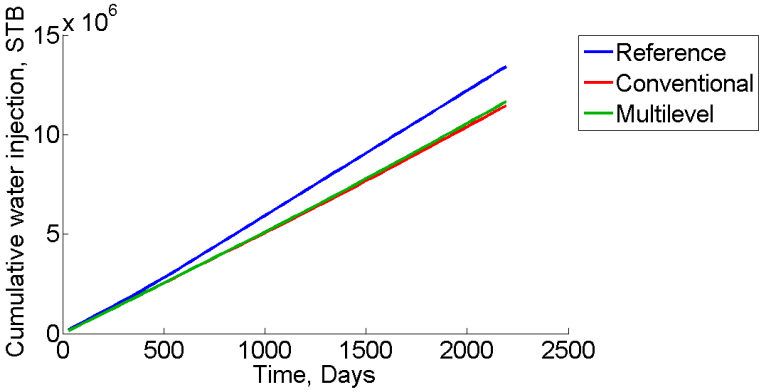

(c) Cumulative water injection

Figure 3.9: Comparison of cumulative production and injection profiles for the optimized and reference solutions (Case 1)

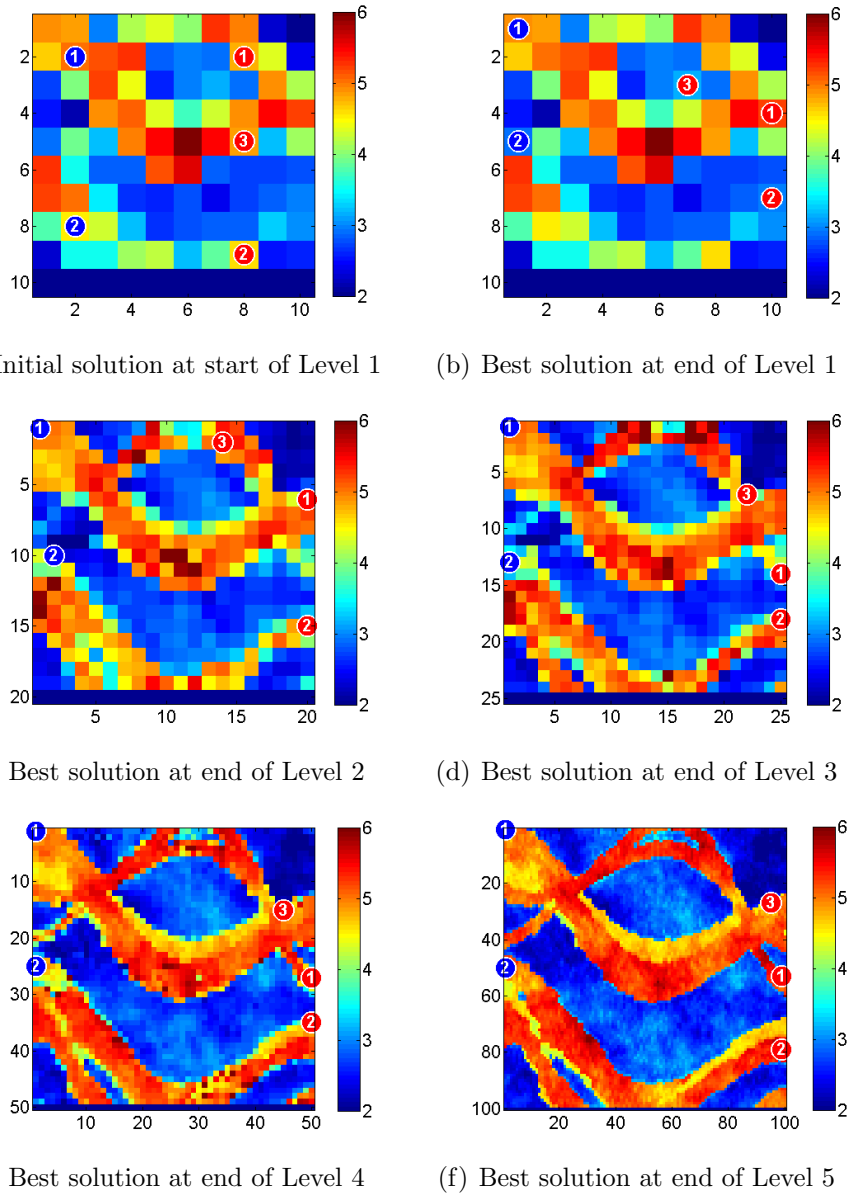

(a) Initial solution at start of Level 1

(b) Best solution at end of Level 1

(c) Best solution at end of Level 2

(d) Best solution at end of Level 3

(e) Best solution at end of Level 4

(f) Best solution at end of Level 5

Figure 3.10: Evolution of optimum well locations in multilevel procedure (Case 1). Background shows $\log_{10} T_y^*$

the various models are shown in Table 3.6. Again we see that a large number of the PSO–MADS iterations are performed using coarse models, and that these models run much faster than the finest-scale model. Results for three runs, again using both the single-level and multilevel approaches, are shown in Table 3.7. Optimized NPVs are

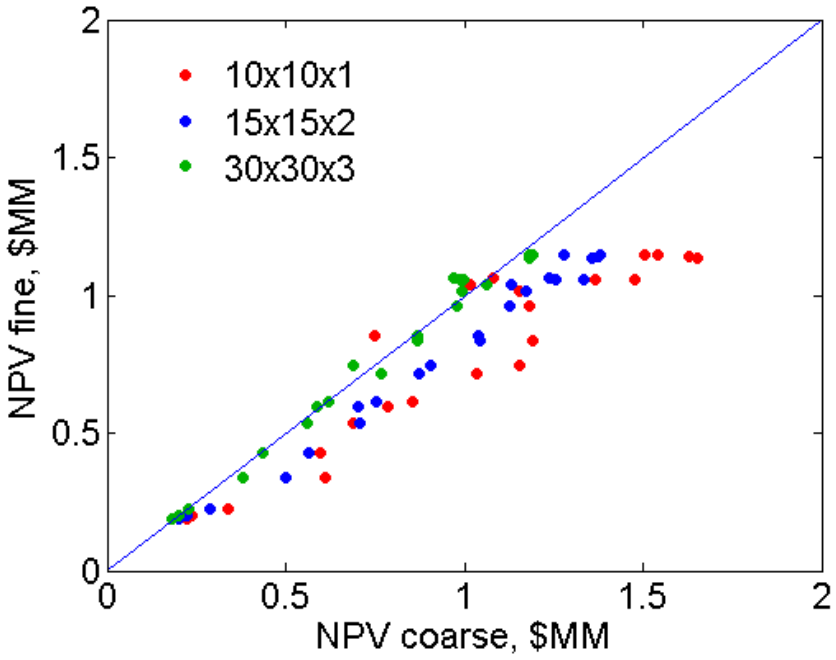

Figure 3.11: Comparison of NPVs evaluated at different grid levels for the 30 well scenarios after 3600 simulation runs (Case 2)

very similar between the two methods, though the multilevel approach provides a factor of eight speedup.

Figures 3.12(a) and (b) show the evolution of NPV with iteration for the conventional and multilevel approaches. The objective function value stops improving after about 4000 simulation runs in the conventional approach. The optimization continues running after 4000 simulation runs because none of the termination criteria is satisfied. If

Table 3.6: Multilevel optimization computations (Case 2)

Model size	$10 \times 10 \times 1$	$15 \times 15 \times 2$	$30 \times 30 \times 3$	$30 \times 30 \times 6$
Total time per model, sec	6	9	39	111
Number of function evaluations	3600	1800	600	180
Elapsed time (fully parallelized), hours	0.20	0.15	0.22	0.19

Table 3.7: Optimization results for three runs. Best result shown in bold (Case 2)

	Run 1 ($B)	Run 2 ($B)	Run 3 ($B)	Average NPV ($B)	Time (hours)
Conventional opt.	**1.21**	1.15	1.20	1.19	6.4
Multilevel opt.	**1.21**	1.19	1.20	1.20	0.8

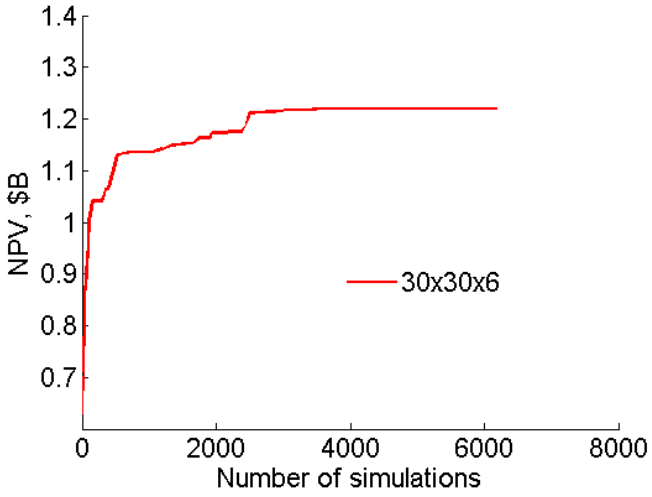

(a) Conventional approach

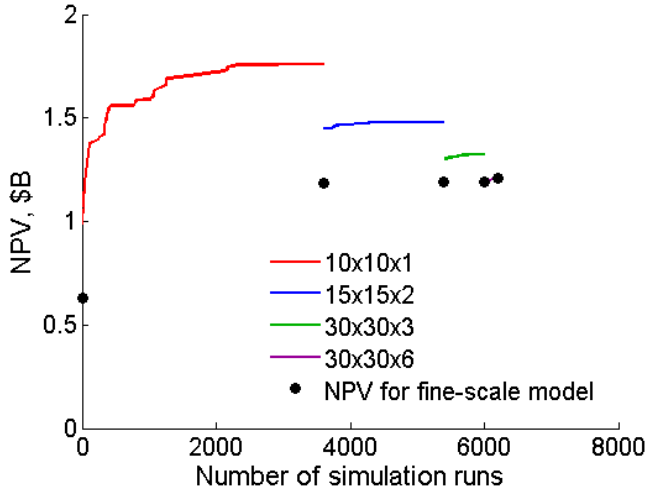

(b) Multilevel approach

Figure 3.12: Evolution of objective function (Case 2)

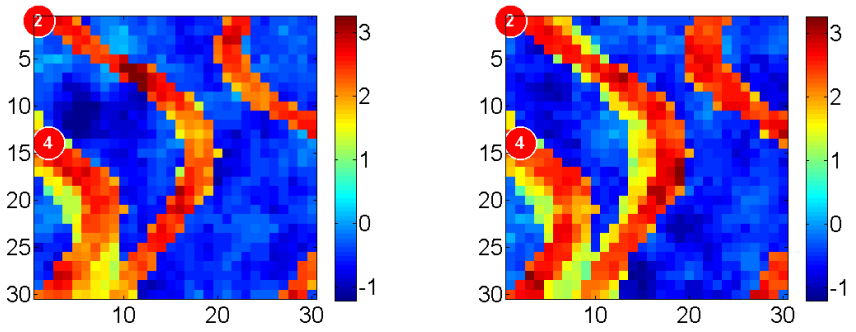

(a) Layer 1

(b) Layer 2

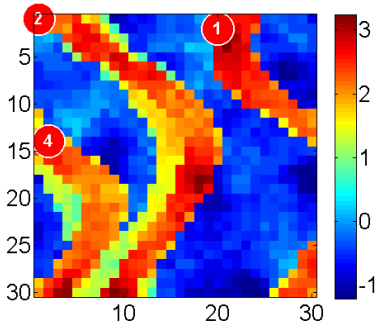

(c) Layer 3

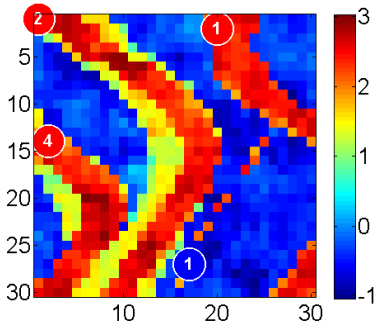

(d) Layer 4

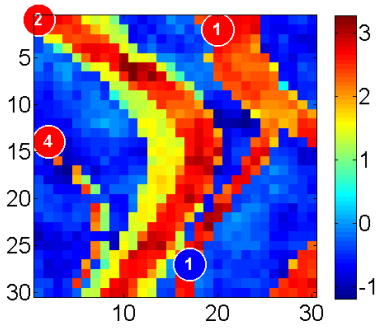

(e) Layer 5

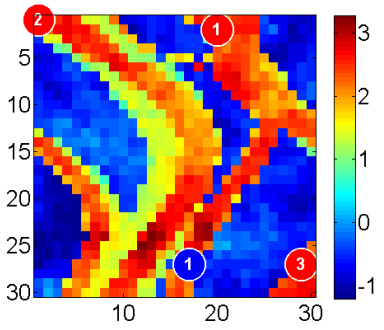

(f) Layer 6

Figure 3.13: Optimum well locations and completion intervals from multilevel procedure (Case 2). Background shows $\log_{10} k$

better termination criteria were used for the conventional approach, then the speedup would be closer to a factor of five.

The behavior of the multilevel approach (Figure 3.12(b)) in this case is similar to

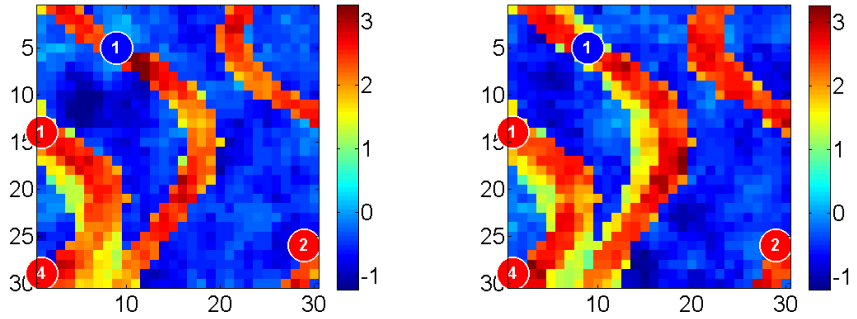

(a) Layer 1

(b) Layer 2

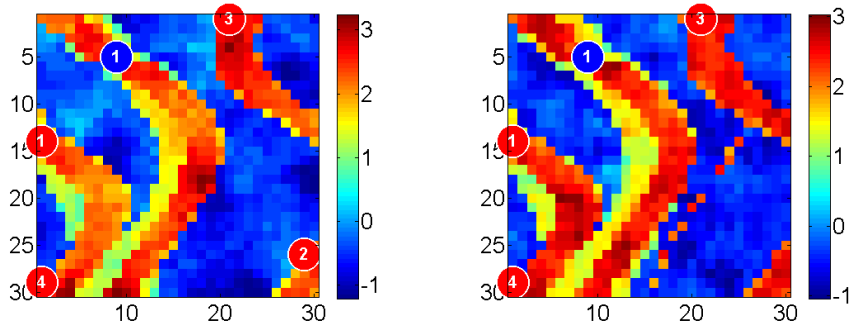

(c) Layer 3

(d) Layer 4

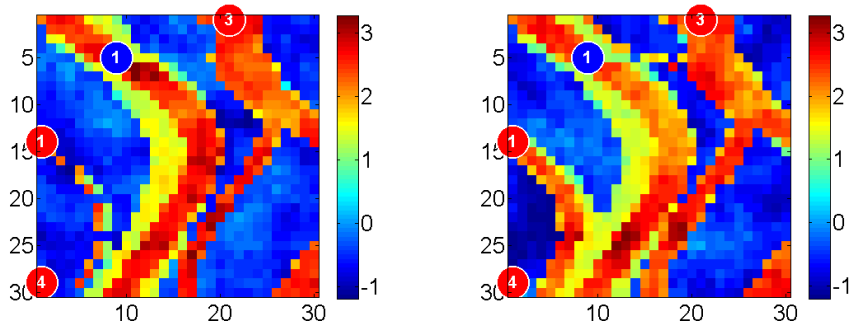

(e) Layer 5

(f) Layer 6

Figure 3.14: Optimum well locations and completion intervals from conventional approach (Case 2). Background shows $\log_{10} k$

that observed in Figure 3.5(b). We again see that the fine-scale NPV continues to improve, though there is only a small amount of improvement after the completion of the Level 1 optimization. The optimized well locations and completion intervals from the best multilevel optimization are shown in Figure 3.13. The injector is completed

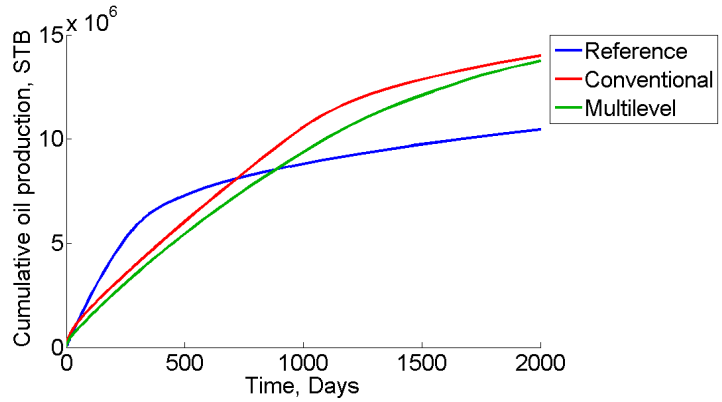

(a) Cumulative oil production

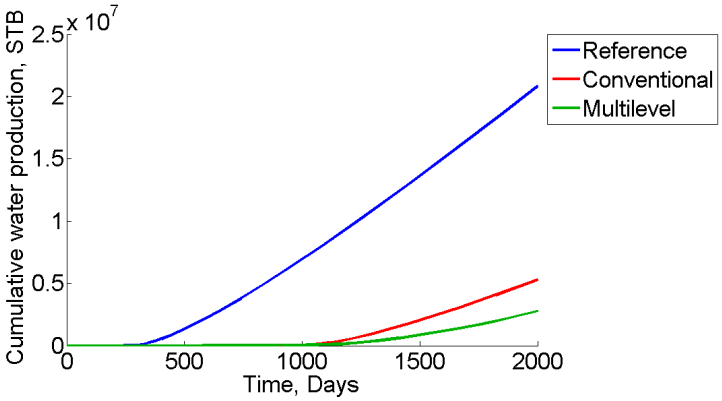

(b) Cumulative water production

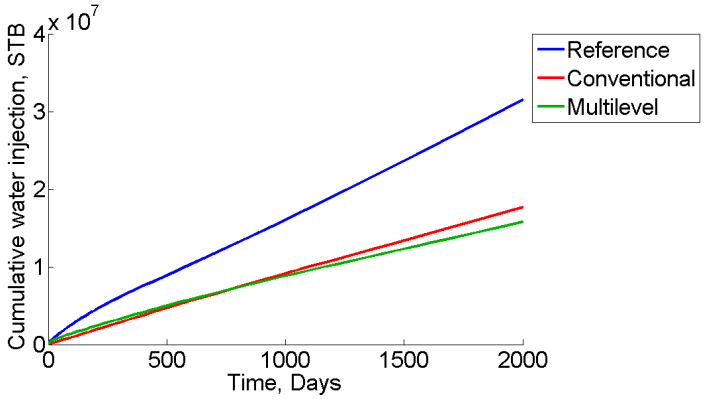

(c) Cumulative water injection

Figure 3.15: Comparison of cumulative production and injection profiles for the optimized and reference solutions (Case 2)

in only the lower three layers. With the exception of P2, all producers are also completed in only some of the layers. The well locations from the best single-level optimization run are shown in Figure 3.14. This well configuration is quite different from that in Figure 3.13, though the NPVs for the two cases are essentially the same.

Cumulative production and injection profiles, for the optimized solutions found by the conventional and multilevel approaches, are presented in Figures 3.15(a)-(c). The reference solution corresponds to the well locations shown in Figure 2.4 and the BHPs prescribed in Table 3.1. Both of the optimized solutions provide similar cumulative oil production, which is about 4 MMSTB more oil than in the reference (initial guess) scenario. This is achieved by injecting significantly less water than in the reference scenario. We also see that water breakthrough occurs about 700 days later in the optimized cases than in the reference case.

3.3 Case 3: Inclusion of categorical variables

The permeability field for this example is shown in Figure 3.16. The model, taken from [56], represents a channelized system with channels oriented in the northwest direction.

In this example, we seek to determine the optimal number of wells (maximum of five wells), their type (injector or producer), locations and controls. For each well, we have one ternary variable (specifying if the well is to be an injector, a producer, or not to be drilled at all) along with two location variables and three BHP variables. The total number of optimization variables is thus 30 (five ternary variables, ten location variables, and 15 control variables). The production time frame is 2190 days, and BHP controls are determined at initial time, at 730 days, and at 1460 days. We have specified well costs to be high in this case ($100 MM, as shown in Table 3.1), so we expect the optimizer to eliminate wells during the course of the optimization. We set $N_p = 60$, and the number of simulation runs per MADS iteration is also 60 ($2n$). We use 60 computer nodes to fully parallelize the PSO-MADS optimization procedure.

