REAL-TIME OPTIMIZATION OF SMART WELLS

A REPORT SUBMITTED TO THE DEPARTMENT OF PETROLEUM ENGINEERING

OF STANFORD UNIVERSITY

IN PARTIAL FULFILLMENT OF THE REQUIREMENTS FOR THE DEGREE OF MASTER OF SCIENCE

By
Inegbenose Aitokhuehi
June 2004

I certify that I have read this report and that in my opinion it is fully adequate, in scope and in quality, as partial fulfillment of the degree of Master of Science in Petroleum Engineering.

_____

Prof. Louis J. Durlofsky
(Principal Advisor)

Abstract

Smart wells are wells that have downhole instrumentation, such as sensors and valves, on the production tubing. These wells provide the ability for both downhole monitoring and control. Downhole monitoring can be achieved through the use of sensors while control is realized with downhole valves. Once a smart well is deployed, valves can be used to independently control each segment / branch of the well in a reactive mode, such as shutting off a zone once it starts producing water, or in a defensive mode, which requires the a priori determination of valve settings. Using the latter approach, which is the method applied in this work, valve settings are determined through an optimization procedure. We show with this procedure that well instrumentation can provide over 50% gain in cumulative oil recovery over the uninstrumented case for systems considered here in which the geology is assumed to be known. Because the geology is not known in real applications, we couple the valve optimization procedure with history matching techniques, in which we use idealized sensor data to update the reservoir description. Up to 90% of the gain attainable with known geology is achieved for the unconditionally and conditionally generated models considered. In addition, we show that it is beneficial to use multiple history-matched models for the optimization in some cases. This is because multiple history-matched models capture the geologic uncertainty better than single history-matched models. We also introduce efficient alternative procedures to improve the speed of the overall technique. These include the use of a Levenberg-Marquardt algorithm for the optimizations.

Acknowledgments

I thank the God of all grace for His mercies and goodness because we all can do nothing without Him. He is my source and help through periods when I travailed.

I appreciate my wife Oyeyemi, and my children, Osedebamen and Ayooluwanitemi, for their love and support. Being with them gave me the strength to continue.

Many thanks go to my advisor, Prof. Lou Durlofsky for his insightful comments and encouragement throughout the course of my MS programme. His guidance, advice and confidence in my work have made this thesis possible. Many thanks also to Prof. Khalid Aziz, for his helpful comments concerning the research. I appreciate Dr. Burak Yeten for his contribution towards the valve optimization module of this study, Prof. Jef Caers, Inanc Tureyen and Satomi Suzuki, for their contributions towards the history matching module of this work.

I thank my friends, especially Dayo Adeogba and family, for their encouragement and support throughout the whole programme. I thank God for surrounding me with helpful people such as Martins Akintunde, Lourdes and Balz Colmeneras, Scot and Christine Lincke and their son Noah, Busola Thomas, Shirley Smith, Marcia Houghton (and the Swamp Asset team of Chevron Nigeria Limited), Abbey Bodunrin, Ezioma Okengwu, Oyinloye Oyelade and a host of other people. I especially acknowledge Pastor and Mrs Muyiwa Maku for the knees that were continually bent in prayers for us.

My gratitude is also expressed towards the entire staff (both teaching and administrative) and students of the Petroleum Engineering department at Stanford, chaired by Prof. Roland Horne, for creating a wonderful environment to learn.

I gratefully acknowledge the financial support from Chevron Nigeria Limited and for the opportunity given me to pursue this MS programme. I also acknowledge the Stanford

University Project on the Productivity and Injectivity of Advanced Wells (SUPRI-HW) and the U.S. Dept. of Energy (contract number DE-AC26-99BC15213) for the opportunity to present the work.

Contents

List of Tables

List of Figures

Chapter 1

1. Introduction

A conventional well is a vertical or slightly deviated well. Horizontal, highly deviated and multilateral wells are generally referred to as nonconventional or advanced wells (NCWs). Smart wells are wells (both conventional and nonconventional) that have been instrumented with downhole sensors and valves on the production tubing. This instrumentation provides the ability to independently monitor and control each segment (or branch) of the well and can therefore be used to maximize oil recovery or to minimize the production of unwanted water or gas. A direct outcome of this is reduced well intervention, particularly in subsea wells. Smart wells are especially beneficial for multiple reservoirs where the main production strategy is to commingle production. Figure 1.1 shows a typical NCW intersecting multiple pay zones. Packers installed in the wellbore separate the zones into segments. These zones are independently controlled by inflow control valves but produce into the same production tubing. Downhole valves can be appropriately (or optimally) located within the drainhole of a horizontal well to provide significant gains. This is especially beneficial for a horizontal drainhole that contacts both sand and shale. With multilateral wells, because the laterals are completed as open holes, it is appropriate to install the valves at the junction of the lateral and the mainbore.

Usually, geologic models generated with sparse data, such as well log, core and seismic data, are used to make decisions about deploying a NCW. However, once the well is deployed, the challenge is then to optimize its production with the available geological, well and acquired production data. With smart wells, downhole sensors provide continuous production data that can be used to update (history match) the geologic model. This updated model can then be used to determine optimal settings for the downhole

valves. The combined use of history matching and valve optimization is the objective of this work.

Figure 1.1: Typical smart well penetrating multiple productive zones.

In the following sections, we present a brief review of previous work on valve optimization and history matching techniques. We then present the problem statement and the solution approach.

1.1. Literature Survey

1.1.1. Optimization Techniques

Optimization algorithms can be categorized into deterministic and stochastic techniques. Deterministic techniques can further be classified into evaluation-only methods and gradient-based methods (Yeten, 2003). The evaluation-only methods do not require any gradient calculations and are therefore simpler but can be extremely slow and inefficient. One such method is the polytope algorithm, which is also known as the Nelder-Mead simplex search or the downhill simplex method (Press *et al.*, 2002). This method has been combined with stochastic methods for well placement problems (Badru and Kabir,

2003; Guyaguler *et al.*, 2002; Bittencourt and Horne, 1997). Gradient-based algorithms, which include the steepest descent method, conjugate gradient method, Newton's method, quasi-Newton method, Gauss-Newton and Levenberg-Marquardt, use the gradient or derivative of the objective function to guide their search. Some of these gradient-based methods also use the Hessian or the second derivative in their search for the optimum solution. However, these methods do not guarantee convergence to a global optimum, except when the objective function is smooth and continuous.