We now assess the performance of the upscaling method for this case. A cross-plot of

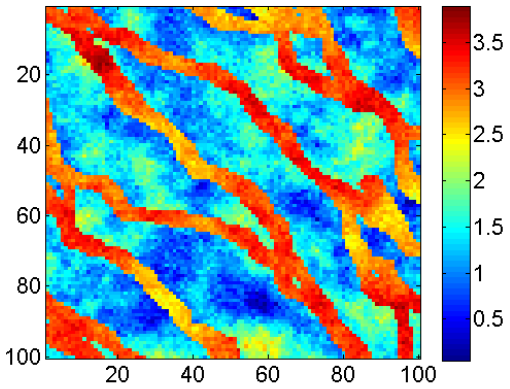

Figure 3.16: Log_{10} permeability (in md) for Case 3. Model from [56]

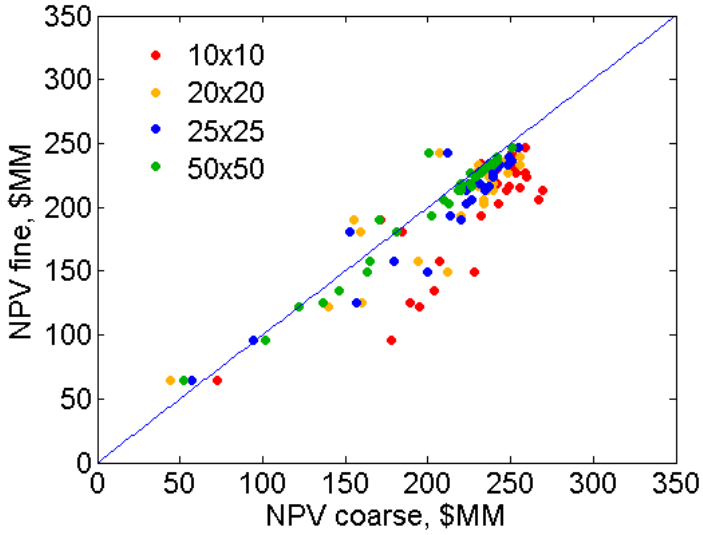

Figure 3.17: Comparison of NPVs evaluated at different grid levels for the 60 well scenarios after 6000 simulation runs (Case 3)

the NPVs for the fine and coarse models for the 60 field development scenarios after 6000 simulation runs is shown in Figure 3.17. The accuracy of the coarse models again increases with refinement, and the general ordering of NPVs, for well scenarios involving different numbers and types of wells, is generally preserved. Thus, coarse models can again be used to replace expensive fine-scale models in the optimization.

Table 3.8: Multilevel optimization computations (Case 3)

Model size	10×10	20×20	25×25	50×50	100×100
Total time per model, sec	4	6	12	24	110
Number of function evaluations	6000	4200	2100	900	300
Elapsed time (fully parallelized), hours	0.11	0.12	0.12	0.1	0.15

Table 3.9: Optimization results for three runs. Best result shown in bold (Case 3)

	Run 1 ($MM)	Run 2 ($MM)	Run 3 ($MM)	Average NPV ($MM)	Time (hours)
Conventional opt.	**332.4**	302.1	275.6	303.4	6.9
Multilevel opt.	273.9	**329.0**	297.2	300.0	0.6

The computational requirements for this case are shown in Table 3.8. Again, most of the flow simulations are performed with the coarser models. Optimization results for three runs are shown in Table 3.9. In this case, the conventional approach provides both a higher NPV and the best individual run. However, for both quantities, the multilevel procedure provides results within about 1%. We again observe an order of magnitude speedup through the use of the multilevel procedure.

The evolution of the objective function for the conventional and multilevel optimizations is shown in Figures 3.18(a) and (b). The conventional approach does not improve the objective function value after about 7000 simulation runs. If we stopped the conventional optimization after 7000 simulation runs, then the speedup from the multilevel procedure would be about a factor of six. In Figure 3.18(b), the solid lines again show the improvement of the objective function value at each grid level, while the black points show the current best well scenario evaluated on the 100×100 grid. In this figure, we see that as the optimization proceeds, the NPV displays significant increase at both the 10×10 and 20×20 levels. This is due to changes in the number of wells, as we now illustrate.

The optimum well locations obtained at the end of each optimization level are shown in Figure 3.19. We begin the optimization with five wells (two injection and three production wells). After the Level 1 optimization, we see that the number of wells

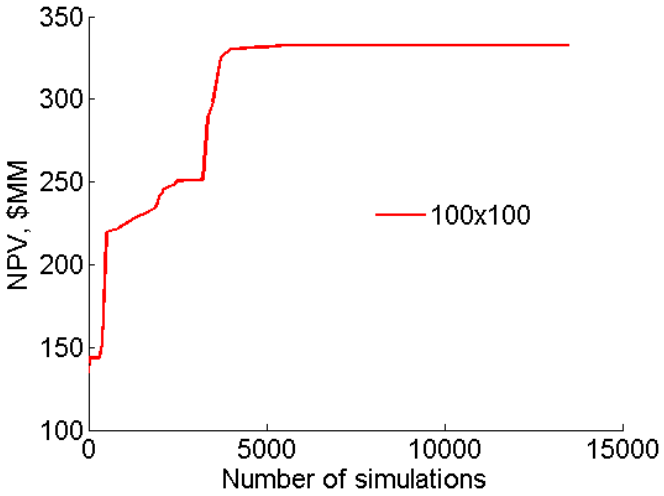

(a) Conventional approach

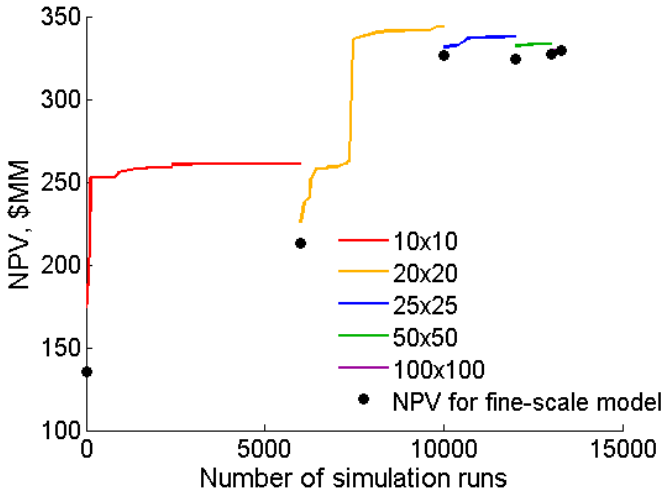

(b) Multilevel approach

Figure 3.18: Evolution of objective function (Case 3)

has decreased to three. The number of optimum wells is further reduced to two after the Level 2 optimization. The number of wells stays at two for the rest of the optimization. This indicates the ability of the coarser models to find the optimal number of wells, and is consistent with expectations given the very high cost of wells in this case.

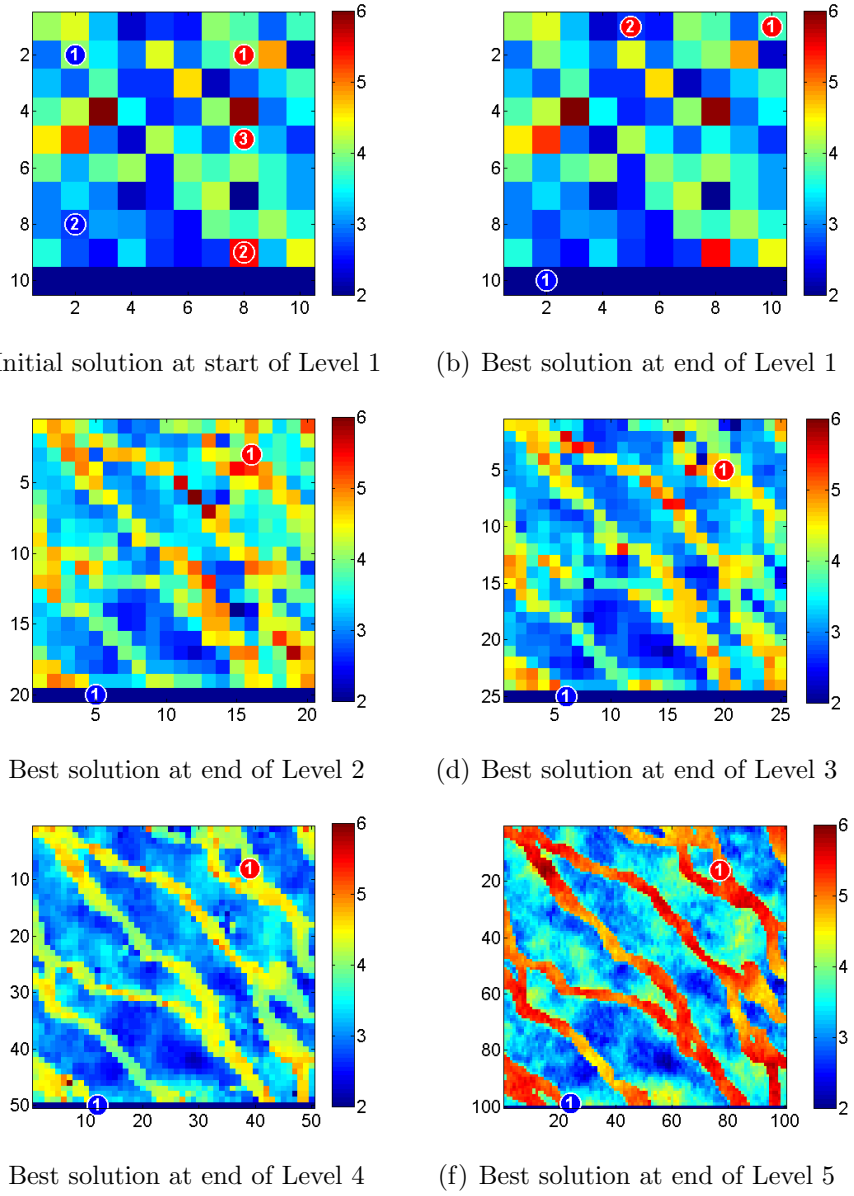

(a) Initial solution at start of Level 1

(b) Best solution at end of Level 1

(c) Best solution at end of Level 2

(d) Best solution at end of Level 3

(e) Best solution at end of Level 4

(f) Best solution at end of Level 5

Figure 3.19: Evolution of optimum well locations in multilevel procedure (Case 3). Background shows $\log_{10} T_y^*$

The best well locations found by the conventional and multilevel methods are compared in Figure 3.20. Both approaches find the optimal number of wells to be two, with one an injector and one a producer. The well locations are actually quite similar, though the well types are opposites. The BHP profiles for the two wells are shown

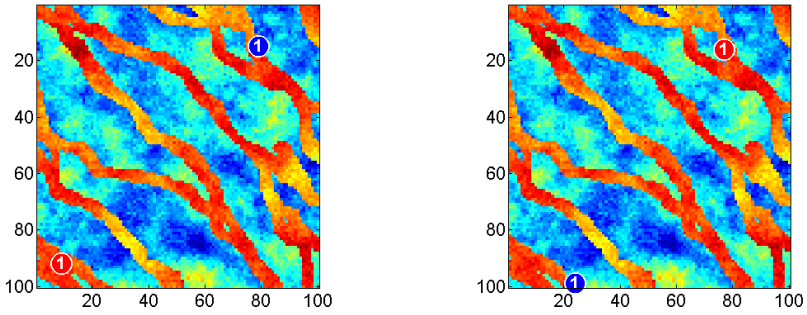

(a) Conventional approach (b) Multilevel approach

Figure 3.20: Best solutions found by the two methods (Case 3). Background shows $\log_{10} k$

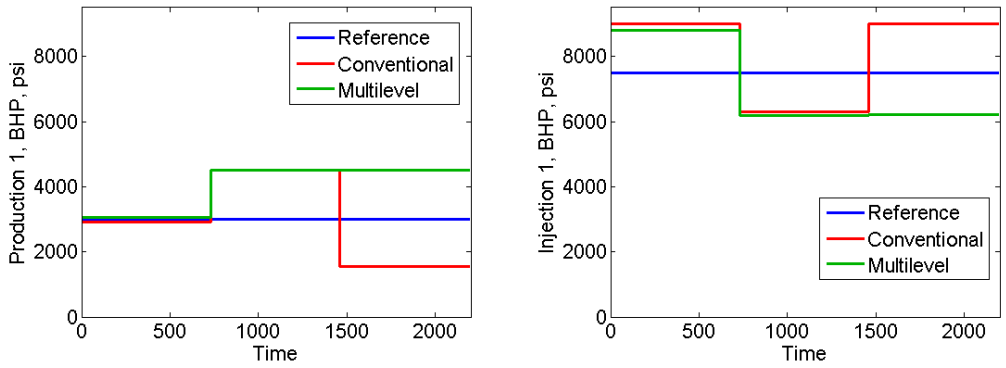

(a) Production well P1 (b) Injection well I1

Figure 3.21: Optimum BHPs found by the two methods (Case 3)

in Figure 3.21. The optimized solutions provide similar BHP values except at the last time period. The final saturation maps of these optimized solutions are shown in Figure 3.22. The areal sweep associated with both solutions is very similar even though the wells are located in opposite positions.

Cumulative production and injection profiles for the optimized and reference (initial guess) solutions are shown in Figure 3.23. The optimized solutions provide almost identical production and injection profiles (the optimal NPVs for the two cases differ by only about 1%). The reference solution produces slightly less oil, while injecting and producing much more water, than the optimized cases.

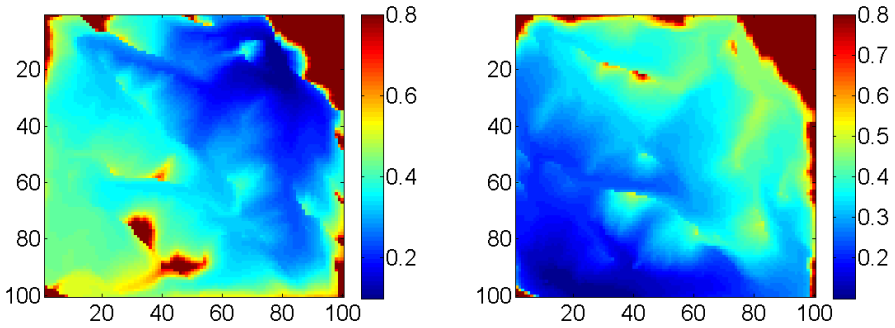

(a) Conventional approach (b) Multilevel approach

Figure 3.22: Comparison of the final oil saturation maps (red indicates oil and blue water), for the optimized solutions found by the two methods (Case 3)

3.4 Summary

In this chapter, we applied the multilevel optimization procedure for three different field development problems involving a single realization. We showed that, although the coarse models provided objective function values that differed from those for the fine models, the ordering of the objective functions for the various well scenarios was essentially preserved. Replacing fine models with coarse models in the multilevel optimization procedure thus provided comparable optimal objective function values to those from the conventional (single-level) approach. Specifically, optimization results for both average optimal NPV over three runs, and best NPV obtained in any run, were consistently within about 1% for the two approaches. We achieved a factor of 5–10 speedup using the multilevel optimization procedure.

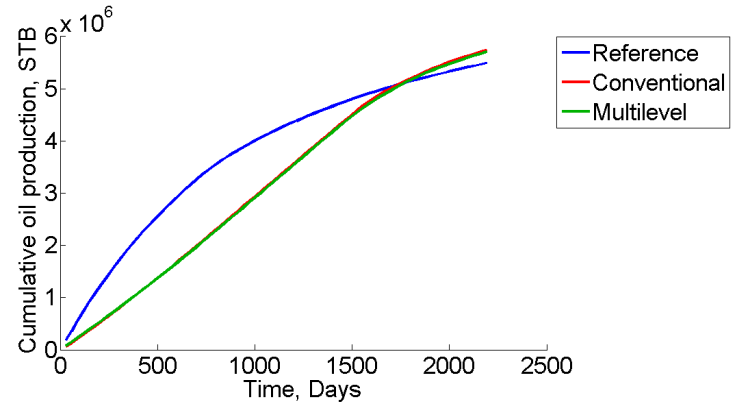

(a) Cumulative oil production

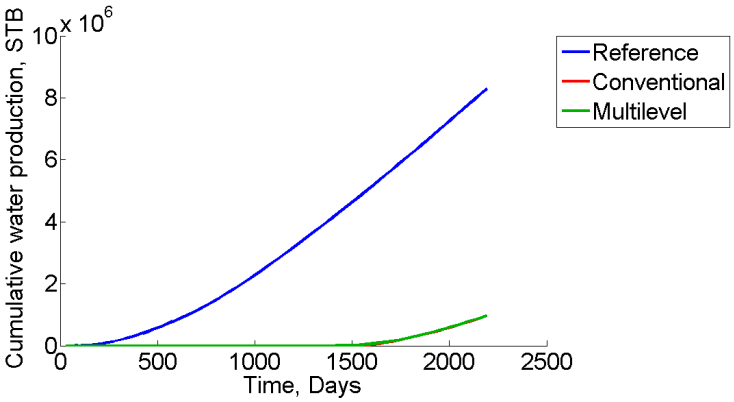

(b) Cumulative water production

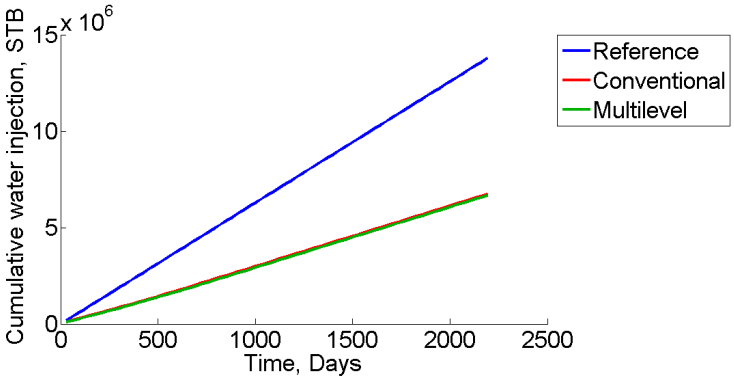

(c) Cumulative water injection

Figure 3.23: Comparison of cumulative production and injection profiles for the optimized and reference solutions (Case 3)

Chapter 4

Optimization under Geological Uncertainty

In reservoir modeling and optimization studies, multiple geological realizations are generated to represent subsurface geological uncertainty. Although the use of multiple realizations in optimization substantially increases the computational demand, the optimum solution found by considering multiple realizations is expected to be more robust than that found using a single realization. In the previous chapter, we have shown that the multilevel optimization procedure can reduce the computational expense by up to an order of magnitude. However, even using the multilevel procedure, optimization under geological uncertainty can still be very expensive. In this chapter, we apply modified versions of the multilevel optimization procedure that lead to reduced computational expense for optimization under geological uncertainty.

In Chapter 2, we described two modifications to improve the efficiency of the multilevel optimization procedure. The first modified approach is referred to as the 'accelerated' multilevel procedure. In this accelerated procedure, we apply only the local MADS optimizer at later stages of the multilevel procedure. The second modification is the incorporation of the 'optimization with sample validation' (OSV) procedure into the accelerated multilevel optimization framework.

We consider three example cases in this chapter, each of which involves 100 realizations intended to capture geological uncertainty. The first example is a two-dimensional model that is not conditioned to any hard data and does not have any existing wells. The second and third examples are two- and three-dimensional systems that are conditioned to hard data and have existing wells.

We divide each example into two sets of runs. In the first set, we use ten geological realizations to compare the conventional, multilevel, and accelerated multilevel optimization procedures. In the second set, we compare the accelerated multilevel procedure using all 100 realizations to the accelerated multilevel procedure with the OSV approach. We do not run the conventional and (standard) multilevel procedures in the second set of runs because, with 100 realizations, the optimization process is very expensive with these approaches. In addition, in the first set of runs we will show that the accelerated multilevel procedure provides results close to those from the conventional and standard multilevel procedures. We additionally present optimization results using MLMC for the first example. This enables comparison of MLMC to multilevel optimization procedures.

The PSO–MADS hybrid algorithm [31, 32] is used in the conventional (single-level) procedure and at all levels of the multilevel optimization approach. In the accelerated multilevel optimization, standalone MADS is used at all levels except at the first level, in which PSO–MADS is used. Optimization with sample validation (OSV) [50], described in Chapter 2, is employed to reduce the number of geological realizations in the accelerated multilevel optimization framework.

We use Stanford's Automatic Differentiation-based General Purpose Research Simulator (AD-GPRS) [63] for all simulations in this chapter. Tables 4.1 and 4.2 present the simulation and optimization parameters used in this study.

4.1 Case 1: Two-dimensional channelized model

This model is taken from Isebor and Durlofsky [29]. The fine grid in this case is 100×100. The 100 realizations are not conditioned to any hard data, though all

Table 4.1: Simulation parameters used in the example cases

Parameter	Value
Grid cell dimensions	32.8 ft × 32.8 ft × 32.8 ft
Initial pressure p_i, at datum	2900 psi at 3280 ft
Rock compressibility c_R	0.5 $\times 10^{-5}$ psi^{-1}
μ_o and μ_w at p_i	3 cp and 1 cp
ρ_o and ρ_w	49.6 and 62.4 lbm/ft^3
B_o and B_w at p_i	1.00 RB/STB (both)

Table 4.2: Optimization parameters used in the example cases

Parameter	Value
P_{po}, P_{pw} and P_{iw}	\$100, \$5 and \$5/STB
Well drilling cost	\$5 million per well
Injection BHP range	6000 – 9000 psi
Production BHP range	1500 – 4500 psi

realizations correspond to a channelized system with channels oriented roughly in the x-direction. We optimize the locations of three production wells and two injection wells. A constant BHP is also determined for each well. We additionally specify nonlinear constraints in the form of well minimum oil ($q_{o,min}$) and maximum water production rate constraints of 62.9 STB/day and 2000 STB/day, respectively. If either constraint is violated by a particular well, we shut the well in until the end of the simulation. We then compute the penalty R_s in Eq. 2.6 using P_{pen} = \$110/STB (note that P_{po} = \$100/STB). The value of Q_{unmet} in Eq. 2.6 is simply (62.9 STB/day)$\times t_{shutin}$, where t_{shutin} is the total number of days during which the well is shut in. The total simulation time is 3650 days. There are 15 optimization variables for this case, and we again take $N_p = 30$. The maximum number of available cores is 30 in this example.

4.1.1 Case 1a: Optimization with ten geological models

In the first set of runs, we use ten realizations and compare the conventional (single-level) approach to the multilevel and accelerated multilevel optimization approaches. These ten realizations are shown in Figure 4.1. For each level of the multilevel optimization algorithm, the maximum number of simulation runs is defined. Optimization at each level is terminated when the maximum number of simulation runs is reached.

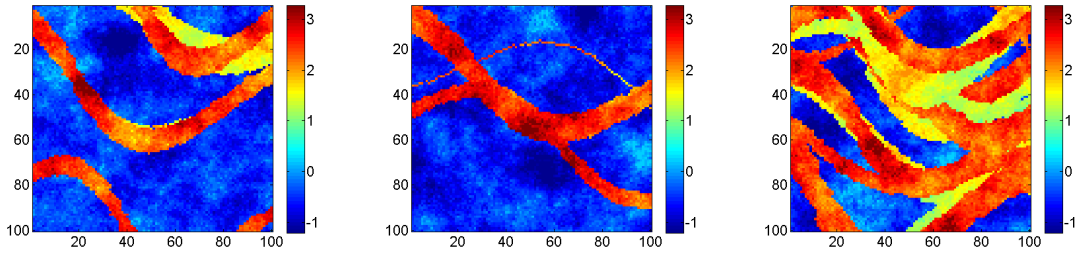

(a) Realization 1 (b) Realization 2 (c) Realization 3

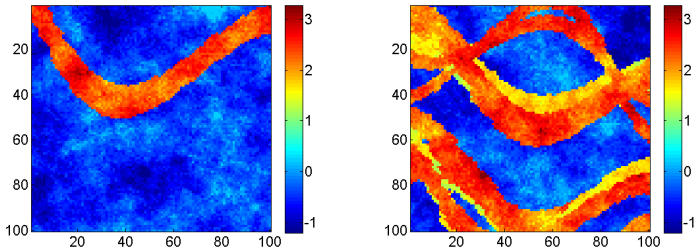

(d) Realization 4 (e) Realization 5

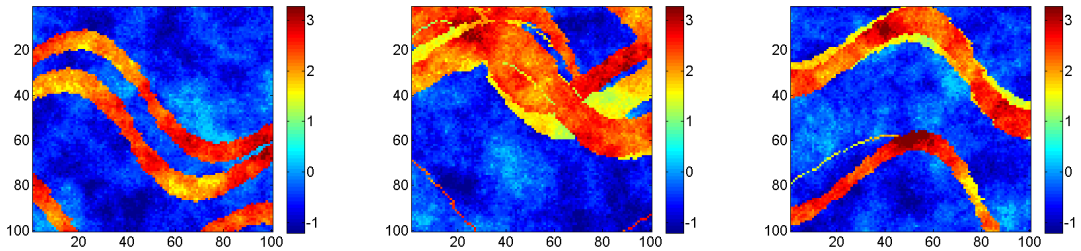

(f) Realization 6 (g) Realization 7 (h) Realization 8

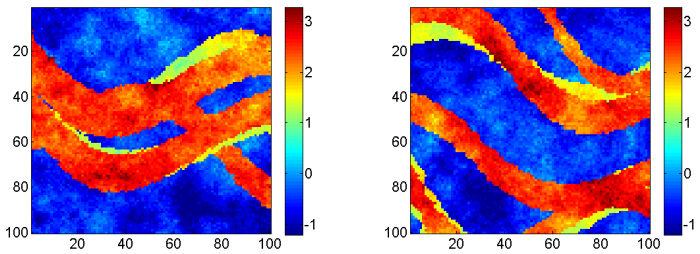

(i) Realization 9 (j) Realization 10

Figure 4.1: Ten geological realizations used for Case 1a. Log_{10} permeability (in md) is shown. Model from [29]

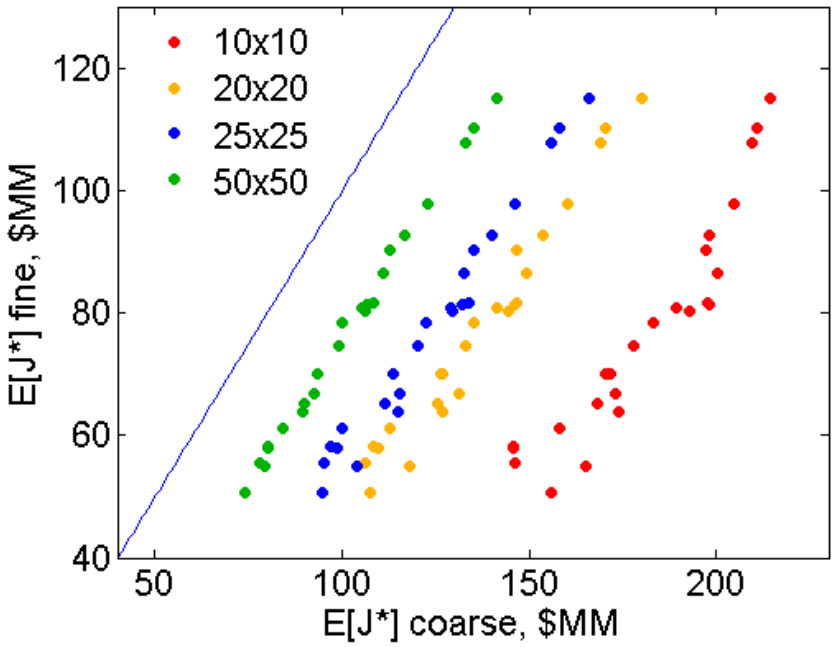

Figure 4.2: Comparison of expected objective function values over 10 realizations evaluated at different grid levels for the 30 candidate well scenarios after 3600 function evaluations using the multilevel approach (Case 1a)

For the conventional (single-level) approach, the number of simulation runs is specified to be the sum of the maximum numbers of simulation runs at all levels in the multilevel optimization. Note that a single function evaluation in Case 1a requires ten flow simulations (one for each realization).

We first assess the accuracy of the upscaled models for this case. Figure 4.2 displays the objective function values of 30 well scenarios (after 3600 function evaluations in the optimization run) for different coarsening levels. The objective function plotted in Figure 4.2 is computed using Eq. 2.4. Note that the x and y scales differ for clarity ($y = x$ line is shown). As in the single-realization cross-plots in the previous chapter, we see that the ordering of the objective function values is very well preserved, and the best well scenarios for coarse models correspond to the best well scenarios for the actual models. This indicates that, for optimization over multiple realizations, coarse models can again be used to evaluate the relative performance of the various

scenarios considered in the optimization. It is again apparent that the offset (error) decreases systematically as the grid is refined.

Table 4.3: Multilevel optimization computations (Case 1a)

Model size	10×10	20×20	25×25	50×50	100×100
Total time per model, sec	6	9	14	28	120
Number of function evaluations (PSO–MADS)	3600	2400	1200	600	180
Elapsed time, hours (using 30 nodes)	2.0	2.0	1.6	1.6	2.0

Table 4.4: Accelerated multilevel optimization computations (Case 1a)

Model size	10×10	20×20	25×25	50×50	100×100
Total time per model, sec	6	9	14	28	120
Number of function evaluations (PSO–MADS/MADS)	3600	810	420	210	60
Elapsed time, hours (using 30 nodes)	2.0	0.7	0.5	0.5	0.7

Table 4.5: Optimization results for three runs. Best result shown in bold (Case 1a)

	Run 1 ($MM)	Run 2 ($MM)	Run 3 ($MM)	Average ($MM)	Time (hours)
Conventional opt.	170.2	**186.3**	164.6	173.7	88.7
Multilevel opt.	**186.5**	159.2	162.6	169.4	9.2
Accelerated multilevel opt.	**183.3**	159.2	161.2	167.9	4.4

Tables 4.3 and 4.4 show the performance of the multilevel and the accelerated multilevel procedures. In these tables, the total computation time per model includes the time required to upscale a fine-scale realization and run the flow simulation. The number of simulation runs and the elapsed time for each level of the multilevel procedures are also presented. Recall that the difference between the multilevel and

accelerated multilevel procedures is that we use MADS instead of PSO–MADS after the first optimization level in the accelerated procedure. Therefore, the number of function evaluations in the accelerated multilevel optimization (5100) is less than in the multilevel optimization (7980). For the conventional (single-level) approach, the total number of function evaluations is 7980, which is equal to the total number of function evaluations in the (standard) multilevel optimization. Note that the accelerated procedure requires about one-third the number of function evaluations at Levels 2–5 of that required by the standard multilevel approach. This leads to significant savings in elapsed time.

Optimization results for three runs using all procedures are presented in Table 4.5. We see that all three approaches achieve fairly similar results. The conventional approach provides the best average result, though the best individual run is obtained by the multilevel approach. The accelerated multilevel procedure is about two times faster than the multilevel approach and about 20 times faster than the conventional approach. Optimization results using this procedure are about 3% less on average than those with the conventional approach.

The improvement in the objective function for the best solution over the course of the optimization for the conventional, multilevel and accelerated multilevel approaches is shown in Figures 4.3(a)-(c). In the conventional approach, the objective function improves very little after 40,000 simulation runs. This indicates that, with an appropriate termination criterion, this optimization might require only 40,000 or 50,000 simulations. In this case, the timing reported in Table 4.5 would reduce to about 45 hours (instead of 88.7 hours), and the speedups for the multilevel and accelerated multilevel procedures would be about 5 and 10 (rather than 10 and 20).

As in Chapter 3, for the multilevel and accelerated multilevel approaches, the solid lines in Figures 4.3(b) and (c) indicate the objective function values evaluated using coarse models and the black points indicate results for the actual (fine-scale) models. Although there are errors in the objective function values of the coarse models (as is evident in Figure 4.2), the optimizer still improves the actual objective function value. We also see that after the first level, the number of simulation runs is reduced significantly with the accelerated multilevel procedure. Since flow simulations are

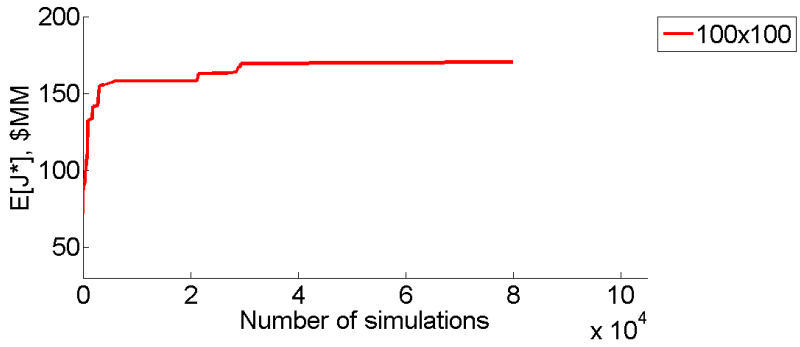

(a) Conventional approach

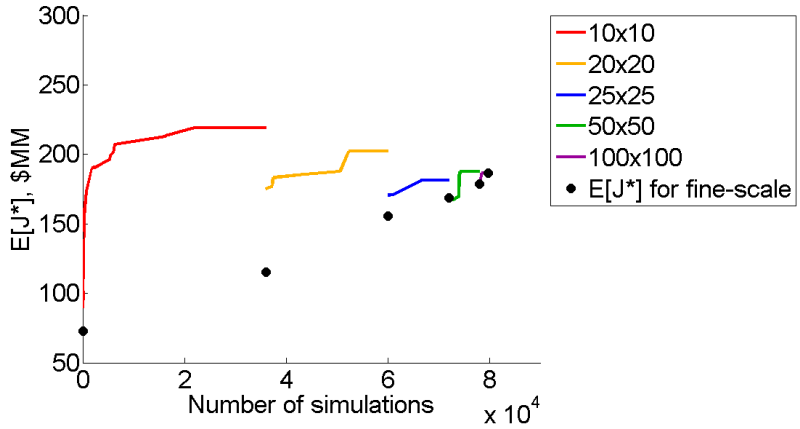

(b) Multilevel approach

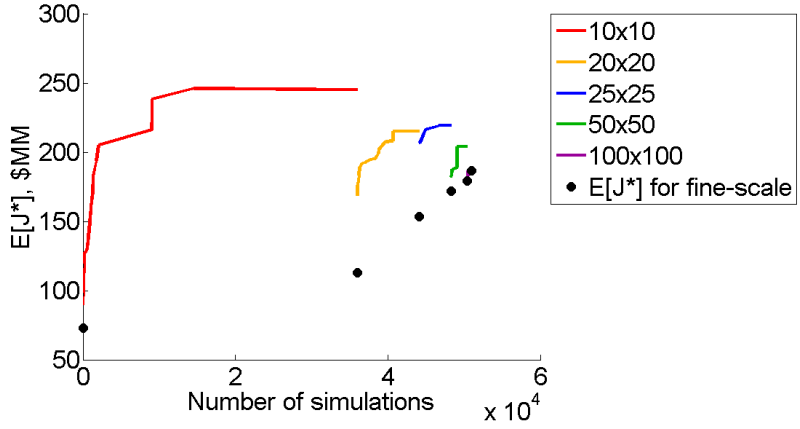

(c) Accelerated multilevel approach

Figure 4.3: Evolution of objective function (Case 1a)

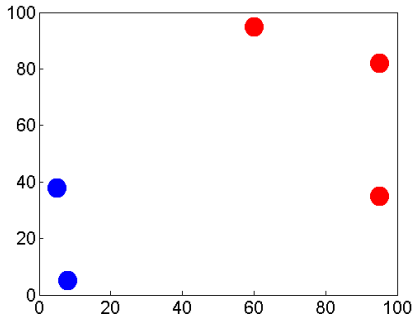

(a) Best solution using conventional approach

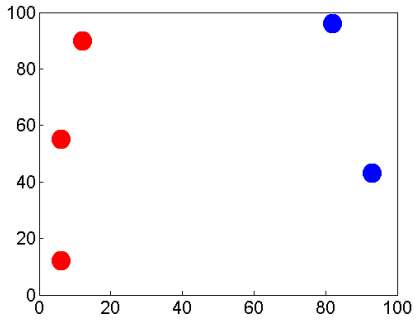

(b) Best solution using multilevel approach

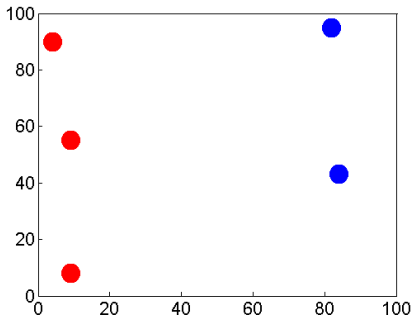

(c) Best solution using accelerated multilevel approach

Figure 4.4: Best solutions found by the three methods. Injection and production wells are shown as blue and red circles respectively (Case 1a)

more expensive in later levels, this reduces considerably the cost of the optimization. The optimum well locations obtained by the conventional, multilevel and accelerated multilevel approaches are shown in Figure 4.4. It is evident that the multilevel and accelerated multilevel procedures provide similar well configurations. Figure 4.5 shows how the well locations change as the optimization proceeds in the (standard) multilevel approach. The well locations change less after the Level 1 optimization, which motivates the switch, in the accelerated procedure, from the global optimizer (PSO–MADS) to the local optimizer (MADS) after Level 1. This approach is effective because PSO has relatively little effect on the improvement of the objective function after the Level 1 optimization is completed.

The optimized (constant) BHPs for the production well P1 and injection well I1 are shown in Figures 4.6(a) and (b). The reference value is the initial guess BHP used

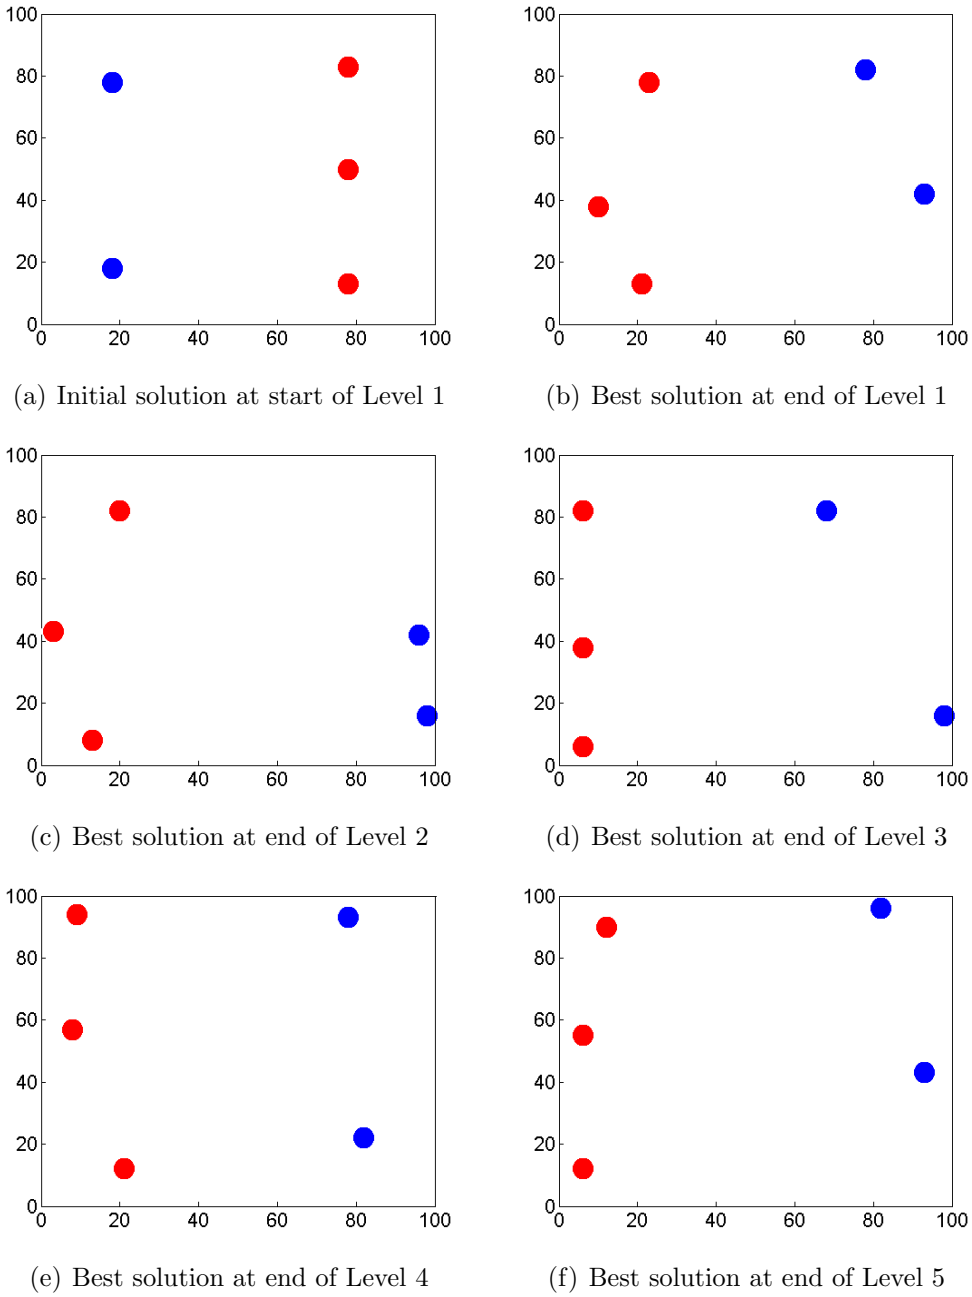

(a) Initial solution at start of Level 1

(b) Best solution at end of Level 1

(c) Best solution at end of Level 2

(d) Best solution at end of Level 3

(e) Best solution at end of Level 4

(f) Best solution at end of Level 5

Figure 4.5: Evolution of optimum well locations in multilevel (PSO–MADS at all levels) procedure. Injection and production wells are shown as blue and red circles respectively (Case 1a)

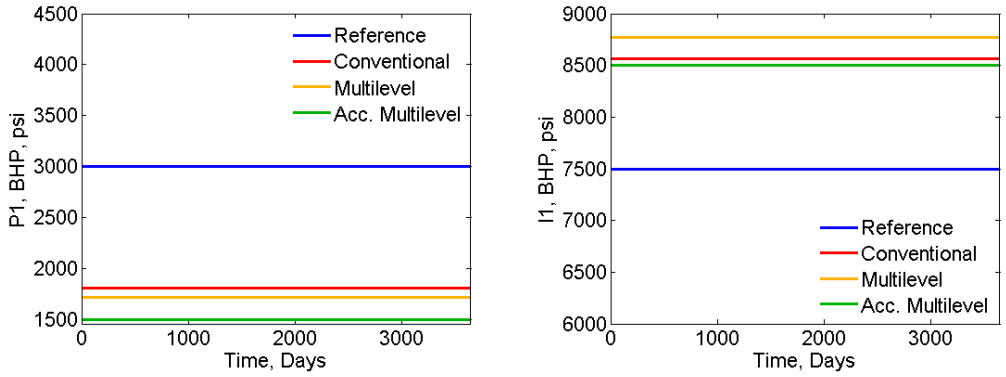

(a) Production well P1 (b) Injection well I1

Figure 4.6: Optimum BHPs found by the three methods (Case 1a)

in the optimization. It is evident that the optimization shifts the BHPs towards the bounds (1500 psi for P1, 9000 psi for I1). The optimized BHPs for the other production and injection wells display similar results.

Cumulative production and injection profiles, for the optimized solutions obtained by the conventional, multilevel and accelerated multilevel approaches, are compared to the reference solution in Figures 4.7(a)-(c). The well locations and controls used in the reference solution are shown in Figures 4.5(a), 4.6(a) and 4.6(b). In Figures 4.7(a)-(c), the thin broken lines indicate the simulation results for each of the ten realizations, and the thick solid line is the mean of the results for the ten realizations. It is evident that the three optimized solutions provide significantly larger ($\sim$ 1 MMSTB) expected cumulative oil recovery than the reference scenario. The optimized solutions all display similar expected cumulative oil recoveries. The optimized expected cumulative water injection in Figure 4.7(c) exceeds that of the reference case. This is because, given the water production and injection costs ($5/STB), additional water use is beneficial in terms of NPV. Recall that the nonlinear constraint violation term in these optimizations enters as a penalty function, so these constraints may be violated in some realizations.