Stochastic methods, such as genetic algorithms (GA) and simulated annealing (SA), attempt to find the global optimum. These methods do not require gradient information. Rather, they utilize function evaluations to seek the global optimum. They can be quite slow and may require many function evaluations. Stochastic methods have been applied to a wide range of problems in the oil and gas industry including production optimization and, quite recently, well placement problems (Yeten *et al*., 2003). In this study, we will use the conjugate gradient method basically because of the super-linear convergence and the nature of the objective function specified for the problem.

1.1.2. Smart Well Control Optimization

Once a smart well has been deployed, maximizing its recovery (or net present value, NPV) requires finding the optimum settings of the installed valves. Different methods exist for finding the optimum valve settings. Because valves can be operated either in a reactive mode or in a defensive (or predictive) mode, previous methods can be classified under these two modes of control.

The reactive mode of control (Yeten, 2003; Khargoria *et al*., 2002) is a type of control where zones (well segments / branches) with high water cut and / or high gas-oil ratio (GOR) problems are isolated by continuous reduction or outright closure of valve openings. With this mode of control, downhole sensor data is used directly to identify zones producing unwanted water and / or gas. The static optimization method used by Brouwer *et al*. (2001) to waterflood a reservoir can be classified as a reactive mode of control. They considered an oil reservoir completed with fully penetrating horizontal

injection and production wells that were rate constrained. In their approach, the valves were operated in an on-off manner and were closed once water breakthrough occurred at the producers. The flow to the closed valves was then redirected by allocating the lost production to the other producing segments of the well. The graphical means used by Gai (2001) to optimize valve settings for a multilateral multi-zone completion can also be classified as reactive control. Gai generated straight line inflow performance relationships and valve performance relationships to optimally produce from the laterals. One requirement of this method is a means to update the inflow and valve performance relationships whenever well conditions change. Arenas and Dolle (2003) applied a pressure cycling principle to waterflood a tight fractured reservoir. They essentially considered a two-dimensional model completed with a horizontal producer and injector. The reservoir contained a set of fractures spanning the two wells. Valves were installed in the injector only and used in the reactive mode to optimize recovery. They were operated in an on-off fashion and were closed or opened once water cut measured at the producer reached a predefined upper or lower limit, respectively.

Reactive control has been applied in real applications. The operation of smart well valves to maximize recovery in two Asian fields (reported by Snaith *et al*., 2003) essentially utilized reactive control. Also, Al-Khodhori (2003) reported the implementation of a reactive smart quad-lateral well technology to enhance production and recovery in an Oman field.

Instead of reacting to problems occurring at the wells, a defensive mode of control (Yeten, 2003; Yeten *et al*., 2002) can be adopted. This mode of control entails an optimization and requires that valve settings be determined a priori using a predictive reservoir and well model. The aim of defensive control is an early mitigation of problems that might otherwise occur at the wells. Dolle *et al*. (2002) utilized optimal control theory to dynamically determine the valve settings. They considered the same reservoir and rate-constrained wells as in their previous work (Brouwer *et al*., 2001). An adjoint method of optimization, which is gradient-based, was then applied to find the optimum valve settings. With this method of defensive control, they were able to improve sweep

and recovery over their previous method of static optimization. In subsequent work (again using the adjoint method of optimization), Brouwer and Jansen (2002) further explored the effects of well constraints on the scope of optimization.

More recently, Yeten (2003) and Yeten *et al*. (2002) used a nonlinear conjugate gradient method in conjunction with a commercial simulator that contains a detailed multi-segment well model (Holmes *et al*., 1998) to optimize valve settings defensively. Their approach entailed dividing the simulation period into equal optimization steps or time periods. Valve settings were then optimized in time for each of the optimization steps. They investigated the effect of geological uncertainty and showed that it directly impacted the gains of downhole control in terms of cumulative recovery. Their method was based on a fixed reservoir model during the optimization process, although sensor data allow for continuous model updating. As a result, their method was most appropriate for use as a smart well screening tool. In their well placement work, Yeten *et al*. (2004) modeled the operational reliability of downhole valves. They applied a utility risk function to model downhole valve failure and explored its effect on recovery.

The downside of using the conjugate gradient method to determine optimum valve settings is that it can be slow. Gradients are computed numerically and therefore require as many function evaluations as there are control variables. The adjoint method used by Dolle *et al*. (2002) and Brouwer *et al*. (2002), on the other hand, has the advantage of being intimately coupled with the reservoir simulator. As a result, the gradient information can be computed directly and efficiently. However, the implementation of the adjoint method is complicated. Ongoing research is focused on streamlining the link between the simulation equations and the adjoint model.

The work presented in this report follows the approach of Yeten (2003) and Yeten *et al*. (2002) but utilizes the sensor data to update (history-match) the geological model, which is then used to optimize the valve settings in real time.

1.1.3. History Matching Techniques

History matching is an ill-conditioned inverse problem resulting in nonunique solutions. History matching (HM) techniques can be grouped into traditional HM methods (THM) and automatic HM methods (AHM). The traditional methods examine the HM problem in a series of scales. They first match the pressure by tuning the directional transmissibilities and the pore volume and then adjust the relative permeability curves to obtain a saturation match. These approaches assume a single (or very few) models. They are very time consuming and rely on experience. Automatic methods were therefore developed to address these and many other issues.

Kabir *et al*. (2003) compared the traditional HM method with two forms of the automatic method using a real oil and gas field. They used the Gauss-Newton optimization algorithm in the two AHMs to minimize a least-square type error between the observed data and the simulated data. Based on their results, they concluded that the AHM approaches reduced the time required to complete a HM study even with multiple parameter adjustments. In addition, they stated that the AHMs preserved the HM objectivity.

Caers (2003) proposed another form of AHM, the probability-perturbation method. The method is facies based and conditions a model to production data by perturbing an initial geologic realization iteratively until a match is achieved. The geologic models are generated with multiple-point geostatistics using a training image, which is an unconstrained conceptual image of the reservoir. As a result, the probability-perturbation method is suitable for diverse geological structures including channels and fractures. The method honors the geology because the models are also conditioned to a specified training image event. Caers modeled the magnitude of the facies perturbation with a single parameter determined through a one-dimensional (1D) optimization.

Hoffman and Caers (2003) extended the probability-perturbation method to handle regional facies distributions with known within-facies petrophysical properties. They divided the reservoir model into regions and perturbed each of them independently to

determine the overall facies distribution that matches the history data. Suzuki (2003) also extended the probability-perturbation method to treat channel systems and developed a method of perturbing the unknown within-facies petrophysical properties in conjunction with the facies perturbation.

The work presented in this report utilizes the combined works of Hoffman and Caers, with the reservoir specified as a single region, and that of Suzuki. Both of these algorithms are from the Stanford PE SCRF research group.

1.2. Statement of the Problem

As stated earlier, Yeten (2003) and Yeten *et al*. (2002) presented an optimization procedure for smart well valve settings. Their strategy was to maximize the cumulative recovery by controlling the valve settings defensively. This approach was based on an a priori geological model that remained fixed throughout the entire optimization process. However, sensors installed in smart wells provide the ability to continuously monitor in situ flow rates, temperature and pressure. These sensor data can be utilized to regularly update the geologic model, which can in turn be used in the valve optimization. The challenge therefore is to use the sensor data to update the geologic model and to apply it to optimize the valve settings. The implementation of a procedure that combines valve optimization and history matching has not been previously reported.

1.2.1. Solution Approach

We approach this problem by coupling smart well valve optimization and control with history matching or geologic model updating. We condition an initial geologic model to a training image event and update it using sensor data in conjunction with the probability-pertubation method of HM. We then utilize the updated model to optimize the downhole valve settings at different stages of production. We assume that downhole sensors for measuring in situ phase flow rates are available. Although these sensors exist for conventional wells, they are not yet available for NCWs (though, it is worth noting that this is an active area of research). For NCWs, downhole pressure and temperature data can be used to estimate phase flow rates.

Two particular drawbacks of HM are that the solutions are not unique and that the match may not be sufficiently close. We address both problems by optimizing the valve settings over multiple history-matched models to explore the potential benefits of using multiple updated models.

1.3. Report Outline

This report proceeds as follows. Chapter 2 covers the solution method in greater detail. The previous approach for optimizing valve settings and its limitations are discussed. The way in which the updated geologic models are integrated into the existing valve optimization approach is explained in depth. The HM algorithm used is also described and validated for a few sample models.

Chapter 3 presents the application of the overall method (smart well valve optimization and HM) to synthetic examples. These examples include a variogram-based model and channel systems. The results for both unconditionally and conditionally generated models are presented and discussed.

Although the overall procedure performs quite well, the downside is that it is time consuming. We address this in Chapter 4 by considering alternative optimization procedures. These techniques include the Levenberg-Marquardt (LM) method. The implementation of the LM method is presented and the results are compared with those of the conjugate gradient method.

Chapter 5 provides a summary and conclusions. Areas for future work are also suggested.

Chapter 2

2. Well Optimization with History Matching

Smart wells provide the ability to monitor and control fluid rates and pressure with sensors and downhole valves. The data from the sensors can be used to periodically or continuously update the geological model. Using the updated model, valve settings can be determined via optimization by maximizing a specified objective function. In order to evaluate the objective function, it is necessary to perform simulations that accurately model the smart well. ECLIPSE (GeoQuest 2001a, b) was used as the function evaluator because of its detailed multi-segment well option, which includes the capability to model smart wells (Holmes *et al*., 1998).

2.1. Modeling of Smart Wells

Figure 2.1 below shows the schematic of a horizontal smart well. The well is segmented into three independent branches or segments by packers. Each segment has a valve installed within it on the production tubing. Fluid enters the well from the reservoir and flows into the valves in each segment. The flow from each valve is commingled in the production tubing. This horizontal smart well can be represented with the ECLIPSE multi-segment well model as a network of segments as shown in figure 2.2. The production tubing is represented as three segments, each with a node that receives flow from the valve segment/node. The annular space within each segment is also divided into a series of segments with a nodal link to each grid block. Fluid from the grid blocks is received in the annular space through the nodal link and subsequently flows into the valve node.

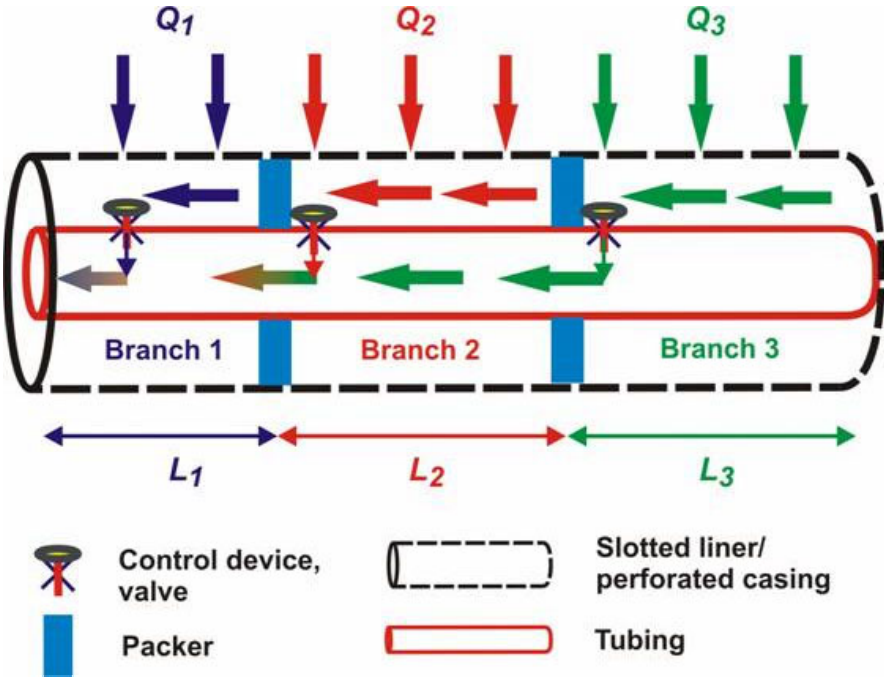

Figure 2.1: Schematic of a horizontal smart well (adapted from Yeten, 2003).