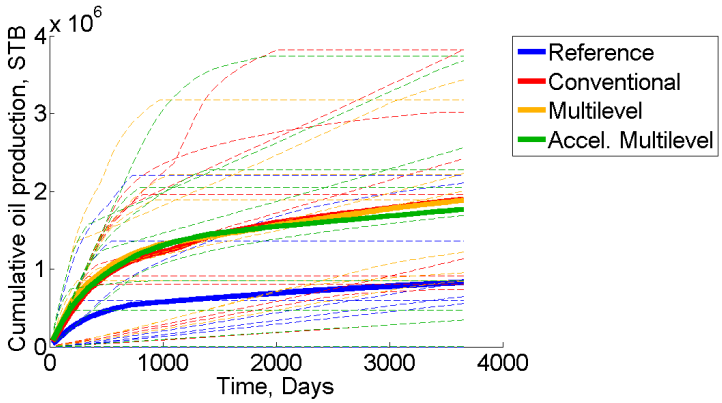

(a) Cumulative oil production

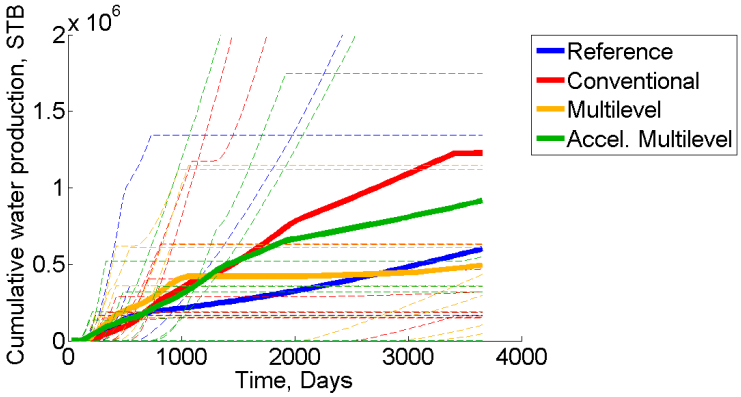

(b) Cumulative water production

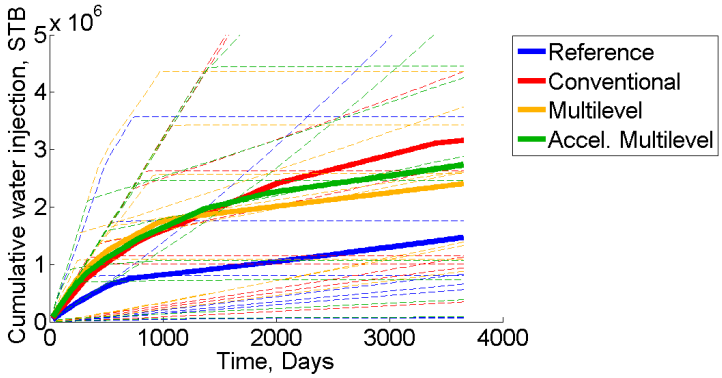

(c) Cumulative water injection

Figure 4.7: Cumulative production and injection profiles for the optimized solutions obtained by the conventional, multilevel and accelerated multilevel approaches (Case 1a)

Table 4.6: Accelerated multilevel optimization computations (Case 1b)

Model size	10×10	20×20	25×25	50×50	100×100
Total time per model, sec	6	9	14	28	120
Number of function evaluations (PSO–MADS/MADS)	1800	420	210	120	60
Elapsed time using all 100 models, hours (using 30 nodes)	10	3.5	2.7	3.1	6.7

4.1.2 Case 1b: Optimization with 100 geological models

We now consider the full set of 100 geological realizations in the optimization. The problem set up is otherwise identical to that of Case 1a. In Case 1a we showed that the accelerated multilevel optimization provided relatively close results to those from the conventional and multilevel approaches with ten geological realizations. Therefore, we will not run the conventional and multilevel approaches with 100 geological models as these would be quite expensive (especially the conventional single-level procedure).

Here, we compare the accelerated multilevel approach using all 100 realizations to the accelerated multilevel approach with OSV. We also run optimization with the MLMC approach. The total computation time per model, number of simulation runs, and elapsed time for the accelerated multilevel optimization using all geological realizations are shown in Table 4.6. Using the OSV approach, we specify the number of realizations for each problem to be either 10, 30, 50 or 100. In other words, optimization is first performed over $N_{rep} = 10$ realizations, and if the validation criterion ($RI > 0.5$) is not satisfied, then N_{rep} is increased to 30, etc. The number of representative realizations used in each subproblem at each level is shown in Table 4.7. We see that an RI of 0.5 or greater is achieved at every level. The time required at each level using OSV is shown in Table 4.8. For both the accelerated multilevel approach using all 100 geological realizations and the OSV approach, 30 cores are used for parallelization.

Table 4.7: Number of representative realizations (determined using OSV) and the corresponding relative improvement values at each level for the best run (Case 1b)

Level	N_{rep}	RI
10×10	10	0.18
10×10	30	0.77
20×20	10	0.13
20×20	30	0.55
25×25	10	0.63
50×50	10	0.34
50×50	30	0.95
100×100	10	-0.33
100×100	30	0.10
100×100	50	0.78

We now analyze the accuracy of the objective function, $E[J^*]$, using different numbers of representative models at different levels of resolution. Results are shown in Figure 4.8. To generate these plots, we first simulate all 100 realizations using 10×10 models. The upscaling required to generate these 10×10 models is performed using the best scenario after 1800 function evaluations have been performed by the optimizer (using the accelerated multilevel with OSV method). The CDF of the objective function is then constructed. Next, we select different numbers of representative models (10, 30, 50, 100) from this CDF, as described in Section 2.4. Finally, we compare the expected values for the representative models to the expected values for all 100 fine-scale models for various well scenarios (the 30 candidate scenarios in the optimization after 1800 function evaluations).

We see in Figure 4.8(d) that when we use all 100 realizations, there are clear trends between $E[J^*]$ for the fine and coarse models at all levels. This is consistent with Figure 4.2. These trends are still evident with 50 representative models (Figure 4.8(c)) but significant scatter appears with 30 and 10 representative models (Figures 4.8(b) and 4.8(a)). Note that the scales differ between these figures and that the y-axis range is fairly narrow. Figure 4.9 shows how the accuracy of the expected objective function value increases as the number of representative models increases for the 100×100

Table 4.8: Accelerated multilevel optimization computations with OSV for the best run (Case 1b)

Model size	10×10	20×20	25×25	50×50	100×100
Total time per model, sec	6	9	14	28	120
Number of function evaluations with 10 models	1800	420	210	120	60
Number of function evaluations with 30 models	1800	420	0	120	60
Number of function evaluations with 50 models	0	0	0	0	60
Number of function evaluations with 100 models	0	0	0	0	0
Elapsed time with OSV, hours (using 30 nodes)	4	1.4	0.3	1.2	6

model. Figures 4.8 and 4.9 highlight the impact of approximations in both model resolution and the representation of uncertainty. For this example, the upscaling errors appear to be more systematic. This should pose less of a problem for the optimizer than the scatter evident when a relatively small number of realizations is used.

Table 4.9: Optimization results for three runs. Best result shown in bold (Case 1b)

	Run 1 ($MM)	Run 2 ($MM)	Run 3 ($MM)	Average ($MM)	Time (hours)
Accelerated multilevel opt. (100 realizations)	**148.2**	143.6	139.7	143.8	26.0
Accelerated multilevel opt. with OSV	**147.8**	136.2	144.0	142.6	12.9
Optimization with MLMC	**146.2**	139.2	141.2	142.2	51.2

Optimization results for three runs using the accelerated multilevel optimization with all geological models and accelerated multilevel optimization with OSV are compared in Table 4.9 (MLMC results will be discussed below). For the runs that use 100 realizations, each function evaluation requires 100 simulations. When OSV is used, the number of simulations per function evaluation varies. We see that the OSV

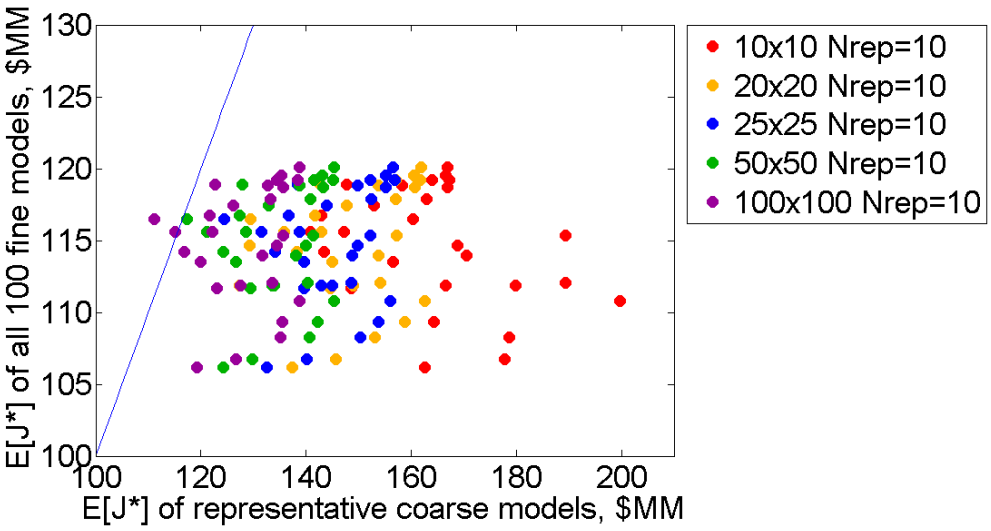

(a) 10 representative models

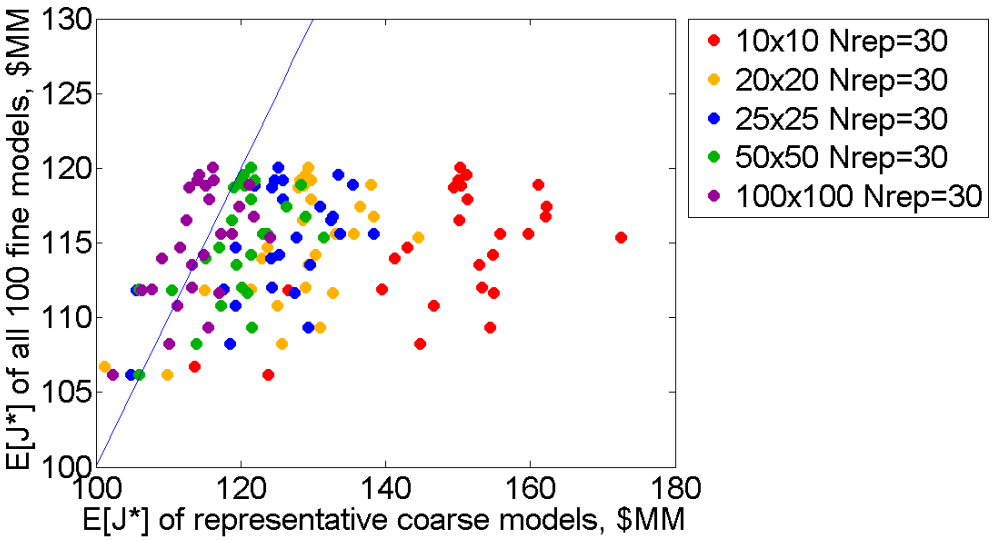

(b) 30 representative models

(Caption on following page)

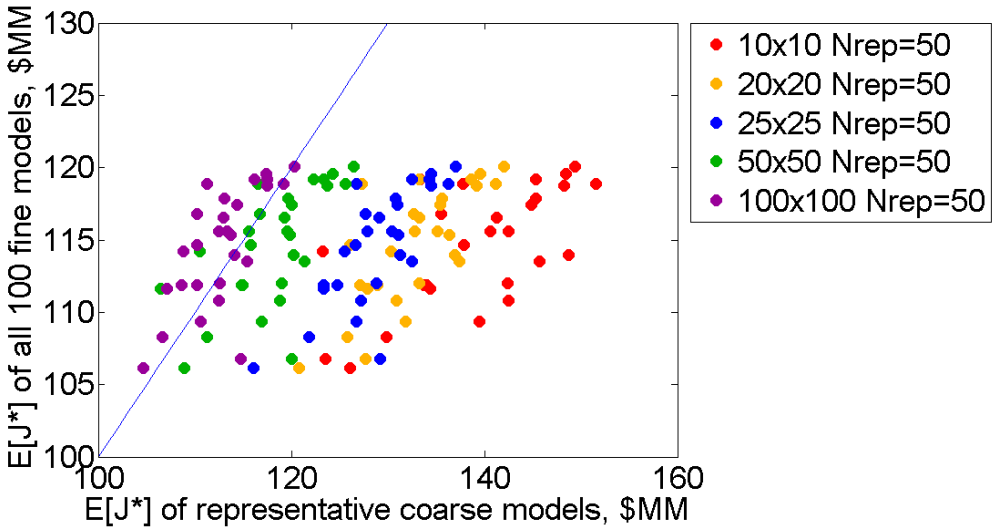

(c) 50 representative models

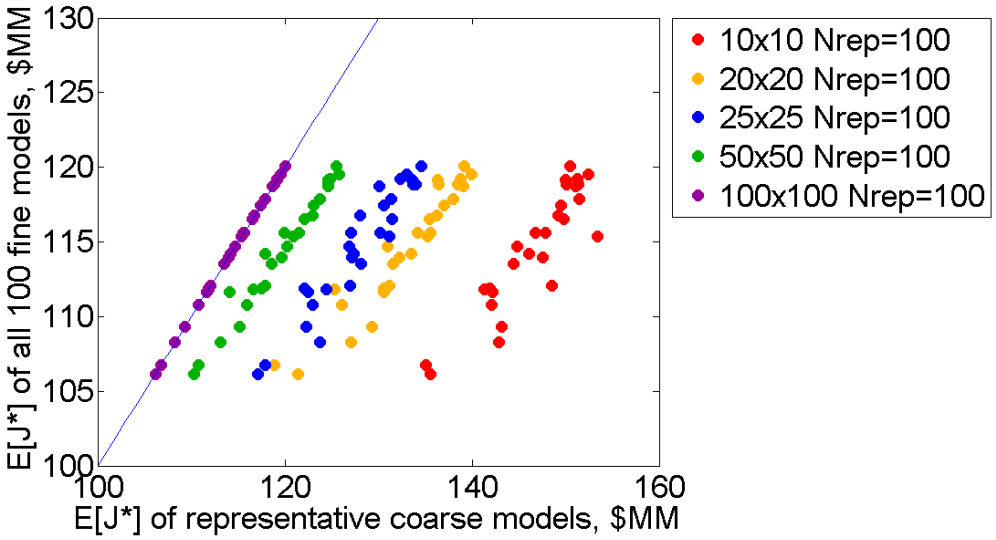

(d) 100 representative models

Figure 4.8: Comparison of expected objective function values evaluated at different grid levels with different numbers of representative models for the 30 candidate well scenarios after 1800 function evaluations using the accelerated multilevel approach with OSV method (Case 1b)

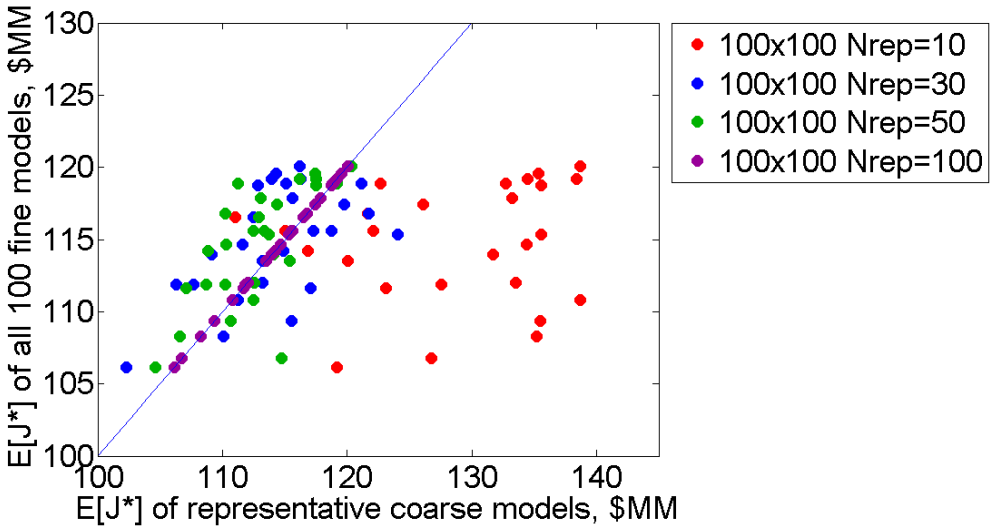

Figure 4.9: Comparison of expected objective function values over 100 realizations evaluated at the fine-grid level with different numbers of representative models for the 30 candidate well scenarios after 1800 function evaluations using the accelerated multilevel approach with OSV method (Case 1b)

procedure provides optimum objective function values, for both the average and the best individual run, that are within 1% of those using all realizations. However, OSV requires only one half of the time needed when all models are used. We estimate that the time required to run the conventional approach, with 100 fine-scale (100×100) models used for every simulation, and a total of 4080 function evaluations (this is the number of function evaluations in the multilevel approach), would be about 453 hours. Relative to this, accelerated multilevel optimization with OSV provides a speedup of about 35. Speedup would be closer to a factor of 20 if only around 2000 function evaluations were required.

Figures 4.10(a) and (b) display the best well locations found by the accelerated multilevel optimization with all geological models and accelerated multilevel optimization with OSV. Both approaches locate injection and production wells on opposite sides of the reservoir. Cumulative production and injection profiles for the optimized and reference solutions are shown in Figures 4.11(a)-(f). The reference scenario is the same as in Case 1a. Although the variation in the cumulative oil production for the

100 realizations (Figure 4.11(a)) is very large, we see a significant improvement in the expected value (Figure 4.11(b)) compared to the reference scenario. Optimized solutions correspond to more water production and injection compared to the reference scenario (Figures 4.11(c)-(f)).

The evolution of the objective function for the accelerated multilevel procedure using all geological models and with the OSV approach is shown in Figure 4.12. In Figure 4.12(b), we see that OSV requires more than one optimization subproblem, with increasing numbers of representative models, for the 10×10, 20×20, 50×50 and 100×100 models. This is because the relative improvement (RI>0.5) criterion is not satisfied using only ten realizations at these grid levels.

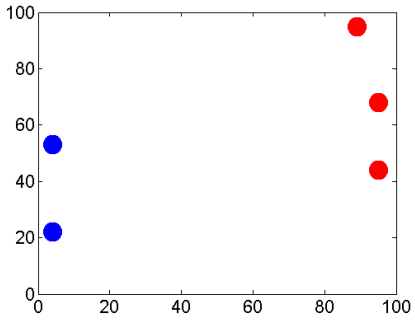

(a) Best solution using accelerated multilevel procedure

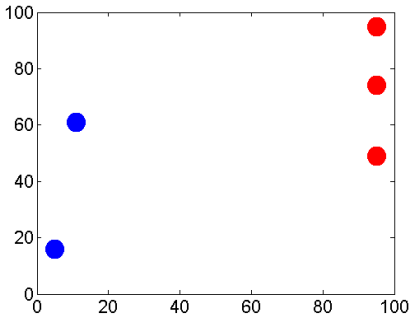

(b) Best solution using accelerated multilevel procedure with OSV

Figure 4.10: Best solutions found by the two methods. Injection and production wells are shown as blue and red circles respectively (Case 1b)

Finally, we consider the use of the multilevel Monte Carlo (MLMC) method for this example. In the MLMC method presented by Müller et al. [40], different subsets of realizations are randomly sampled from the full set of realizations at every coarsening level. In this study, we compare random sampling to the CDF approach [50], discussed in Section 2.4, to select representative models in MLMC. Note that, even though the sampling is no longer random when we use the CDF approach, we still refer to the procedure as multilevel Monte Carlo.

We also compare MLMC with a single and multiple coarsening levels. In MLMC

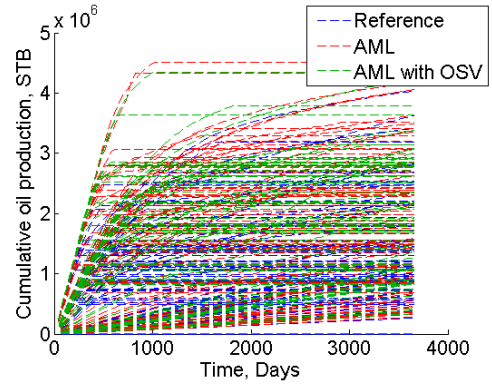

(a) Cumulative oil production

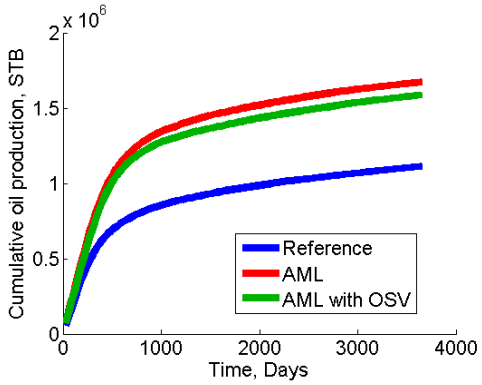

(b) Expected cumulative oil production

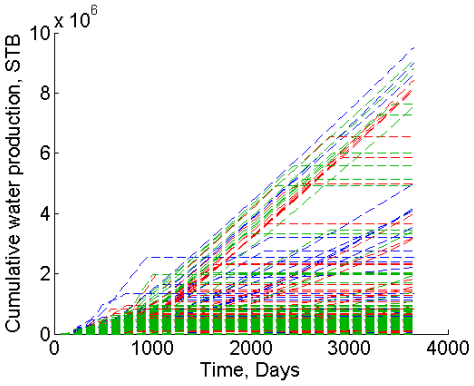

(c) Cumulative water production

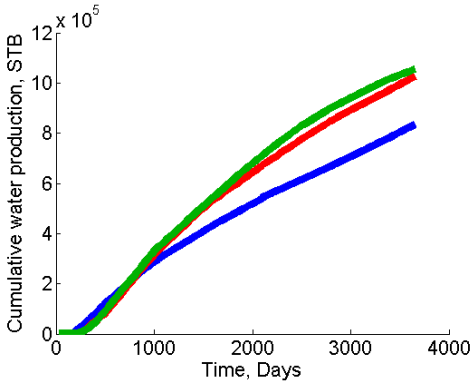

(d) Expected cumulative water production

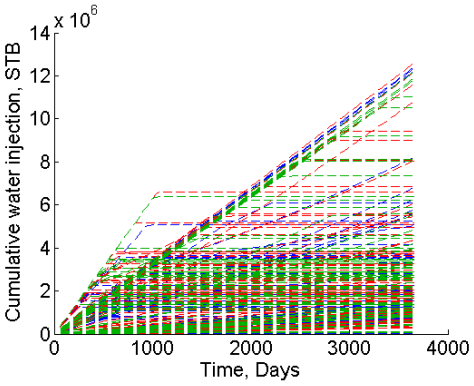

(e) Cumulative water injection

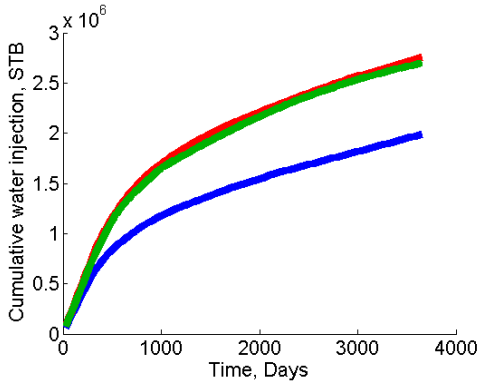

(f) Expected cumulative water injection

Figure 4.11: Cumulative production and injection profiles for the optimized solutions obtained by the accelerated multilevel and accelerated multilevel with OSV approaches. Left plots show results for 100 realizations, right plots show expected values. Note difference in scales (Case 1b)

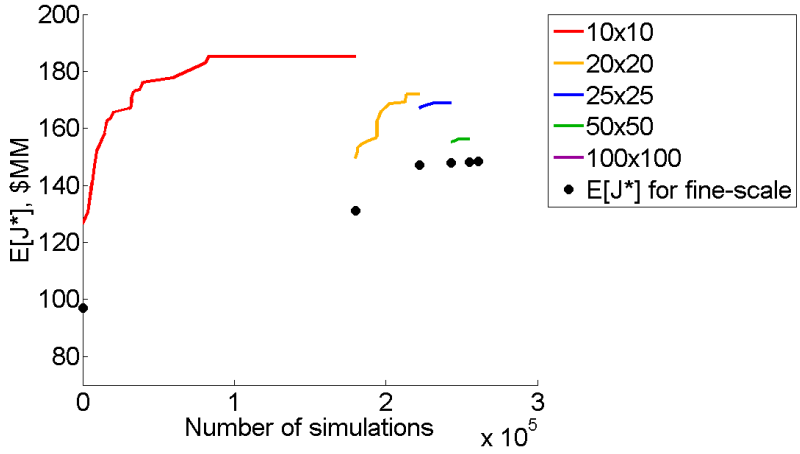

(a) Accelerated multilevel approach with all 100 realizations

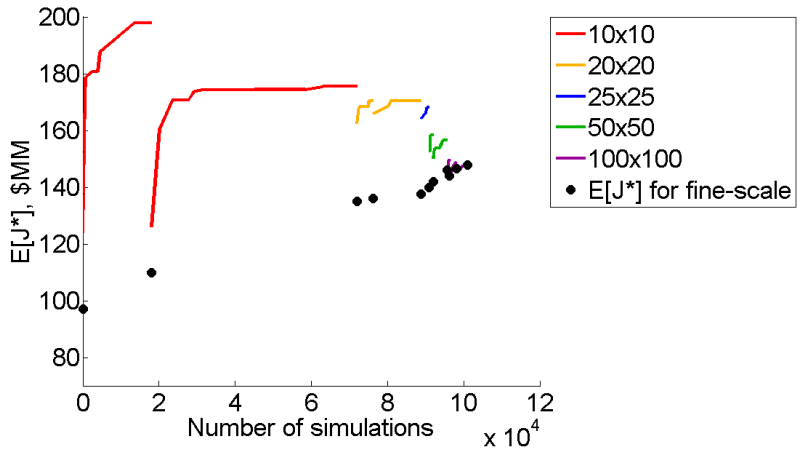

(b) Accelerated multilevel approach with OSV

Figure 4.12: Evolution of objective function (Case 1b)

with multiple coarsening levels, we run all 100 realizations at the 10×10 level, 40 realizations at the 20×20 level, 20 realizations at the 25×25 level, ten realizations at the 50×50 level, and five realizations at the 100×100 level. The expected value of the objective function is then estimated using Eq. 2.14. The total computation time for one function evaluation performed in this way is 2120 seconds, compared to 12,000 seconds in the conventional approach (100 runs with 100×100 models). In MLMC with a single coarsening level, we use 100 realizations at the 10×10 level and five realizations at the 100×100 level. The total time required for a function

evaluation using this treatment is 1200 seconds.

Results using the four combinations of the sampling and coarsening-level treatments are compared to the expected value of the objective function evaluated using all 100 fine models in Figure 4.13. Each point represents one of the 30 well configurations considered at a particular point in the optimization (after 1800 function evaluations). Although MLMC with a single coarsening level is computationally more efficient, MLMC with multiple coarsening levels appears to provide better results (e.g., the blue points show less scatter than the green points). We therefore use MLMC with multiple coarsening levels in subsequent optimizations. In addition, the use of the CDF approach for model selection in MLMC can be seen in Figure 4.13 to provide slightly more accurate results than random selection (e.g., the blue points show less scatter than the red points). Therefore, we use the CDF approach with multiple coarsening levels in our MLMC-based optimizations.

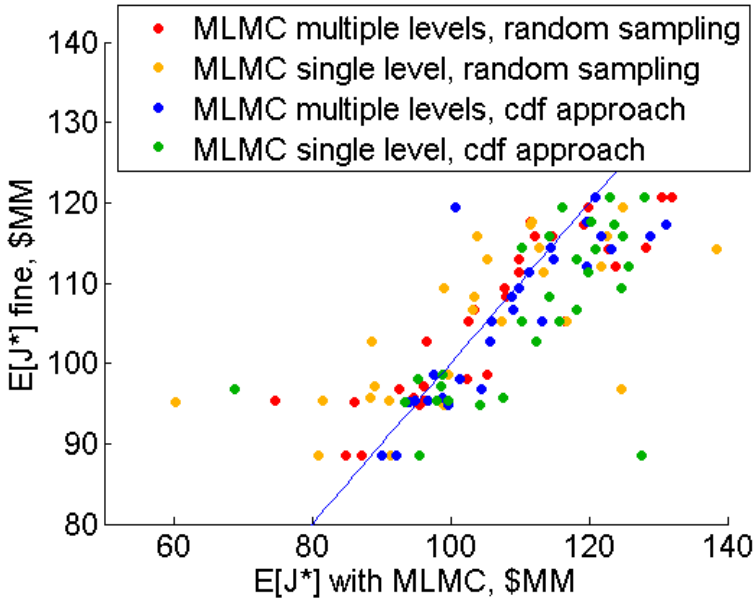

Figure 4.13: Comparison of expected objective function values over 100 realizations evaluated at the fine-grid level with MLMC estimation for the 30 candidate well scenarios after 1800 function evaluations using the MLMC optimization method (Case 1b)

Optimization results using MLMC with multiple coarsening levels and the CDF approach for model selection are shown in Table 4.9 above. The elapsed time required for these runs (2610 function evaluations) is 51.2 hours. This approach provides optimum objective function values that are comparable to those from the two accelerated multilevel optimization procedures. Although optimization with the MLMC approach is more efficient than the conventional approach, it is still more expensive than the accelerated multilevel optimization procedures for this example. It is possible that other MLMC implementations may be more efficient, and this should be explored in future work.

4.2 Case 2: Oriented channel models

In Case 2, we begin with an initial set of 1500 geological realizations of a bimodal channelized system generated by Vo [56]. Ten of these geological realizations are shown in Figure 4.14. All realizations are conditioned to hard data that are available from the nine wells shown as black, red and blue points in Figure 4.15 (the wells in black can be viewed as exploration wells).

Of these nine wells, two production wells and one injection well, shown as red and blue circles in Figure 4.15, are assumed to have been producing for two years. The production history is synthetically generated by taking one of the 1500 models as the 'true model' (realization 1052), which is shown in Figure 4.15, and then performing flow simulation. We then select the 100 geological realizations from the remaining 1499 prior realizations that provide the closest production response to the production history. 'Closeness' is quantified here by the usual least-squares objective function used in [42]. These 100 geological realizations are then used in the optimization.

The optimization problem now involves finding the locations of one production well and one injection well and the BHP control values for all five wells (three existing and two new wells) for three time intervals. The simulation is performed for a total of eight years (two years of which is viewed as production history). The BHP values are updated at the end of the second, fourth and sixth years. We also specify minimum oil and maximum water production rate constraints of 62.9 STB/day and 2000 STB/day,

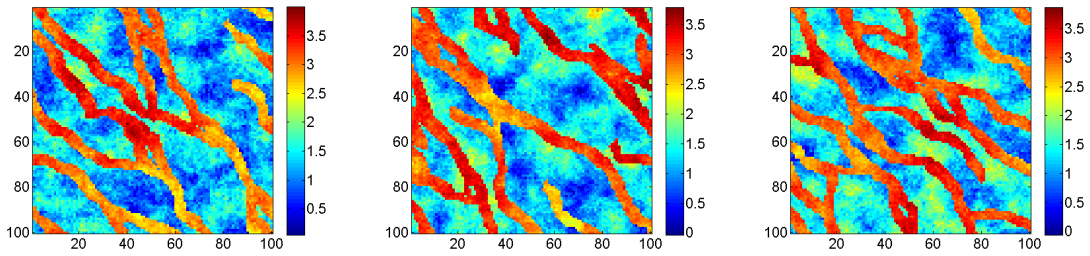

(a) Realization 1 (b) Realization 2 (c) Realization 3

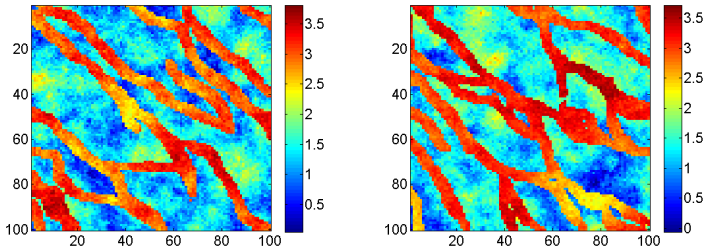

(d) Realization 4 (e) Realization 5

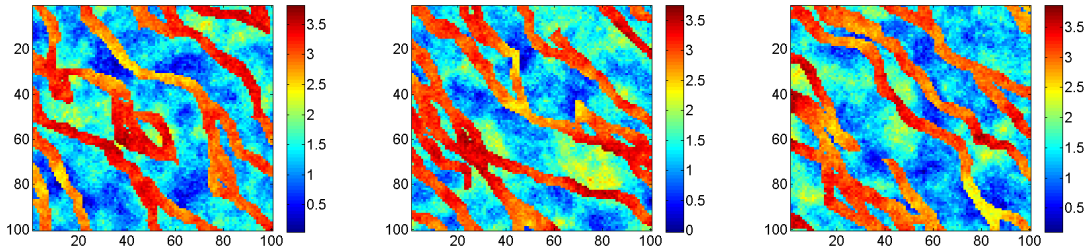

(f) Realization 6 (g) Realization 7 (h) Realization 8

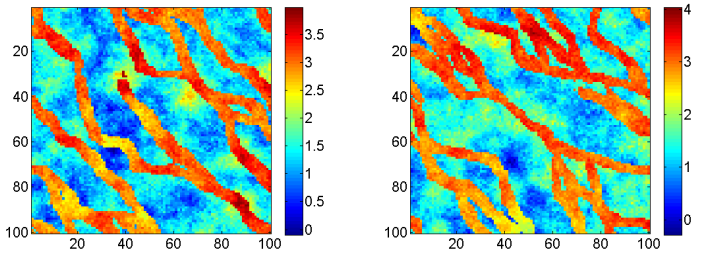

(i) Realization 9 (j) Realization 10

Figure 4.14: Ten geological realizations used for Case 2a. Log_{10} permeability (in md) is shown. Models from [56]

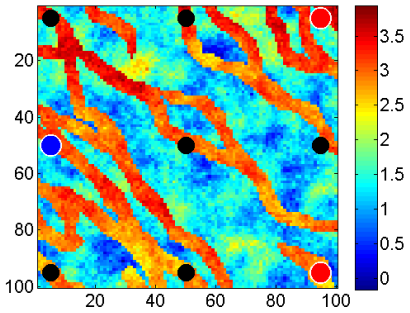

Figure 4.15: Log_{10} permeability (in md) of the true model (realization 1052) for Case 2. Black points indicate hard data (exploration well) locations. Existing production and injection wells shown as red and blue circles, respectively. Model from [56]

respectively. The constraints are again handled with the penalty approach described in Chapter 2.

There are a total of 19 optimization variables, of which four are location and 15 are control variables. We set $N_p = 38$, and the number of simulation runs per MADS iteration is also 38. The optimizations are run with 38 cores.

Table 4.10: Multilevel optimization computations (Case 2a)

Model size	10×10	20×20	25×25	50×50	100×100
Total time per model, sec	4	6	12	24	110
Number of function evaluations (PSO–MADS)	1824	1216	608	380	228
Elapsed time, hours (using 38 nodes)	0.53	0.53	0.53	0.67	1.83

4.2.1 Case 2a: Optimization with ten geological models

In Case 2a, in the conventional (single-level) approach, we simulate ten 100×100 models at every function evaluation. The PSO–MADS algorithm is terminated after 4256 function evaluations. One simulation run using the 100×100 model takes about 110 seconds. Therefore, the total elapsed time required to run the conventional

Table 4.11: Accelerated multilevel optimization computations (Case 2a)

Model size	10×10	20×20	25×25	50×50	100×100
Total time per model, sec	4	6	12	24	110
Number of function evaluations (PSO–MADS/MADS)	1824	456	228	152	76
Elapsed time, hours (using 38 nodes)	0.53	0.20	0.20	0.27	0.61

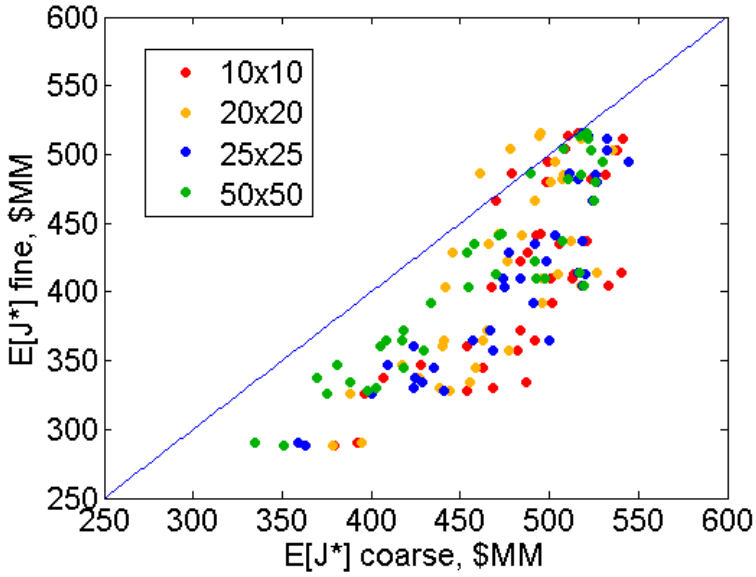

Figure 4.16: Comparison of expected objective function values over 10 realizations evaluated at different grid levels for the 38 candidate well scenarios after 1824 function evaluations using the accelerated multilevel approach (Case 2a)

approach with 38 cores is about 34 hours. In the multilevel approach, most of the 4256 function evaluations are performed using upscaled models. The total computation time per model, number of function evaluations per level, and elapsed time are shown in Table 4.10. The total elapsed time required for the multilevel approach is about 4.1 hours. We also run the accelerated multilevel approach, with standalone MADS in the later levels. The elapsed time required for the accelerated multilevel approach is 1.8 hours, as indicated in Table 4.11.

The objective function values for 38 different well configurations, after 1824 function evaluations using the accelerated multilevel approach, are shown in Figure 4.16. The objective function value for each scenario is the expected value over the ten geological realizations. Results for different coarsening levels are shown. The cross-plot for this case displays more scatter than in the previous example (Figure 4.2). This might be due to the large number of thin high-permeability channels that characterize the models in this case. Such systems can be more challenging to upscale accurately. The upscaling is also complicated in this case because new wells are introduced after production is underway.

Table 4.12: Optimization results for three runs. Best result shown in bold (Case 2a)

	Run 1 ($MM)	Run 2 ($MM)	Run 3 ($MM)	Average ($MM)	Time (hours)
Conventional opt.	**522.5**	500.8	492.3	505.2	34.2
Multilevel opt.	513.5	508.6	**530.7**	517.6	4.1
Accelerated multilevel opt.	**514.8**	509.9	501.8	508.8	1.8

Optimization results for three runs of the conventional, multilevel and accelerated multilevel optimization approaches are presented in Table 4.12. The multilevel approach provides the highest average of the three runs and the best individual run, though the other two methods provide results that are within about 3% for both quantities. The accelerated multilevel approach is 19 times faster than the conventional approach and about twice as fast as the standard multilevel procedure. Figures 4.17(b)-(d) show the best well locations obtained by the three approaches (the initial guess is shown in Figure 4.17(a)). All three approaches locate the new wells (open circles) away from the existing wells (solid circles). The optimized BHP controls found by the three methods for the two new wells are shown in Figure 4.18. The optimized BHP controls differ substantially between the three methods. Note that these wells are drilled at 730 days, so BHPs are not defined before this time.

We present results for cumulative production and injection profiles in Figures 4.19(a)-(c) for the optimized solutions. These results display less variation than the previous case. All three optimized solutions have similar expected cumulative oil production profiles and provide about 1.4 MMSTB more oil than the reference scenario. Reference

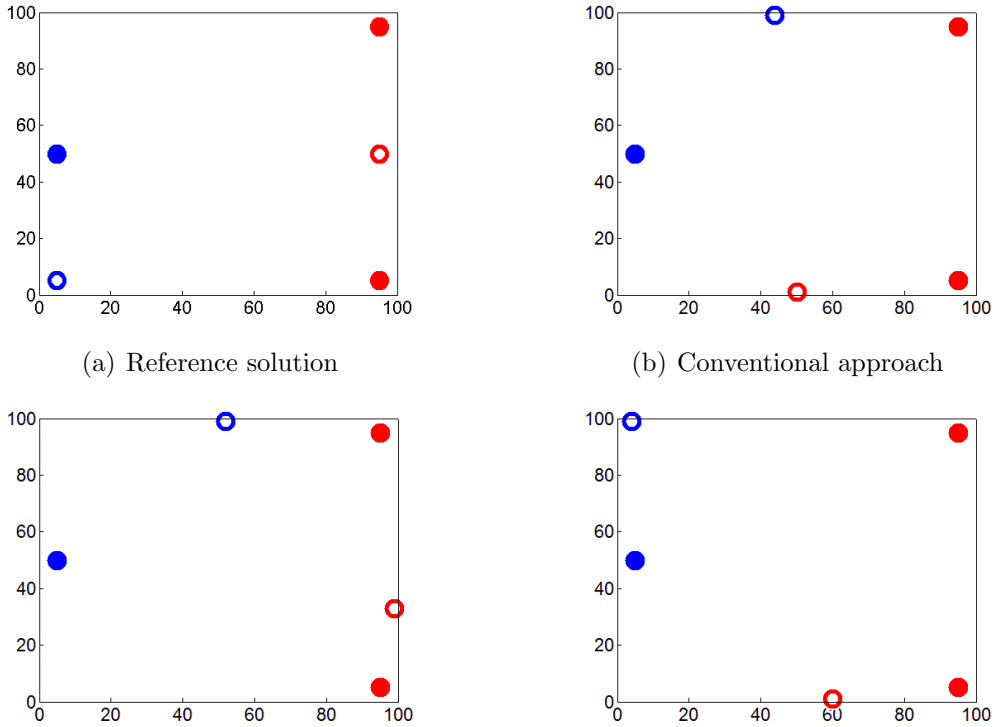

(a) Reference solution

(b) Conventional approach

(c) Multilevel approach

(d) Accelerated multilevel approach

Figure 4.17: Reference (initial guess) solution and best solutions found by the three methods. Existing production and injection wells shown as filled red and blue circles respectively. Optimized production and injection wells are shown as open red and blue circles (Case 2a)

well locations and controls are shown in Figures 4.17(a) and 4.18(a) and (b). Note also that the cumulative oil production for each individual realization is larger than that in the reference scenario. We see in Figure 4.19(b) that the optimized solution from the multilevel procedure produces slightly less water than the solutions from the other two optimization methods. This solution provides the highest NPV.

In Figure 4.20, we show the evolution of the objective function with the conventional, multilevel and accelerated multilevel approaches. In the conventional approach the objective function does not improve after about 25,000 simulation runs. Thus, with better termination criteria, the time required for this approach could be reduced from about 34 hours to perhaps 20 hours. The multilevel procedures also perform more

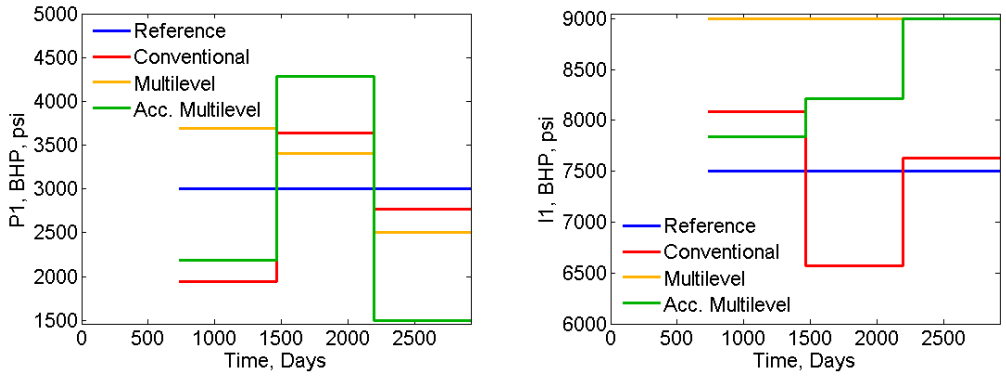

(a) Production well P1 (b) Injection well I1

Figure 4.18: Optimum BHPs found by the three methods (Case 2a)

simulations than are necessary, but most of these are at coarse grid levels, so savings in elapsed time through use of better termination criteria would be fairly modest for these methods.