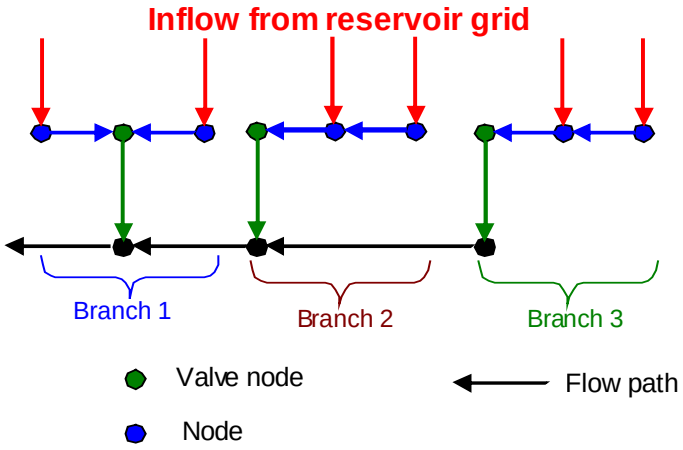

Figure 2.2: ECLIPSE multi-segment network model representation of horizontal smart well.

2.2. Control Strategies

The valves on the smart well can be controlled in a reactive mode or in a defensive mode. The following sections describe both of these control strategies in greater detail.

2.2.1. Reactive Control Strategy

As explained in the previous chapter, this type of valve control isolates well segments or branches that produce high water and / or gas rates harmful to recovery by reducing the valve openings continuously or closing them entirely. This strategy, which is the approach most commonly employed in the oil and gas industry, generally involves trial and error. It can be implemented in ECLIPSE through the use of frictional pressure drop multipliers. The relationship between the desired frictional pressure drop multiplier, water and gas production is prescribed by the following equation (Yeten, 2003; GeoQuest 2001a, b):

$$P_f^m = max\left[A + B(WOR)^C + D\left(\frac{R}{R_{\min}}\right)^E, 1.0 \right] \tag{2.1}$$

P_f^m is the frictional pressure multiplier for each valve segment, while *WOR* and *R* are the water-oil and gas-oil ratios at the valve segment. $R_{\min}$ is a scaling factor for the gas-oil ratio (*R*) and is specified as an input. The parameters *A*, *B*, *C*, *D* and *E* are constants that can vary in time. They can be chosen through test runs or can be determined through an optimization algorithm (Yeten, 2003).

2.2.2. Defensive Control Strategy

With this control strategy, optimum valve settings are determined a priori using a predictive reservoir and well model so that future problems can be mitigated early. These problems can be harmful to the recovery and may be caused by heterogeneities such as high permeability channels or fractures within a reservoir. The intent therefore is to delay water and / or gas breakthrough and to accelerate production. Defensive control has been implemented using a nonlinear conjugate gradient method of optimization (Yeten, 2003;

Yeten *et al*., 2002). The work presented in this report uses the defensive control strategy in conjunction with history matching to optimize smart well recovery.

2.3. Overall Valve Control Optimization and History Matching Procedure

Usually, sub-surface geology is modeled with sparse well data combined with seismic information. The resulting model is only an approximation of the actual geology. Therefore, to implement defensive control, history matching will need to be linked with valve control optimization. Figure 2.3 shows the work flow diagram describing the linkage of both techniques.

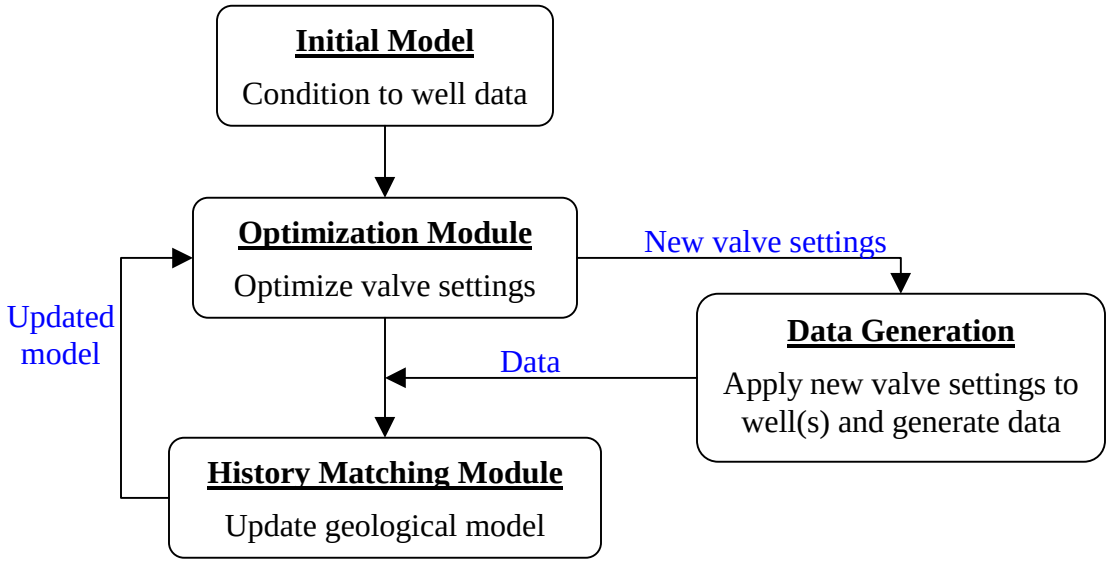

Figure 2.3: Work flow diagram of the overall valve control optimization and history matching procedure.

We can summarize the overall valve control optimization and history matching procedure as follows:

a) Divide the entire simulation period into a number of steps (these steps are different than the simulation time steps). Specify initial geological model(s).

This can be a single model or an ensemble of models and can be conditioned to well data.

b) Use this model(s) in the optimization module to find a single set of valve settings that optimizes recovery (in an average sense for the ensemble case). The module returns the valve settings and the optimized recovery.

c) Apply the valve settings from the optimization module to wells in a reference model (this would be the sub-surface reservoir in real applications and is here assumed to be unknown). Run the reference model for a period of time with these new valve settings to generate data. Pass the generated data to the history matching module.

d) Run the history matching module with the data obtained from the reference model (using the actual valve settings) to generate single or multiple history-matched (or updated) models.

e) Loop through steps (b) to (d) for each optimization step until the entire simulation period is covered.

The initial model is used only for the first step of the entire simulation period. Subsequent steps use updated model(s) obtained from the history matching module. These models are expected to better reflect the characteristics of the actual (unknown) model. Multiple initial models can also be specified. In all cases, a single set of valve settings is generated. These settings optimize the recovery on average for the specified models.

The data generation module is used to simulate the sub-surface reservoir. The target model is treated as unknown. Valve settings determined through optimization are implemented in the wells. We run this true model for a period of time to generate sensor data for use within the history matching module. In the following sections, the two main components (valve control optimization and history matching) of the overall procedure are discussed.

2.4. Valve Control Optimization Module

Due to the nature of the problem, this module utilizes a nonlinear conjugate gradient (CG) algorithm to determine the valve settings that maximize cumulative recovery (or that minimize the negative of the cumulative recovery) for any given model(s). This is achieved by maximizing an objective function defined for the problem. The optimization can be accomplished subject to production and injection constraints. The gradient of the objective function is used to provide the direction (incrementally opened or closed) of the valve settings. In the following sections, we describe the objective function, the CG implementation, and the modifications introduced to incorporate history matching.

2.4.1. Objective Function

The objective function utilized in this work is based on the cumulative oil recovery. NPV, though not considered here, is also a relevant objective function. Using a single geological model, the objective function is stated as:

$$\underset{x_i \in [0,1]}{\text{maximize}}\, F(\mathbf{x}) \quad = \quad f \tag{2.2}$$

where $f \in [0, 1]$ is the recovery factor and $\mathbf{x}$ is the scaled vector of valve settings. For multiple models, the objective function is here stated as (Yeten *et al*., 2003):

$$\underset{x_i \in [0,1]}{\text{maximize}}\, F(\mathbf{x}) \;=\; < f > \;+\; r\,\sigma \tag{2.3}$$

where r is a risk factor and σ is the standard deviation of the recovery over a given number of models. $< f >$ is a weighted average recovery factor defined as:

$$< f > \;=\; \sum_{j=1}^{N} \lambda_j \; f_j(\mathbf{x}) \tag{2.4}$$

N is the total number of realizations considered, and $\lambda \in [0, 1]$ is the weight assigned to each f_j such that $\sum_{j=1}^{N} \lambda_j = 1$. We note that $\mathbf{x}$ (the valve settings) is the same for all N

realizations. The case of $r = 0$ is a risk neutral case and $F(\mathbf{x})$ is then simply a weighted average of the recovery factors. The case of $r = -1$ is a risk averse case where $F(\mathbf{x})$ is also a function of σ (in which case σ will tend to be smaller than in the case of $r \geq 0$). A value of $r = -1$ was used for the work presented in this report.

For the initial model specification, $< f >$ is set to be the arithmetic average of the f_j. Subsequently, λ_j is either set to $1/N$ (equal weighting) or is taken as the rank transform of the inverse objective function value from the history matching module. This latter approach more heavily weighs models that better match the history. In the results, we will indicate which approach was used.

2.4.2. Conjugate Gradient Algorithm

The conjugate gradient (CG) algorithm is a gradient based method that can be used to find the minimum of a quadratic function. The CG algorithm requires the gradient of the function being minimized. For any function approximated by a quadratic function, the gradient is given as:

$$\mathbf{g}(\mathbf{x}) = -\mathbf{b} + \mathbf{H}\mathbf{x} \tag{2.5}$$

where $\mathbf{H}$ is the Hessian matrix and $\mathbf{b}$ is an arbitrary vector. Hence the objective of the CG method is to find the $\mathbf{x}$ that gives $\boldsymbol{g}(\mathbf{x}) \approx 0$. Since the gradient direction is opposite to the direction of descent, the residual, $\mathbf{r}$, is set to the negative of the gradient. For our objective function in the previous section, this becomes:

$$\mathbf{r}^k = -\mathbf{F}'(\mathbf{x}^k) \tag{2.6}$$

where the superscript k indicates the iteration level with $k = 0$ referring to the initial guess. We define the vector $\mathbf{d}$ as the search direction, α is the step size in the search direction and β is the Gram-Schmidt constant. The outline of the CG algorithm used in this work is stated below (this description is from Yeten *et al*., 2002):

Given an initial guess, $\mathbf{x}^0$

Set $\mathbf{d}^0 = \mathbf{r}^0 = -\mathbf{F}'(\mathbf{x}^0)$

For $k = 0$ until $\mathbf{x}^k$ has desired accuracy

$$\begin{aligned}
&\text{find } \alpha^k \text{ that minimizes } -F(\mathbf{x}^k + \alpha^k \mathbf{d}^k) \\
&\mathbf{x}^{k+1} = \mathbf{x}^k + \alpha^k \mathbf{d}^k \\
&\mathbf{r}^{k+1} = -\mathbf{F}'(\mathbf{x}^{k+1}) \\
&\beta^{k+1} = \max\left\{\frac{(\mathbf{r}^{k+1})^T (\mathbf{r}^{k+1} - \mathbf{r}^k)}{(\mathbf{r}^k)^T \mathbf{r}^k}, 0\right\} \\
&\mathbf{d}^{k+1} = \mathbf{r}^{k+1} + \beta^{k+1} \mathbf{d}^k
\end{aligned} \tag{2.7}$$

end

The algorithm is considered converged when

$$\left\| \mathbf{r}^k \right\| < \varepsilon \left\| \mathbf{r}^0 \right\| \tag{2.8}$$

with ε here taken as 0.01.

The Gram-Schmidt constants in Eq. 2.7 are computed using the Polak-Ribiérie (PR) method. An alternative is to use the Fletcher-Reeves (FR) method. The advantage of using the PR method is that it automatically restarts as the steepest descent method if there is not much change in the gradient between iterations. The FR method, on the other hand, may get trapped in a local valley.