4.2.2 Case 2b: Optimization with 100 geological models

We now consider all 100 realizations. The OSV parameters are the same as in Case 1b. Accelerated multilevel optimization results are presented in Table 4.13. The total time required for each run is about 18 hours with this approach. The numbers of realizations used by the accelerated multilevel approach with OSV are shown in Table 4.14. Timings for the best run are presented in Table 4.15. The total time required for this case is 7.6 hours.

Table 4.13: Accelerated multilevel optimization computations (Case 2b)

Model size	10×10	20×20	25×25	50×50	100×100
Total time per model, sec	4	6	12	24	110
Number of function evaluations (PSO–MADS/MADS)	1824	456	228	152	76
Elapsed time, hours (using 38 nodes)	5.3	2.0	2.0	2.7	6.1

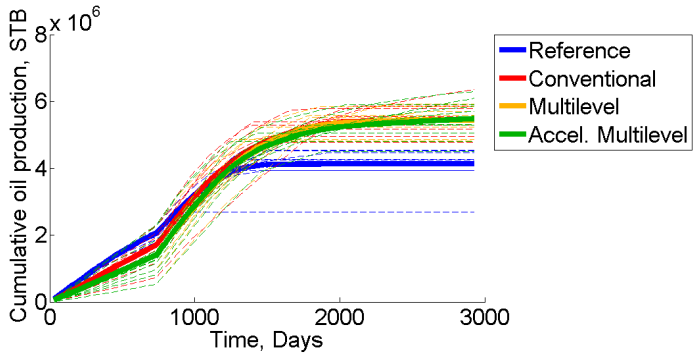

(a) Cumulative oil production

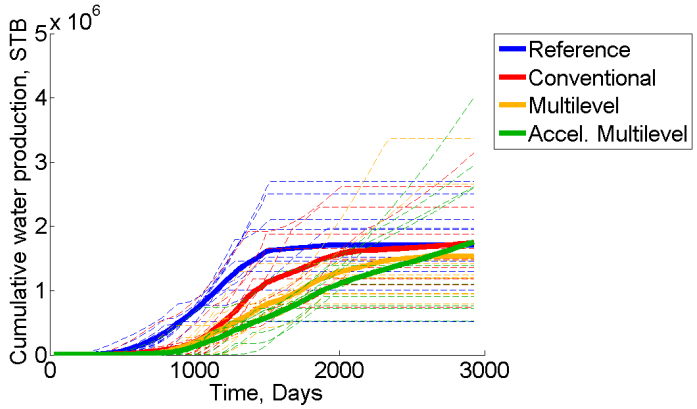

(b) Cumulative water production

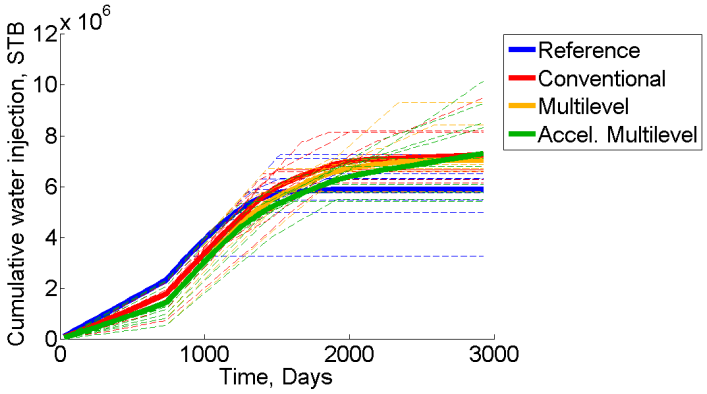

(c) Cumulative water injection

Figure 4.19: Cumulative production and injection profiles for the optimized solutions obtained by the conventional, multilevel and accelerated multilevel approaches (Case 2a)

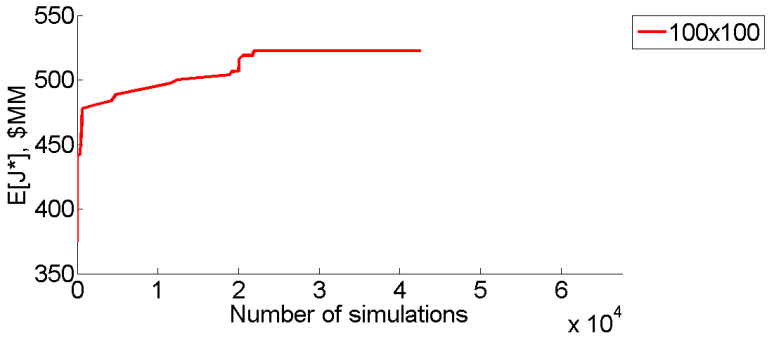

(a) Conventional approach

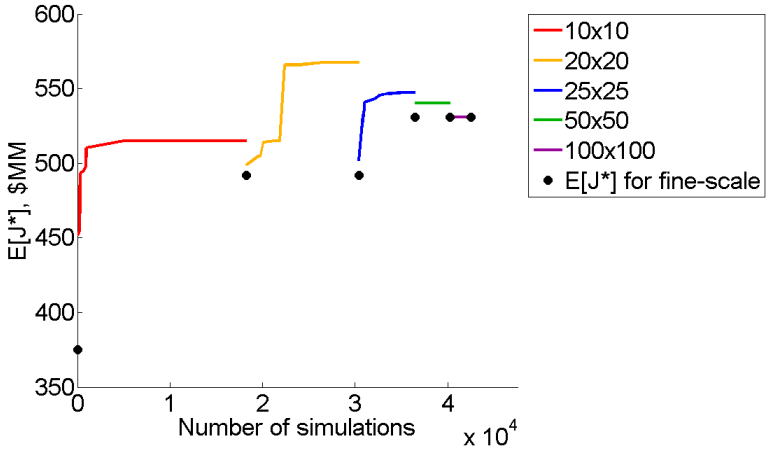

(b) Multilevel approach

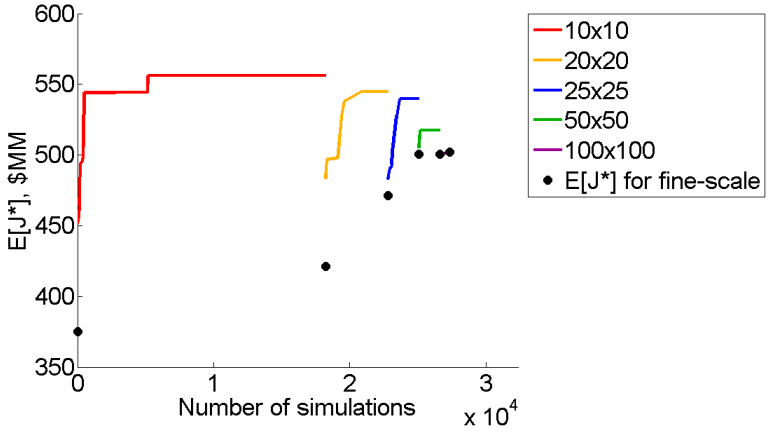

(c) Accelerated multilevel approach

Figure 4.20: Evolution of objective function (Case 2a)

Table 4.14: Number of representative realizations (determined using OSV) and the corresponding relative improvement values at each level for the best run (Case 2b)

Level	N_{rep}	RI
10×10	10	0.69
20×20	10	0.44
20×20	30	0.15
20×20	50	0.72
25×25	10	0.21
25×25	30	0.46
25×25	50	0.89
50×50	10	-1.07
50×50	30	1.36
100×100	10	-0.57
100×100	30	0.52

We compare the expected objective function values with different numbers of representative models to the expected values of all 100 realizations in Figures 4.21(a)-(d) for 38 different well scenarios. As in Case 1b, we see that as the number of representative realizations increases, the estimation of the expected objective function value improves, though there is significant scatter. Figure 4.22 shows how the number of representative realizations for the fine (100×100) models affects the estimation of expected value for all 100 realizations. As in the previous case, accuracy is improved as the number of realizations is increased.

Figures 4.23(a) and (b) show the improvement of the objective function for both approaches. There is a consistent improvement in the actual objective function value in the accelerated multilevel procedure using all 100 realizations. In the accelerated multilevel procedure with OSV, in Levels 2, 3, 4 and 5, the use of a small number of representative realizations does not improve the actual objective function value and the optimization must be restarted with a larger number of realizations.

The optimization results for both approaches are shown in Table 4.16. The accelerated multilevel optimization approach with OSV slightly outperforms accelerated multilevel optimization using all realizations, though the results are very close. The estimated times for the conventional and multilevel approaches with 100 realizations

Table 4.15: Accelerated multilevel optimization computations with OSV for the best run (Case 2b)

Model size	10×10	20×20	25×25	50×50	100×100
Total time per model, sec	4	6	12	24	110
Number of function evaluations with 10 models	1824	456	228	152	76
Number of function evaluations with 30 models	0	456	228	152	76
Number of function evaluations with 50 models	0	456	228	0	0
Number of function evaluations with 100 models	0	0	0	0	0
Elapsed time, hours (using 38 nodes)	0.5	1.8	1.8	1.1	2.4

Table 4.16: Optimization results for three runs. Best result shown in bold (Case 2b)

	Run 1 ($MM)	Run 2 ($MM)	Run 3 ($MM)	Average ($MM)	Time (hours)
Accelerated multilevel opt. (100 realizations)	**498.4**	488.8	489.0	492.1	18.1
Accelerated multilevel opt. with OSV	**504.9**	491.7	487.1	494.6	7.6

are about 342 and 41 hours respectively. The accelerated multilevel with OSV approach would thus provide a speedup of nearly 45 relative to the conventional approach.

The optimum well locations obtained by the two approaches are shown in Figures 4.24(a) and (b). The optimized well locations are again located away from the existing wells, though the locations of the new producers differ in the two cases. Production results for all 100 realizations for the optimized solutions are shown in Figures 4.25(a)-(f). The optimized solutions provide higher expected cumulative oil recovery than the reference case (the reference scenario for Case 2b is the same as in Case 2a). Additional water injection is also evident in the optimized solutions.

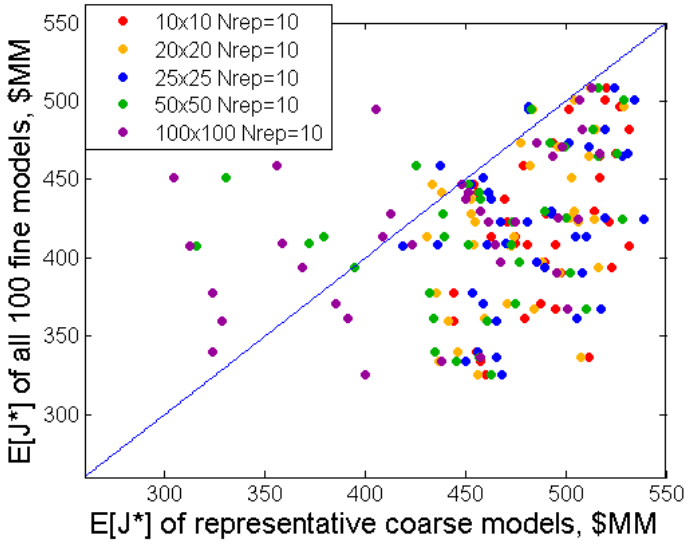

(a) 10 representative models

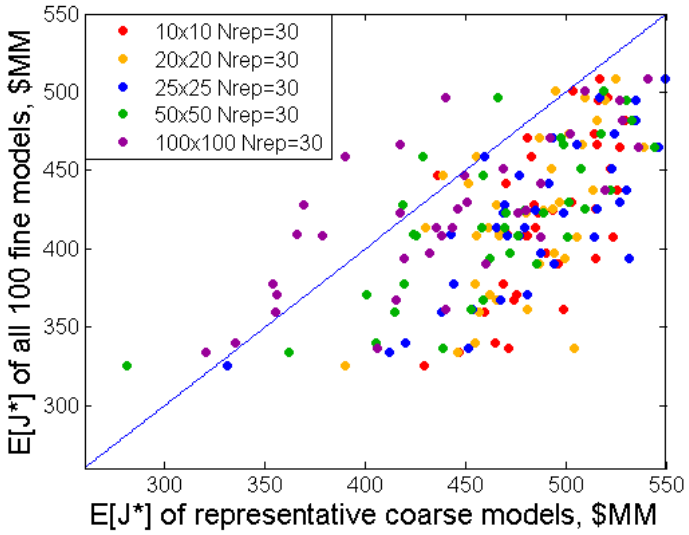

(b) 30 representative models

(Caption on following page)

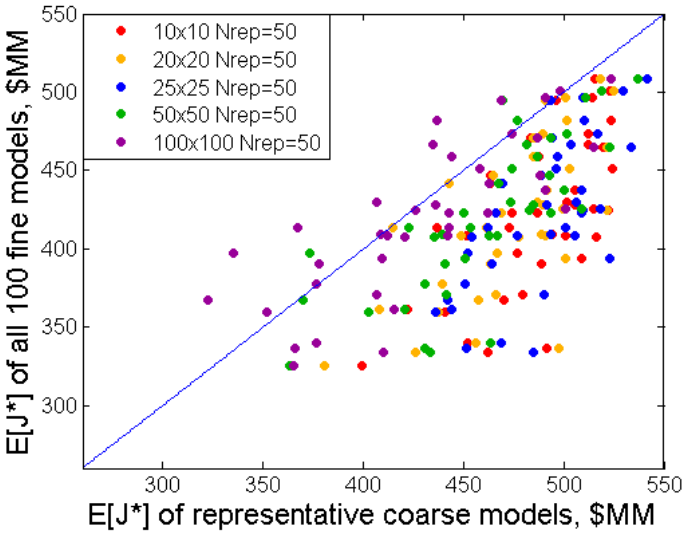

(c) 50 representative models

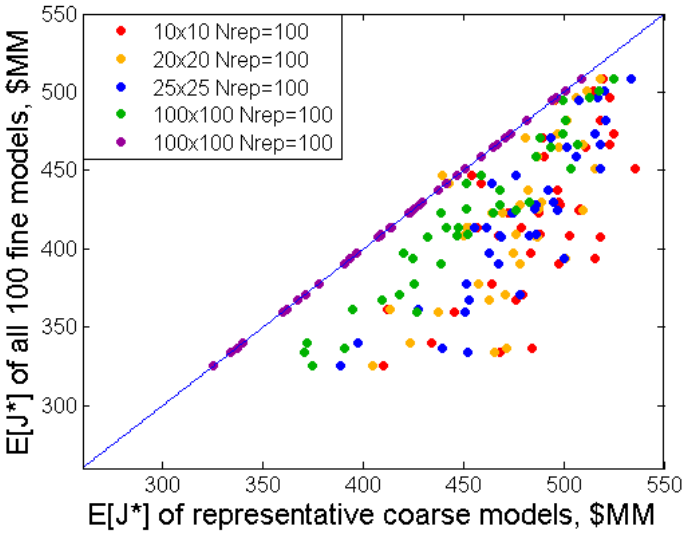

(d) 100 representative models

Figure 4.21: Comparison of expected objective function values evaluated at different grid levels with different numbers of representative models for the 38 candidate well scenarios after 1824 function evaluations using the accelerated multilevel approach with OSV method (Case 2b)

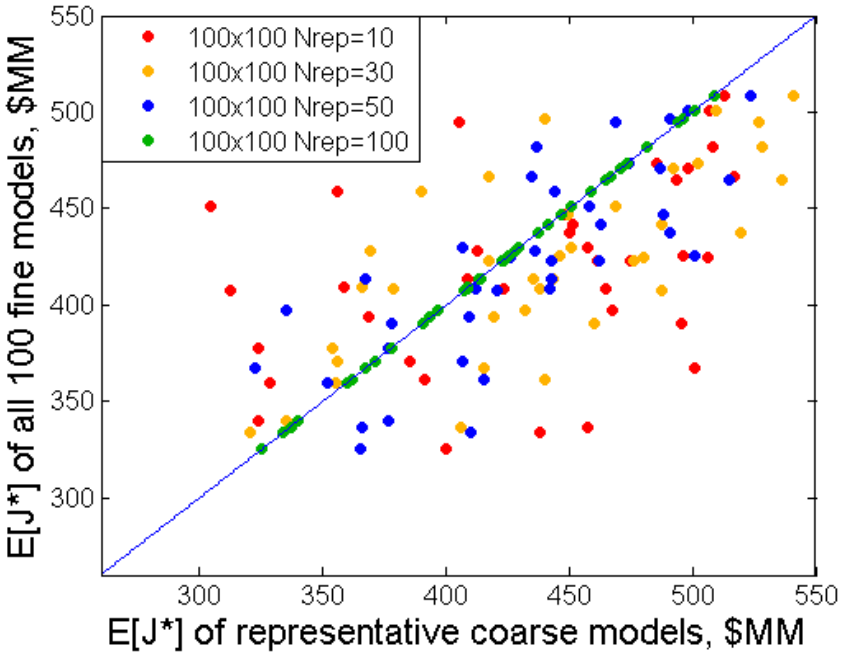

Figure 4.22: Comparison of expected objective function values evaluated at the fine-grid level with different numbers of representative models for the 38 candidate well scenarios after 1824 function evaluations using the accelerated multilevel approach with OSV method (Case 2b)

4.3 Case 3: Three-dimensional channel-levee model

For this case we generated 100 three-dimensional geological models using the Petrel software [53]. The models are of dimensions $30 \times 30 \times 6$. All models are conditioned to hard data from five vertical exploration wells. One production and one injection well that perforate all layers, as shown in Figure 4.26, are considered to have been operating for one year.

In this example, the areal locations of three new vertical wells, their completion intervals, and a single BHP control value for all wells (two existing and three new wells), are to be optimized. The production time frame is nine years. Each new well is defined with four optimization variables (x, y, z_1, z_2), where (x, y) define the areal location and (z_1, z_2) the completion interval. The total number of variables is 17 (12 location and 5 control variables). The PSO swarm size is 34 and MADS evaluates 34

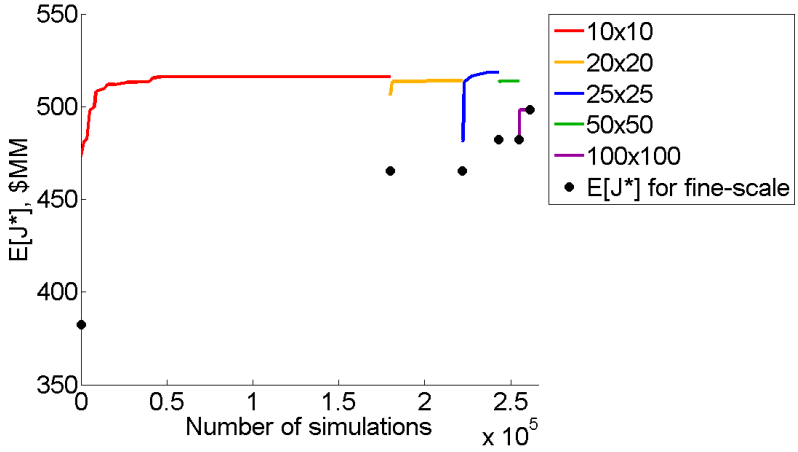

(a) Accelerated multilevel approach with all 100 realizations

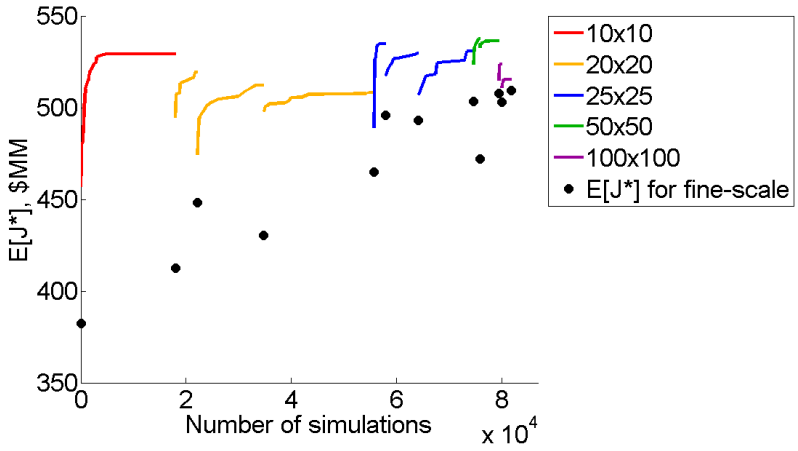

(b) Accelerated multilevel approach with OSV

Figure 4.23: Evolution of objective function (Case 2b)

simulation runs per iteration. The maximum number of available cores is 34. Other simulation and optimization parameters are shown in Tables 4.1 and 4.2. Nonlinear constraints, specifically minimum oil and maximum water production rates, are the same as in the previous cases.