The gradient of the objective function is numerically computed using a forward finite difference approximation:

$$\mathbf{F}'(\mathbf{x}) = \nabla F(\mathbf{x}) = \frac{F(\mathbf{x} + h\mathbf{e}) - F(\mathbf{x})}{h} \tag{2.9}$$

where $\mathbf{e} = \{ \mathbf{e}_i \}$ is a set of unit vectors and h is the step size. When the valve settings approach the upper limit (i.e., $x_i \rightarrow 1$), Eq. 2.9 is replaced with the backward difference approximation to avoid unphysical valve settings.

The choice of the derivative step size *h* was investigated using a simple variogram based model with a horizontal well located in the middle. Yeten *et al*. (2002) took the derivative step size to be 0.05, which we will refer to as the default. Here, we compared three different step sizes – 0.01, 0.02 and the default – to observe their effect on cumulative recovery and water cut. The result is shown in figure 2.4 below, which is a plot of both the recovery factor and water cut against time for the different step sizes. As observed in both plots of figure 2.4, taking *h* as 0.01 and 0.02 yielded practically the same result. Both of these sizes gave higher cumulative recoveries than the default size. The effect on water cut is more pronounced than on cumulative recovery. Both the 0.01 and 0.02 sizes gave lower water cuts than the default. For this work, therefore, a step size of $h = 0.02$ was used.

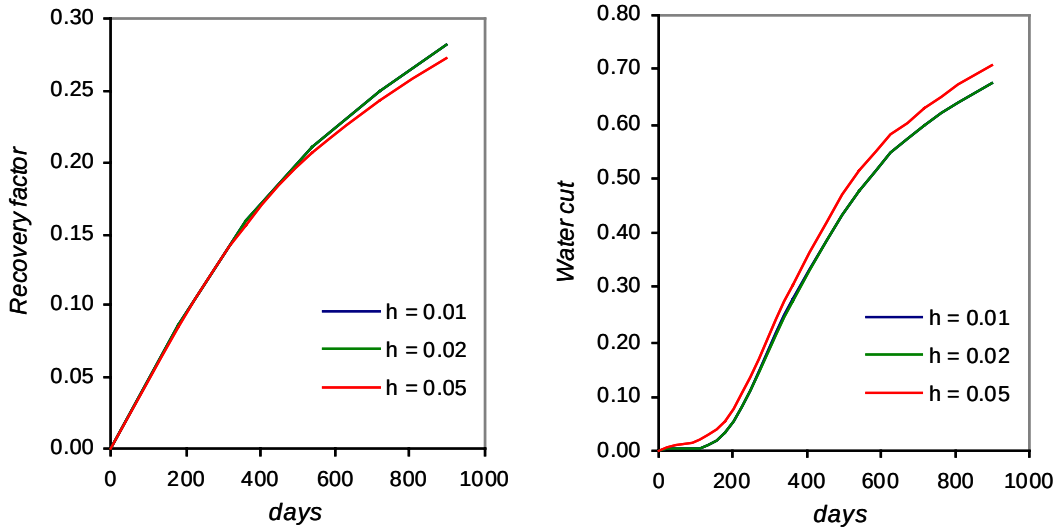

Figure 2.4: Plots comparing recovery factors and water cuts for different values of the step size, *h*.

2.4.3. Previous Implementation of the Defensive Control Strategy

Figure 2.5 below depicts schematically the previous implementation of the defensive strategy by Yeten *et al*. (2002) using the nonlinear CG algorithm as the optimization engine. It shows how the valve settings are optimized in time. The entire simulation period is divided into equal optimization steps at which the valve settings are updated to

optimize the objective function. Valve settings for each optimization step are determined by solving Eq. 2.2 or Eq. 2.3 over the remaining simulation period, represented by the same colour in the horizontal direction in figure 2.5.

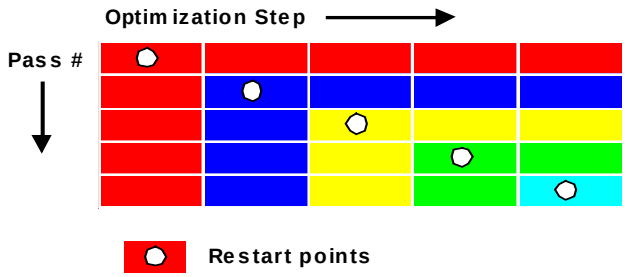

Figure 2.5: Yeten *et al.* (2002) technique for the optimization of valve settings in time.

To proceed to the next optimization step, the simulation is restarted (represented by the circles in the figure) at the end of the previous optimized step (represented vertically by the same colors in the figure). This is continued until the entire simulation period is covered and optimized valve settings are obtained for each step. This "global" optimization approach was applied in order to avoid valve settings at early times that led to deleterious effects at later times. We will explore alternative approaches in Chapter 4.

The key limitation of this technique is that it is based on a fixed geological model throughout the entire process. This approach, however, serves as a basis for the new method that addresses this issue.

2.4.4. Implementation of the Defensive Control Strategy with History Matching

Sensors, which are installed with the valves in a smart well, provide data from each branch or segment of the well. As a result, the geological model can be updated continuously and then used to optimize the valve settings. This is depicted schematically in figure 2.6 below with the nonlinear CG as the optimization engine.

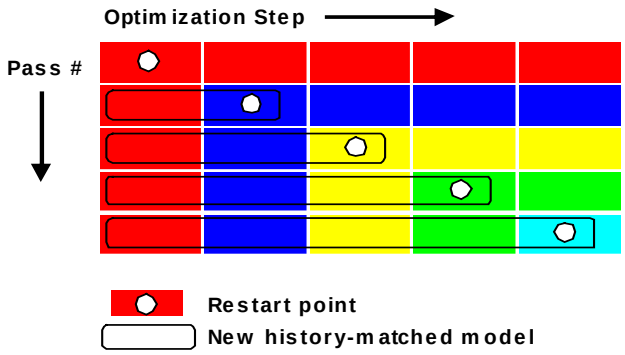

Figure 2.6: Real-time optimization of valve settings with history matching.