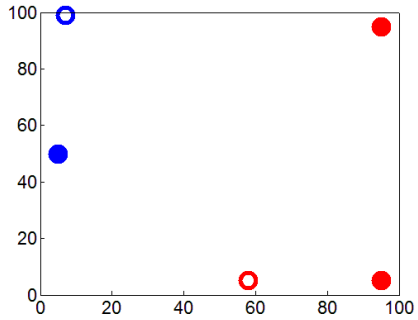

(a) Best solution using accelerated multilevel procedure

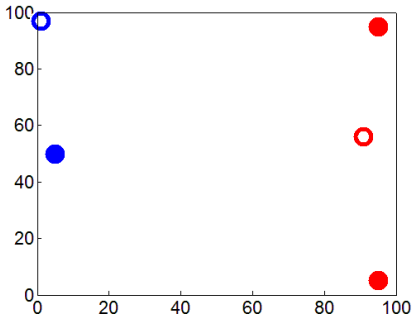

(b) Best solution using accelerated multilevel procedure with OSV

Figure 4.24: Best solutions found by the two methods. Existing production and injection wells shown as filled red and blue circles respectively. Optimized production and injection wells are shown as open red and blue circles (Case 2b)

4.3.1 Case 3a: Optimization with ten geological models

We now compare conventional (single-level) optimization to the multilevel and accelerated multilevel approaches using ten realizations. The number of function evaluations is 4216. Ten simulations are run for every function evaluation in all methods.

Table 4.17: Multilevel optimization computations (Case 3a)

Model size	$10 \times 10 \times 1$	$15 \times 15 \times 1$	$30 \times 30 \times 3$	$30 \times 30 \times 6$
Total time per model, sec	5	7	30	80
Number of function eval. (PSO–MADS)	2040	1224	680	272
Elapsed time, hours (using 34 nodes)	0.8	0.7	1.7	1.8

Details of the computation time for the multilevel and accelerated multilevel approaches are shown in Tables 4.17 and 4.18. The time required to run the multilevel and the accelerated multilevel approaches are 5.0 hours and 1.9 hours respectively. The conventional optimization for this case requires about 28 hours.

The objective function values of the coarse models versus those for the actual models for various well scenarios are shown in Figure 4.27. We again see that the coarse

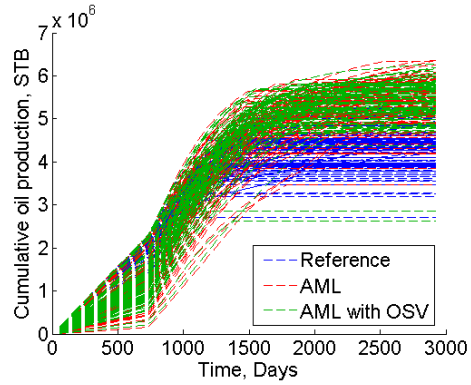

(a) Cumulative oil production

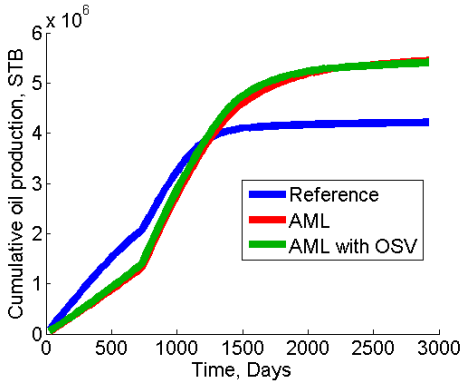

(b) Expected cumulative oil production

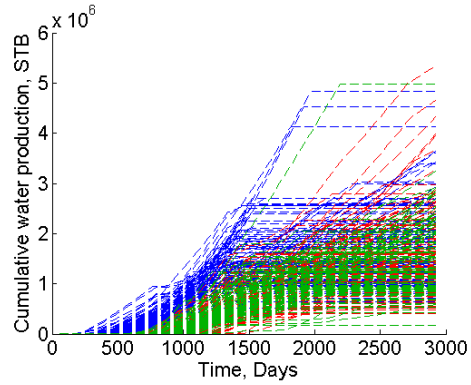

(c) Cumulative water production

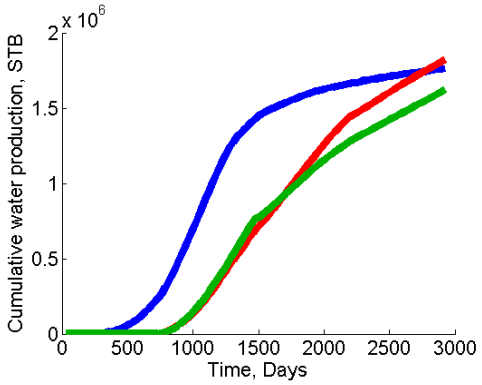

(d) Expected cumulative water production

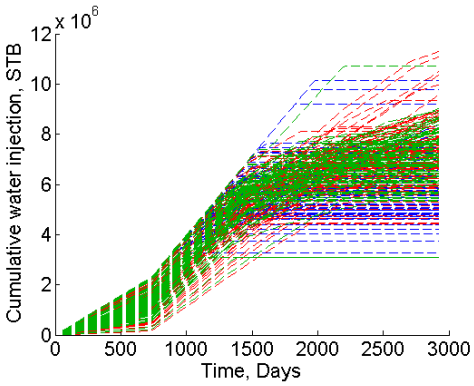

(e) Cumulative water injection

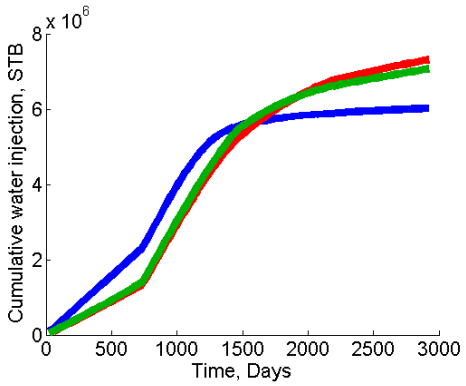

(f) Expected cumulative water injection

Figure 4.25: Cumulative production and injection profiles for the optimized solutions obtained by the accelerated multilevel and accelerated multilevel with OSV approaches. Left plots show results for 100 realizations, right plots show expected values. Note difference in scales (Case 2b)

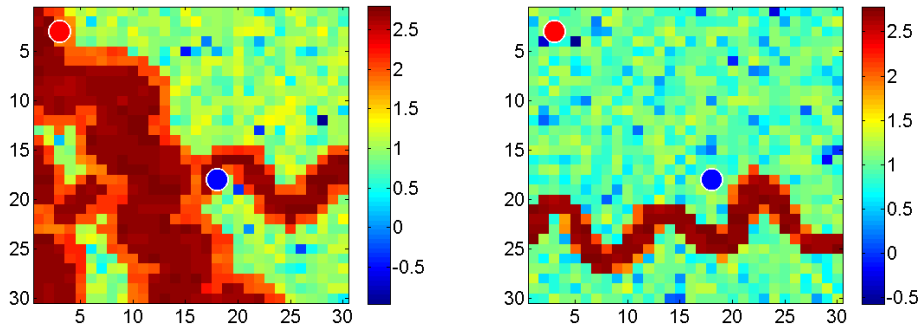

(a) Layer 1

(b) Layer 2

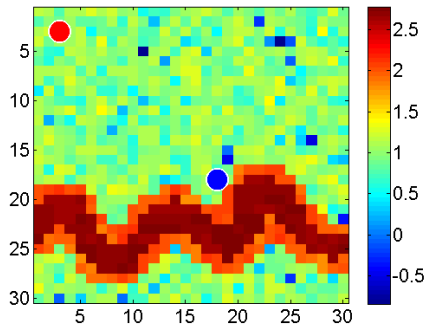

(c) Layer 3

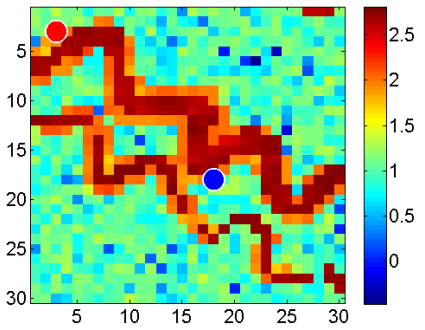

(d) Layer 4

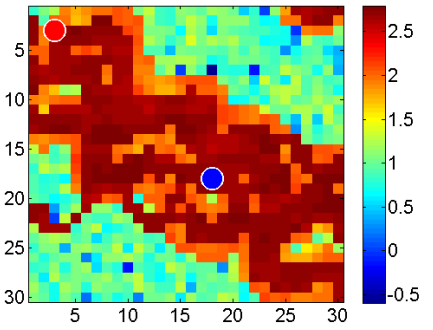

(e) Layer 5

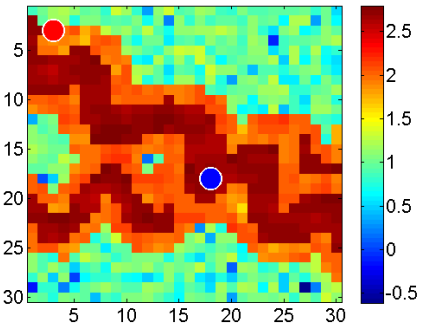

(f) Layer 6

Figure 4.26: Realization 1 of the three-dimensional geological model used for Case 3. Existing production and injection wells shown as red and blue circles respectively. Log_{10} permeability (in md) is shown.

Table 4.18: Accelerated multilevel optimization computations (Case 3a)

Model size	$10 \times 10 \times 1$	$15 \times 15 \times 1$	$30 \times 30 \times 3$	$30 \times 30 \times 6$
Total time per model, sec	5	7	30	80
Number of function evaluations (PSO–MADS/MADS)	2040	306	170	68
Elapsed time, hours (using 34 nodes)	0.8	0.2	0.4	0.4

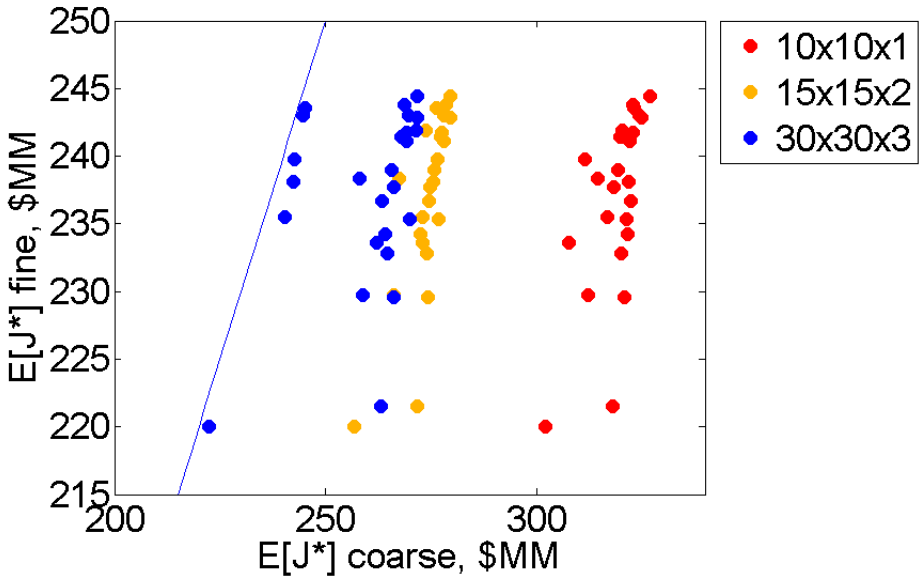

Figure 4.27: Comparison of expected objective function values over 10 realizations evaluated at different grid levels for the 34 candidate well scenarios after 2040 function evaluations using the accelerated multilevel approach (Case 3a)

models (even the $10 \times 10 \times 1$ model) maintain reasonable orderings of the solutions.

Optimization results for three runs using each approach are shown in Table 4.19. Average results are very close for the three methods, though the best run using the accelerated multilevel procedure is 4% less than that using the multilevel procedure. We do not show the production profiles for Case 3a. The optimized solutions obtained by the three methods provide similar cumulative oil and water production profiles, and these correspond to more produced oil and water than in the reference scenario. This

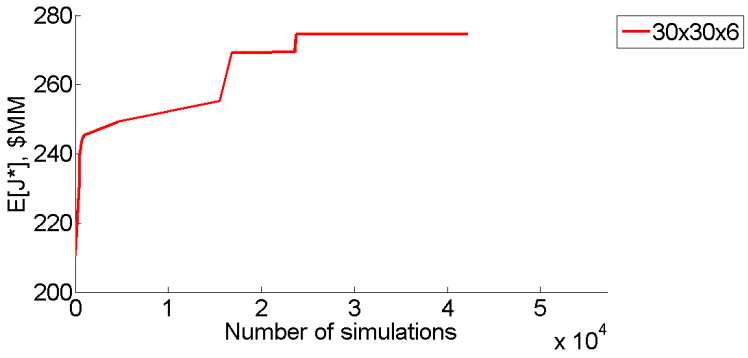

(a) Conventional approach

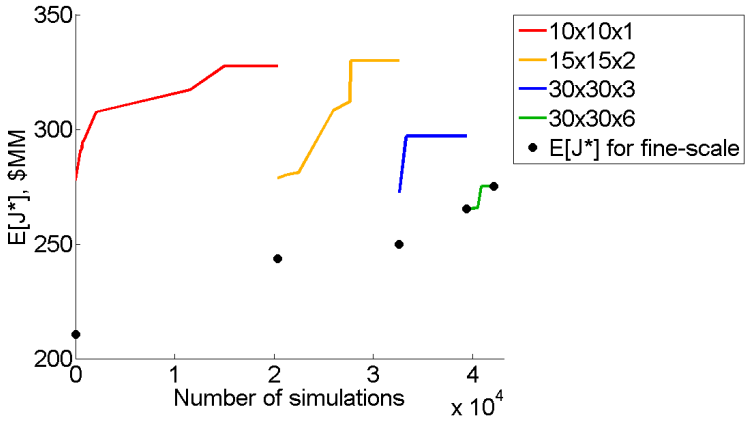

(b) Multilevel approach

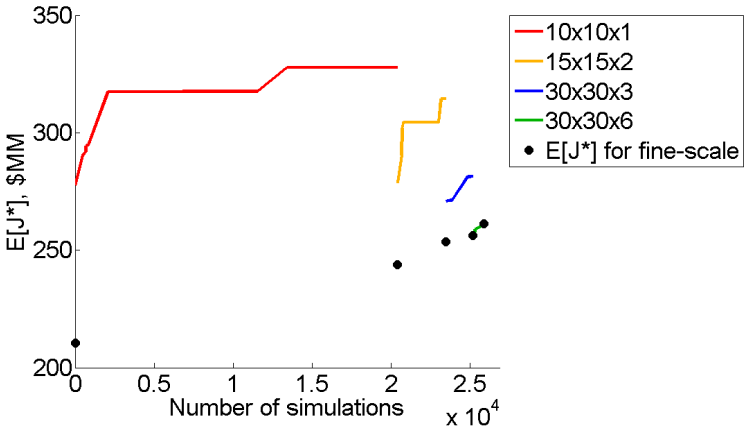

(c) Accelerated multilevel approach

Figure 4.28: Evolution of objective function (Case 3a)

is achieved by injecting more water compared to the reference scenario. The evolution of the objective function values for the three methods are shown in Figure 4.28. The general behavior is consistent with that observed in Cases 1a and 2a.

Table 4.19: Optimization results for three runs. Best result shown in bold (Case 3a)

	Run 1 ($MM)	Run 2 ($MM)	Run 3 ($MM)	Average ($MM)	Time (hours)
Conventional opt.	252.1	248.2	**274.3**	258.2	27.6
Multilevel opt.	**275.2**	251.6	245.1	257.3	5.0
Accelerated multilevel opt.	**264.5**	257.4	251.5	257.8	1.9

Table 4.20: Accelerated multilevel optimization computations with all realizations (Case 3b)

Model size	$10 \times 10 \times 1$	$15 \times 15 \times 1$	$30 \times 30 \times 3$	$30 \times 30 \times 6$
Total time per model, sec	5	7	30	80
Number of function evaluations (PSO–MADS/MADS)	2040	306	170	68
Elapsed time using all 100 models, hours (using 34 nodes)	8.3	1.8	4.2	4.4

4.3.2 Case 3b: Optimization with 100 geological models

The number of models is now increased to 100. Optimization parameters for the accelerated multilevel approach with all 100 geological realizations and with OSV are shown in Tables 4.20 and 4.21. The computation time required for the accelerated multilevel approach with all realizations is 18.7 hours, while that using OSV is 12.4 hours.

Results for three runs using both methods are tabulated in Table 4.22. The two methods provide average and best case results that are within about 2%. The evolution of the objective function value for the best runs are shown in Figures 4.29(a) and (b). In the accelerated multilevel approach with all 100 realizations, there is a consistent improvement in the actual objective function value. For the accelerated multilevel

Table 4.21: Accelerated multilevel optimization computations with OSV for the best run (Case 3b)

Model size	$10 \times 10 \times 1$	$15 \times 15 \times 1$	$30 \times 30 \times 3$	$30 \times 30 \times 6$
Total time per model, sec	5	7	30	80
Number of function eval. with 10 models	2040	306	170	68
Number of function eval. with 30 models	2040	306	170	0
Number of function eval. with 50 models	2040	0	170	0
Number of function eval. with 100 models	0	0	0	0
Elapsed time with OSV, hours (using 34 nodes)	7.5	0.7	3.8	0.4

Table 4.22: Optimization results for three runs. Best result shown in bold (Case 3b)

	Run 1 ($MM)	Run 2 ($MM)	Run 3 ($MM)	Average ($MM)	Time (hours)
Accelerated multilevel opt. (100 realizations)	**318.6**	286.3	303.7	302.9	18.7
Accelerated multilevel opt. with OSV	**312.0**	299.4	301.5	304.3	12.4

approach with OSV, the expected value over all realizations is seen to decrease in some cases when a small number of representative models is used (Figure 4.29(b)). The number of models used for each OSV level for the best run is shown in Table 4.23. We see that the number of realizations must be increased at all levels except for the fine-grid $30 \times 30 \times 6$ level.

The estimated time for the conventional (single-level) approach with 4216 function evaluations, where the 100 actual models are evaluated for every function evaluation, is about 276 hours. The accelerated multilevel approach with OSV is thus about 22 times faster than the conventional approach. The optimized production profiles for this case are similar to those presented above for Case 3a.

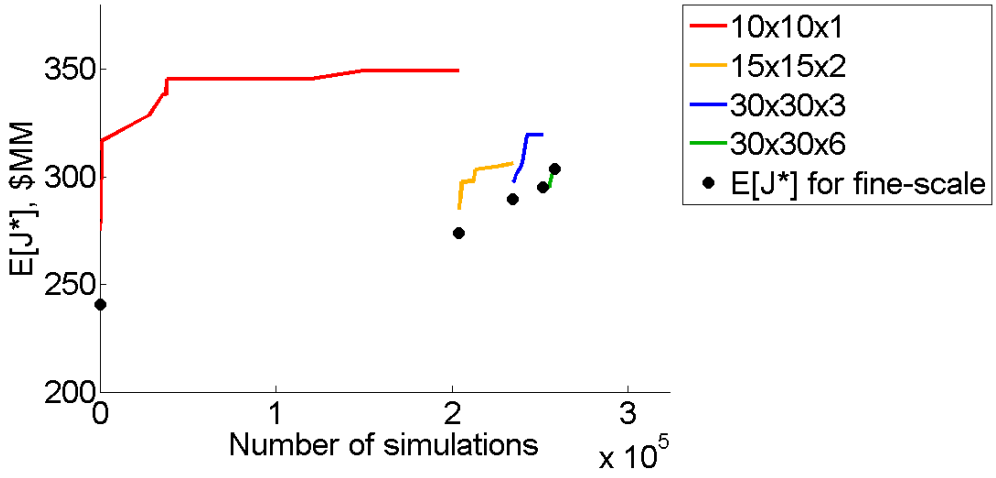

(a) Accelerated multilevel approach with all 100 realization

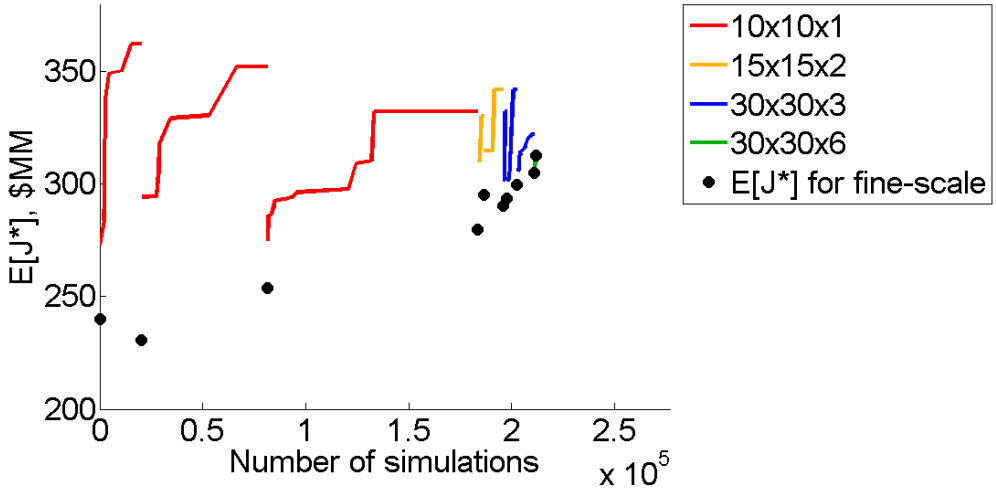

(b) Accelerated multilevel approach with OSV

Figure 4.29: Evolution of objective function (Case 3b)

4.4 Summary

In this chapter, we applied the multilevel procedure for well location and control optimization problems with multiple realizations. Although including multiple realizations in the optimization is computationally expensive, we reduced the computational cost significantly by modifying the multilevel procedure. The modifications

Table 4.23: Number of representative realizations (determined using OSV) and the corresponding relative improvement values at each level for the best run (Case 3b)

Level	N_{rep}	RI
$10 \times 10 \times 1$	10	0.14
$10 \times 10 \times 1$	30	-0.12
$10 \times 10 \times 1$	50	0.92
$15 \times 15 \times 2$	10	-0.39
$15 \times 15 \times 2$	30	0.74
$30 \times 30 \times 3$	10	0.17
$30 \times 30 \times 3$	30	0.42
$30 \times 30 \times 3$	50	0.55
$30 \times 30 \times 6$	10	0.62

introduced were the use of local optimization after the first level and a sample validation procedure to enable the use of an appropriate number of realizations.

Table 4.24: Optimization results for Cases 1a, 2a and 3a

Case	Optimization approach	Best run ($MM)	Average ($MM)	Time (hours)
Case 1a	Conventional opt.	186.3	173.7	88.7
	Multilevel opt.	186.5	169.4	9.2
	Accelerated mult. opt.	183.3	167.9	4.4
Case 2a	Conventional opt.	522.5	505.2	34.2
	Multilevel opt.	530.7	517.6	4.1
	Accelerated mult. opt.	514.8	508.8	1.8
Case 3a	Conventional opt.	274.3	258.2	27.6
	Multilevel opt.	275.2	257.3	5.0
	Accelerated mult. opt.	264.5	257.8	1.9

We presented three optimization problems, each of which involved both ten and 100 geological realizations. The first two problems considered two-dimensional reservoirs, while the third problem involved a three-dimensional reservoir. We divided each problem into two parts. In the first part, we demonstrated that using only MADS in the later levels of the multilevel procedure provides comparable results to the conventional approach for optimization problems with ten realizations. In the second part, we used all 100 realizations in the accelerated multilevel procedure and compared

results to the multilevel procedure with the OSV technique.

The results for optimizations with ten realizations (Cases 1a, 2a, 3a) and 100 realizations (Cases 1b, 2b, 3b) are summarized in Tables 4.24 and 4.25. In Table 4.24, we see that the best individual runs for the three cases were found by the multilevel approach. The conventional method provided the highest average results for Cases 1a and 3a. The accelerated multilevel method provided objective function values that were 1-4% less, for the best individual runs and average results, than those from the conventional and multilevel methods. It is, however, about 20 times faster than the conventional method and two times faster than the multilevel method. From Table 4.25, it is evident that the accelerated multilevel approach provided the best individual runs for Cases 1b and 3b. The highest average results for Cases 2b and 3b were obtained by the accelerated multilevel approach with OSV. The difference between the two methods is 2% or less for best and average runs in all three cases. The accelerated multilevel approach with OSV is two or more times faster than the accelerated multilevel approach in Cases 1b and 2b, though less speedup is observed in Case 3b.

As an alternative to the multilevel approach, we also applied an MLMC procedure in the optimization to reduce the computational cost. This method was applied only in Case 1b. Optimization with the MLMC approach provided optimum objective function values close to those from the accelerated multilevel procedure. Although MLMC is computationally more efficient than the conventional approach, it was more expensive than the multilevel procedure for the case considered.

Table 4.25: Optimization results for Cases 1b, 2b and 3b

Case	Optimization approach	Best run ($MM)	Average ($MM)	Time (hours)
Case 1b	Accelerated multilevel opt. (100 realizations)	148.2	143.8	26.0
	Accelerated multilevel opt. with OSV	147.8	142.6	12.9
Case 2b	Accelerated multilevel opt. (100 realizations)	498.4	492.1	18.1
	Accelerated multilevel opt. with OSV	504.9	494.6	7.6
Case 3b	Accelerated multilevel opt. (100 realizations)	318.6	302.9	18.7
	Accelerated multilevel opt. with OSV	312.0	304.3	12.4

Chapter 5

Summary, Conclusions and Future Work

Existing field development optimization methods are computationally expensive. When geological uncertainty is incorporated into the optimization problem, the computational requirements increase further. In this research, we introduced an optimization framework that improves the efficiency of existing field development optimization algorithms. The key accomplishments and conclusions from this work are as follows:

- We introduced a multilevel optimization framework that entails the use of a sequence of upscaled models. Optimization is performed over models at different levels – from coarsest to finest – and the solution at the previous level is used as an initial guess for the next-level optimization. Essentially any core optimization algorithm could be used in this multilevel framework. In this study, we applied the recently developed PSO–MADS hybrid procedure, which entails global stochastic search (PSO) combined with local pattern search (MADS). A global transmissibility upscaling procedure was applied to generate the coarse-scale models. Because coarse-scale transmissibilities depend on well locations and controls, the upscaling computations must be performed for each candidate solution (in each geological realization) at each iteration.

- We compared multilevel optimization results to results from the conventional

method (which uses only the fine-scale model). The differences in the objective function values between the two approaches were less than 1% for both the best individual run and for the average of three runs, for several example cases. The multilevel approach yielded higher optimum objective function values in some cases. The multilevel procedure provided a factor of 5–10 speedup for the cases considered.

- We extended the multilevel optimization framework to handle optimization under geological uncertainty with nonlinear constraints. This entails upscaling and simulating multiple realizations (rather than a single realization) for every candidate solution. A penalty method was introduced to minimize constraint violations while maximizing the expected value of the objective function over all realizations.

- We incorporated two approaches to improve the efficiency of the multilevel optimization procedure for optimization under uncertainty. First, the PSO-MADS global hybrid algorithm, which requires a large number of simulation runs, is only used at the first level in the multilevel optimization. In subsequent levels, only MADS is used. This approach is referred to as accelerated multilevel optimization. We also applied an optimization with sample validation (OSV) procedure into the multilevel framework to reduce the number of realizations used in the optimization. Reducing the number of realizations in optimization is potentially problematic because the selected realizations might not be representative of the full set of realizations. The OSV procedure allows us to assess the representivity of the selected models, and to increase the number of models used in the optimization when necessary.

- We compared the conventional, multilevel, and accelerated multilevel optimization procedures using ten geological realizations. We showed that the accelerated multilevel approach provided optimum objective function values within about 1–4% of those from the conventional and multilevel approaches. Speedup of about a factor of 2–3 relative to (standard) multilevel optimization was

achieved. We also performed optimization using 100 geological realizations and compared the accelerated multilevel and accelerated multilevel with OSV methods. Results were very similar, though the accelerated multilevel with OSV method is about two times faster than the accelerated multilevel approach.

- We investigated the application of the multilevel Monte Carlo (MLMC) approach for field development optimization under geological uncertainty as an alternative to the multilevel optimization approach. We assessed the performance of MLMC with different model selection strategies and coarsening levels before using it in the optimization. Results for one example demonstrated that optimization using MLMC can provide similar optimum objective function values to those from the other optimization procedures. Although the optimization with MLMC is computationally more expensive than the accelerated multilevel optimization procedures, it is still more efficient than the conventional approach.

The multilevel optimization framework introduced in this work is quite general but there are several additional issues that should be addressed in future work. The following topics are suggested for further study:

- Relative permeability upscaling, in addition to the single-phase upscaling used in the current implementation, could be incorporated. This will require additional upscaling computations, but it would lead to more accurate coarse models and might enable us to perform significantly fewer fine-scale simulations.

- The choice of grid levels, and the optimization parameters to use at these levels, should also be considered. In the examples presented here, we observed that the objective function value improved only slightly at some grid levels. Also, the objective function was seen to stop improving after some number of function evaluations at some levels. This indicates that we could skip or reduce the number of function evaluations at some grid levels. Thus, it will be useful to develop better termination criteria, as well as techniques to determine which grid levels to use in the multilevel optimization.

- Treatments for nonlinear constraints in coarse models should also be studied. Given the general biases of coarse models, it may be useful to modify the constraint handling (or penalty function) as we proceed from level to level. Other constraint handling treatments such as the filter method could also be investigated within the multilevel optimization framework.

- Some problems many involve very large numbers (hundreds or thousands) of realizations. Approaches for such cases, which do not require simulating all realizations, should be developed.

- The use of MLMC for optimization under uncertainty should be further studied. It may be possible to effectively combine this approach with the multilevel procedures developed in this work. This might be accomplished by applying MLMC instead of optimizing over the finer grid levels.

- The introduction of the multilevel optimization procedure into a multiobjective optimization framework should be considered. Multiobjective problems can be very demanding computationally, so the use of our methods could enable two or more (possibly conflicting) objectives to be minimized.

- Finally, the overall methodology should be tested on realistic field cases. In addition, the multilevel optimization algorithm can be considered for use in areas outside of petroleum engineering in which simulation-based optimization is performed.

Nomenclature

Abbreviations

BHP	bottomhole pressure
CDF	cumulative distribution function
GPRS	general purpose research simulator
MADS	mesh adaptive direct search
MINLP	mixed-integer nonlinear programming
MLMC	multilevel Monte Carlo
MM	million
NPV	net present value
OSV	optimization with sample validation
PSO	particle swarm optimization
PSO-MADS	hybrid of particle swarm optimization and mesh adaptive direct search
STB	stock tank barrels

Greek Symbols

μ	viscosity
ρ	density
θ	validation criterion

Variables

B	formation volume factor
C_{drill}	drilling cost per well
$E[.]$	expected value
f_l	flow rate through fine-scale interface l
J	objective function without penalty term

J^*	objective function with penalty term
n	number of optimization variables
N_p	number of PSO particles
N_{real}	number of realizations
N_{rep}	number of representative realizations
q	flow rate
$q_{o,min}$	minimum oil rate
Q_{iw}	injected water rate
Q_{po}	produced oil rate
Q_{pw}	produced water rate
Q_{unmet}	unmet oil production
p	pressure
P_{iw}	price of injected water
P_{pen}	penalty price of oil
P_{po}	price of produced oil
P_{pw}	price of produced water
RI	relative improvement
R_s	penalty function for realization s
T^*	coarse-scale transmissibility
$\mathbf{u}$	vector of continuous well control variables
$\mathbf{v}$	vector of integer (pseudo-continuous) well location variables, or velocity in PSO update equation
WI^*	coarse-scale well index
$\mathbf{x}$	vector of optimization variables
$\mathbf{z}$	vector of categorical variables

Subscripts

I	refers to set of injection wells
l	lower bound
P	refers to set of production wells
o	oil
u	upper bound
w	water

Superscripts

$*$	indicates upscaled quantity
c	coarse
k	iteration level
w	well

Bibliography

[1] A. Y. Abukhamsin. Optimization of well design and location in a real field. 2009. Master's thesis, Department of Energy Resources Engineering, Stanford University.

[2] I. Aitokhuehi and L. J. Durlofsky. Optimizing the performance of smart wells in complex reservoirs using continuously updated geological models. *Journal of Petroleum Science and Engineering*, 48(3-4):254–264, 2005.

[3] E. Aliyev and L. J. Durlofsky. Multilevel field-development optimization using a sequence of upscaled models. Paper SPE 173198 presented at the SPE Reservoir Simulation Symposium, Houston, Texas, USA, 2015.

[4] V. Artus, L. J. Durlofsky, J. E. Onwunalu, and K. Aziz. Optimization of non-conventional wells under uncertainty using statistical proxies. *Computational Geosciences*, 10(4):389–404, 2006.

[5] C. Audet and J. E. Dennis Jr. Mesh adaptive direct search algorithms for constrained optimization. *SIAM Journal on Optimization*, 17(1):188–217, 2006.

[6] C. Audet, G. Savard, and W. Zghal. Multiobjective optimization through a series of single-objective formulations. *SIAM Journal on Optimization*, 19(1):188–210, 2008.

[7] J. W. Bandler, Q. S. Cheng, S. A. Dakroury, A. S. Mohamed, M. H. Bakr, K. Madsen, and J. Søndergaard. Space mapping: The state of the art. *IEEE Transactions on Microwave Theory and Techniques*, 52(1):337–361, 2004.

[8] M. C. Bellout, D. Echeverría Ciaurri, L. J. Durlofsky, B. Foss, and J. Kleppe. Joint optimization of oil well placement and controls. *Computational Geosciences*, 16(4):1061–1079, 2012.

[9] Z. Bouzarkouna, D. Y. Ding, and A. Auger. Well placement optimization with the covariance matrix adaptation evolution strategy and meta-models. *Computational Geosciences*, 16(1):75–92, 2011.

[10] D. R. Brouwer and J. D. Jansen. Dynamic optimization of waterflooding with smart wells using optimal control theory. *SPE Journal*, 9(4):391–402, 2004.

[11] D. A. Cameron and L. J. Durlofsky. Optimization and data assimilation for geological carbon storage. In R. Al-Khoury and J. Bundschuh, editors, *Computational Models for CO_2 Sequestration and Compressed Air Energy Storage*, pages 355–388. Taylor & Francis Group/CRC Press, 2014.

[12] H. Cao. *Development of Techniques for General Purpose Simulators*. PhD thesis, Department of Petroleum Engineering, Stanford University, 2002.

[13] M. A. Cardoso and L. J. Durlofsky. Use of reduced-order modeling procedures for production optimization. *SPE Journal*, 15(2):426–435, 2010.

[14] Y. Chen and L. J. Durlofsky. Efficient incorporation of global effects in upscaled models of two-phase flow and transport in heterogeneous formations. *Multiscale Modeling and Simulation*, (5):445–475, 2006.

[15] Y. Chen, L.J. Durlofsky, M. Gerritsen, and X.H. Wen. A coupled local–global upscaling approach for simulating flow in highly heterogeneous formations. *Advances in Water Resources*, 26(1):1041–1060, 2003.

[16] Y. Chen, B. T. Mallison, and L. J. Durlofsky. Nonlinear two-point flux approximation for modeling full-tensor effects in subsurface flow simulations. *Computational Geosciences*, 12(3):317–335, 2008.

[17] R. Dimitrakopoulos. Stochastic optimization for strategic mine planning: a decade of developments. *Journal of Mining Science*, 47(2):138–150, 2011.

[18] J. F. M. Doren, R. Markovinovic, and J. D. Jansen. Reduced-order optimal control of water flooding using proper orthogonal decomposition. *Computational Geosciences*, 10(1):137–158, 2006.

[19] L. J. Durlofsky and Y. Chen. Uncertainty quantification for subsurface flow problems using coarse-scale models. In I.G. Graham, T.Y. Hou, O. Lakkis, and R. Scheichl, editors, *Numerical Analysis of Multiscale Problems, Lecture Notes in Computational Science and Engineering*, pages 163–202. Springer, 2012.

[20] R. C. Eberhart and J. Kennedy. A new optimizer using particle swarm theory. In *Proceedings of the Sixth International Symposium on Micromachine and Human Science*, pages 39–43, 1995.

[21] D. Echeverría Ciaurri. *Multi-Level Optimization: Space Mapping and Manifold Mapping*. PhD thesis, University of Amsterdam, 2007.

[22] D. Echeverría Ciaurri and P. W. Hemker. Space mapping and defect correction. *Computational Methods in Applied Mathematics*, 5(2):107–136, 2005.

[23] D. Echeverría Ciaurri, O. J. Isebor, and L. J. Durlofsky. Application of derivative-free methodologies for generally constrained oil production optimization problems. *International Journal of Mathematical Modelling and Numerical Optimisation*, 2(2):134–161, 2011.

[24] B. Guyaguler and R. N. Horne. Uncertainty assessment of well-placement optimization. *SPE Journal*, 7(1):24–32, 2004.

[25] B. Guyaguler, R. N. Horne, L. Rogers, and J. J. Rosenzweig. Optimization of well placement in a Gulf of Mexico waterflooding project. *SPE Reservoir Evaluation & Engineering*, 5(3):229–236, 2002.

[26] R. L. Haupt and S. E. Haupt. *Practical Genetic Algorithms*. John Wiley & Sons, 2004.

[27] J. He and L. J. Durlofsky. Reduced-order modeling for compositional simulation by use of trajectory piecewise linearization. *SPE Journal*, 19:858–872, 2014.

[28] T. D. Humphries, R. D. Haynes, and L. A. James. Simultaneous and sequential approaches to joint optimization of well placement and control. *Computational Geosciences*, 18(3-4):443–448, 2014.

[29] O. J. Isebor. *Derivative-free Optimization for Generalized Oil Field Development.* PhD thesis, Department of Energy Resources Engineering, Stanford University, 2013.

[30] O. J. Isebor and L. J. Durlofsky. Biobjective optimization for general oil field development. *Journal of Petroleum Science and Engineering*, 119:123–138, 2014.

[31] O. J. Isebor, L. J. Durlofsky, and D. Echeverría Ciaurri. A derivative-free methodology with local and global search for the constrained joint optimization of well locations and controls. *Computational Geosciences*, 18(3-4):463–482, 2014.

[32] O. J. Isebor, D. Echeverría Ciaurri, and L. J. Durlofsky. Generalized field development optimization with derivative-free procedures. *SPE Journal*, 19:891–908, 2014.

[33] Y. Jiang. *Techniques for Modeling Complex Reservoirs and Advanced Wells.* PhD thesis, Department of Energy Resources Engineering, Stanford University, 2007.

[34] S. Krogstad, H.M. Nilsen, and X. Raynaud. Reservoir management optimization using calibrated transmissibility upscaling. In *Proceedings of the 14th European Conference on the Mathematics of Oil Recovery, Catania, Sicily, Italy*, 2014.

[35] S. Le Digabel. Algorithm 909: NOMAD: Nonlinear optimization with the MADS algorithm. *ACM Transactions on Mathematical Software*, 37(4):44:1–44:15, 2011.

[36] R. M. Lewis and S. G. Nash. Model problems for the multigrid optimization of systems governed by differential equations. *SIAM Journal on Scientific Computing*, 26(6):1811–1837, 2005.

[37] H. Li. *Compositional Upscaling for Individual Models and Ensembles of Realizations*. PhD thesis, Department of Energy Resources Engineering, Stanford University, 2014.

[38] L. Li and B. Jafarpour. A variable-control well placement optimization for improved reservoir development. *Computational Geosciences*, 16(4):871–889, 2012.

[39] M. L. Litvak and P. F. Angert. Field development optimization applied to giant oil fields. Paper SPE 118840 presented at the SPE Reservoir Simulation Symposium, The Woodlands, Texas, USA, 2009.

[40] F. Müller, P. Jenny, and D. W. Meyer. Multilevel Monte Carlo for two phase flow and Buckley–Leverett transport in random heterogeneous porous media. *Journal of Computational Physics*, 250(1):685–702, 2013.

[41] S. G. Nash and R. M. Lewis. Assessing the performance of an optimization-based multilevel method. *Optimization Methods and Software*, 26(4):693–717, 2011.

[42] D. S. Oliver, A. C. Reynolds, and N. Liu. *Inverse Theory for Petroleum Reservoir Characterization and History Matching*. Cambridge University Press, 2008.

[43] J. E. Onwunalu and L. J. Durlofsky. Application of a particle swarm optimization algorithm for determining optimum well location and type. *Computational Geosciences*, 14(1):183–198, 2010.

[44] D. W. Peaceman. Interpretation of well-block pressures in numerical reservoir simulation with nonsquare grid blocks and anisotropic permeability. *SPE Journal*, 23(3):531–543, 1983.

[45] N. Remy, A. Boucher, and J. Wu. *Applied Geostatistics with SGeMS: A Users Guide*. Cambridge University Press, 2009.

[46] R. K. Romeu and B. Noetinger. Calculation of internodal transmissivities in finite difference models of flow in heterogeneous porous media. *Water Resources Research*, 31:943–959, 1995.

[47] P. Sarma, K. Aziz, and L. J. Durlofsky. Implementation of adjoint solution for optimal control of smart wells. Paper SPE 92864 presented at the SPE Reservoir Simulation Symposium, Houston, Texas, USA, 2005.

[48] P. Sarma, L. J. Durlofsky, K. Aziz, and W. H. Chen. Efficient real-time reservoir management using adjoint-based optimal control and model updating. *Computational Geosciences*, 10(1):3–36, 2006.

[49] S. Shan and G. G. Wang. Survey of modeling and optimization strategies to solve high-dimensional design problems with computationally-expensive black-box functions. *Structural and Multidisciplinary Optimization*, 41(1):219–241, 2009.

[50] M. G. Shirangi and L. J. Durlofsky. Closed-loop field development optimization under uncertainty. Paper SPE 173219 presented at the SPE Reservoir Simulation Symposium, Houston, Texas, USA, 2015.

[51] S. N. Sivananadam and S. N. Deepa. *Introduction to Genetic Algorithms*. Springer-Verlag, 2008.

[52] H. Su and D. S. Oliver. Smart well production optimization using an ensemble-based method. *SPE Journal*, 13(6):884–892, 2010.

[53] Seismic to Simulation Software Petrel Introduction Course. Technical report, Schlumberger, 2008.

[54] G. van Essen, M. Zandvliet, P. Van den Hof, O. Bosgra, and J. D. Jansen. Robust waterflooding optimization of multiple geological scenarios. *SPE Journal*, 14(1):202–210, 2009.

[55] A. I. F. Vaz and L. N. Vicente. A particle swarm pattern search method for bound constrained global optimization. *Journal of Global Optimization*, 39(2):197–219, 2007.

[56] H. Vo. *New Geological Parameterizations for History Matching Complex Models*. PhD thesis, Department of Energy Resources Engineering, Stanford University, 2015.

[57] C. Wang, G. Li, and A. C. Reynolds. Production optimization in closed-loop reservoir management. *SPE Journal*, 14(3):506–523, 2009.

[58] H. Wang, D. Echeverría Ciaurri, L. J. Durlofsky, and A. Cominelli. Optimal well placement under uncertainty using a retrospective optimization framework. *SPE Journal*, 17(1):112–121, 2012.

[59] T. Wen, M. R. Thiele, D. Echeverría Ciaurri, K. Aziz, and Y. Ye. Waterflood management using two-stage optimization with streamline simulation. *Computational Geosciences*, 18(3):483–504, 2014.

[60] K. C. Wilson and L. J. Durlofsky. Optimization of shale gas field development using direct search techniques and reduced-physics models. *Journal of Petroleum Science and Engineering*, 108(1):304–315, 2013.

[61] B. Yeten, L. J. Durlofsky, and K. Aziz. Optimization of nonconventional well type, location and trajectory. *SPE Journal*, 8(3):200–210, 2003.

[62] P. Zhang, G. Pickup, and M. Christie. A new practical method for upscaling in highly heterogeneous reservoir models. *SPE Journal*, 13(1):68–76, 2008.

[63] Y. Zhou. *Parallel General-Purpose Reservoir Simulation with Coupled Reservoir Models and Multisegment Wells*. PhD thesis, Department of Energy Resources Engineering, Stanford University, 2012.