The simulation period is again divided into a number of optimization steps. The valve settings are optimized for the first step by solving Eq. 2.2 (or Eq. 2.3 for multiple models) over that step and the remaining simulation period (represented by the same color in the horizontal direction). This is done with the specified initial model (figure 2.3). The optimized valve settings are applied to wells in the reference model (sub-surface reservoir in real applications) to generate data for history matching. Once the model has been updated with the sensor data, it replaces the existing history matched model (shown as the boxes in the figure). The simulation is then restarted by running the updated model with the previously optimized valve settings from the beginning of the entire simulation to the end of the last optimized step. The end of the last optimized step marks the beginning of the next step to optimize with the updated model. This procedure is repeated for each subsequent optimization step until the entire simulation period is covered. This procedure can be summarized as follows:

a) Divide the simulation period into a number of optimization steps at which the valve settings will be updated.

b) Optimize the valve settings for each optimization step such that the objective function is maximized over the current step and the remaining simulation period. Use the initial model for the first step. Subsequent optimization steps use the current history-matched model(s).

c) Generate sensor data (production) from the reference model using the historical and current valve settings.

d) Update the model by history matching the sensor data.

e) Replace the previous model with the new history-matched model.

f) Restart the simulation by running the new history-matched model with previously optimized valve settings from the beginning of the entire simulation to the end of the last optimized step.

g) Repeat step (b) through step (f) until the entire simulation period is covered.

Note that the sensor data from all previous steps are used for the history matching.

2.5. History Matching Module

The module is a combination of two approaches within the probability-perturbation procedure. It applies a facies-based procedure that uses multiple-point (mp) statistics (training images) to populate the geologic model. The approach is capable of generating multiple history-matched models. The following sections describe the probability-perturbation technique and the objective functions that the procedure seeks to minimize. The module is demonstrated by applying it to a variogram-based model and to a channel reservoir case.

2.5.1. Probability-Perturbation Method

This method, proposed by Caers (2003), allows geologic models to be conditioned to production data. It uses multiple-point statistics (through the use of training images) to generate intricate reservoir structures that could not be generated using variogram (two-point geostatistical) methods. The mp statistics can, however, be used to reproduce structures from the variogram method.

Consider the case of a binary random function defined as:

$$I(\mathbf{u}, s_l) = \begin{cases} 1 \ \textit{if a certain facies } s_l \textit{ occurs at location } \mathbf{u} \\ 0 \ \textit{otherwise} \end{cases} \tag{2.10}$$

The event $I(\mathbf{u}, s_l) = 1$ is denoted by A while D indicates the production data event and B is the training image event. The usual model generated by conditioning to B only is the probability of A given B, written as $P(A|B)$. Therefore, the method seeks to perform the transformation:

$$P(A \mid B) \xrightarrow{\text{perturb with } P(A|D)} P(A \mid B, D) \tag{2.11}$$

where $P(A|D)$ is defined as:

$$P(A \mid D) = (1 - r_D)\, i^{(k-1)}(\mathbf{u}, s_l) + r_D\, P(A) \tag{2.12}$$

$P(A)$ is the global proportion of each facies, s_l, within the model and is assumed to be known. The conditional probability $P(A|D)$ is modeled with a free parameter $r_D \in [0, 1]$ that quantifies the change of the current simulated value $i(\mathbf{u}, s_l)$ required to match the data D. This free parameter r_D is constant over $\mathbf{u}$ and is obtained using a 1D optimization method such as the Brent method (Press *et al.*, 2002). Starting from an initial model, $i^{(0)}(\mathbf{u}, s_l)$, subsequent models, $i^{(k)}(\mathbf{u}, s_l)$, are generated iteratively until the simulated data matches the data D. The `SNESIM` algorithm (Strebelle, 2000) is used to generate $P(A|B)$. It is also used to generate $P(A|B,D)$ with $P(A|D)$ and a training image as input.

2.5.2. Objective Functions

The module performs two levels of optimization if the petrophysical properties (permeability and porosity) are to be optimized along with the facies distribution. The two levels are the outer and the inner loops. For the case where the petrophysical properties are assumed known, the module performs only the outer loop. The outer loop is a 1D facies optimization (as described above) to determine the parameter r_D that minimizes the outer objective function:

$$\underset{r_D \in [0,1]}{\text{minimize}}\, g(r_D) = \sum_j \left(D_j(i^{(k)}_{r_D}(\mathbf{u}, s_l)) - D_{obs,\,j} \right)^2 \tag{2.13}$$

where D is the model production data and D_{obs} is the observed production data.

The inner loop is a multivariate optimization of the permeability and porosity distributions. Basically, the mean and standard deviation of the distributions are determined to minimize the inner objective function:

$$\underset{0 \le \alpha_i \le 1}{\text{minimize}}\ f(\boldsymbol{\alpha}) \;=\; \sum_j \left(D_j(\boldsymbol{\alpha}) \;-\; D_{obs,\,j}\right)^2 \tag{2.14}$$

The multivariate parameter $\boldsymbol{\alpha} = \{\mu_{log\,k}, \mu_\phi\}$ has all of the elements scaled by maximum and minimum input values of the $\mu_{log\,k}$ and μ_ϕ. This inner objective function is minimized with the nonlinear Levenberg-Marquardt method (Suzuki, 2003; Press *et al*., 2002). A nonlinear CG method can also be used to minimize this objective function.

The advantage of installing sensors at each branch is to provide data for each lateral. Therefore, the history matching module matches the model to well pressure, cumulative oil production and water cut for each branch of the well. It is also possible to match only the overall well water cut, though this does not guarantee that the water cut from each of the branches will be matched.

2.5.3. History Matching Algorithm

The history matching algorithm can be summarized as follows:

- Generate an initial realization $i^{(k)}(\mathbf{u}, s_l)$, $l = 1, \ldots, L, \forall \mathbf{u}$, $k = 0$, with `SNESIM` conditioned to the training image event given $P(A)$.
- Perform 1D facies optimization to determine the r_D that minimizes the outer objective function while fixing the permeability and porosity distribution.
- Fix the facies distribution and perform multivariate permeability-porosity optimization to determine the $\boldsymbol{\alpha}$ that minimizes the inner objective function.
- Repeat all of the above points for $k = 1, \ldots, K$ and keep those with the least overall objective function to generate multiple history-matched models.

2.5.4. Application of the History Matching Module to Synthetic Models

In this section, our focus is on the quality of the history match produced by the HM module. In order to demonstrate this, we applied the overall valve optimization and history matching procedure to two synthetically generated models. The first model (M1) is a simple variogram based model while the second model (M2) is a more complicated channel reservoir. Both models are two facies (low and high permeability sands) systems.

Variogram based model, M1:

The model parameters are given in Table 2.1. M1 contains only oil and an analytical aquifer acts at its bottom edge to maintain pressure (two-phase flow). The facies were unconditionally generated using indicator geostatistics (Deutsch and Journel, 1998). Dimensionless correlation lengths (l_x, l_y, l_z) for the low and high permeability facies were (0.1, 0.1, 0.4) and (0.4, 0.2, 1.0), respectively. The permeability within each facies was independently and unconditionally populated with sequential Gaussian simulation, `SGSIM` (Deutsch and Journel, 1998). The top view of the model is shown in figure 2.7 and the statistics of the permeability distribution for each facies are summarized in table 2.2. A single horizontal well was introduced into the model. The well has five segments and each was instrumented with a valve and sensor (shown in figure 2.7 as red dots on the white line centred in the middle).

Constant liquid rate control was specified for the well with a minimum bottomhole pressure of 1,500 psi. The initial liquid production rate was 6 MSTB/day. The simulation proceeded for 960 days and the valve settings were updated every 240 days. This implies that the entire simulation period was divided into four optimization steps. The intent here is to demonstrate the quality of the history match obtained.

Table 2.1: Simulation Model Properties for M1

drainage area	3000 x 3000 ft^2
oil thickness	50 ft
grid dimension	50 x 50 x 5
ϕ_{ave}	0.19 (σ = 0.03)
p_i	4000 psi
c_o at p_i	15.00 x 10^{-6} psi^{-1}
c_w at p_i	3.04 x 10^{-6} psi^{-1}
k_{ro}	0.80 at S_{wc} = 0.20
k_{rw}	0.40 at S_{or} = 0.30
μ_o *at* p_i	4.00 cp
μ_w *at* p_i	0.30 cp
B_o at p_i	1.15 RB/STB
B_w at p_i	1.02 RB/STB
k_v/k_h	0.1

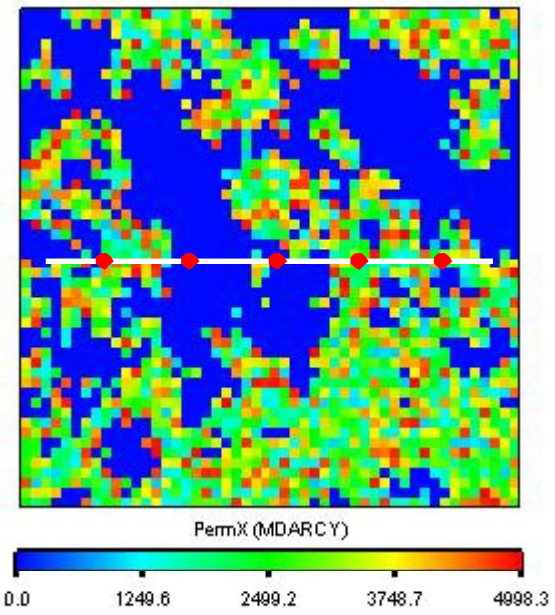

Figure 2.7: Permeability field of variogram-based model M1.

Table 2.2: Model M1 Permeability Statistics

Facies	Average (md)	Standard Deviation (md)
High Perm Sand	2933	1178
Low Perm Sand	5.04	2.86

Figure 2.8 shows the cumulative oil and water cut plots for the M1 and HM models. As shown in the figure, the final HM model showed good agreement with the actual data (circles). Note that the initial model did not provide a close match to the actual data.

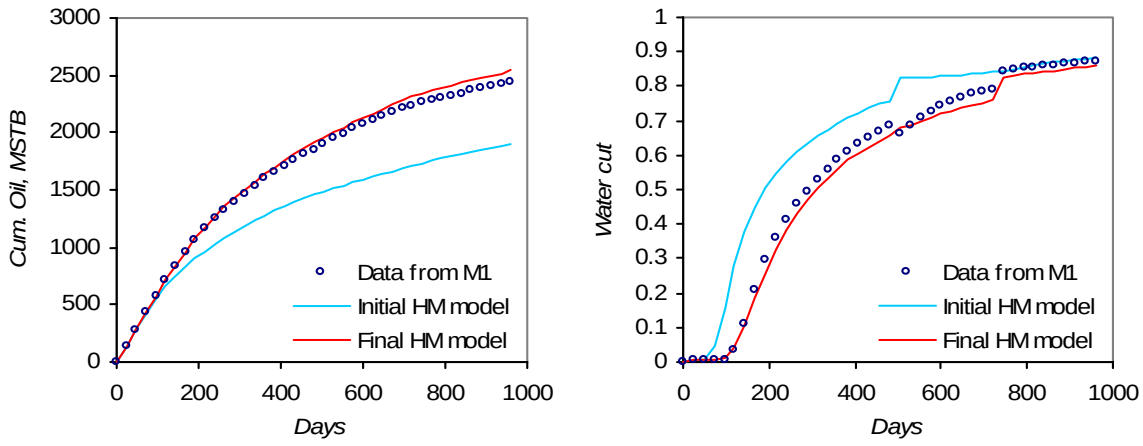

Figure 2.8: History match plot for M1 model.

Channel reservoir case, M2:

The model parameters for this case are given in Table 2.3. M2 contains oil and gas with an analytical aquifer acting at its bottom edge (three-phase flow). The SNESIM algorithm was used to unconditionally generate the facies model with a training image. Permeability within each facies was independently and unconditionally populated with SGSIM. Table 2.4 summarizes the statistics of the permeability distribution for each facies. The model is shown in figure 2.9 below. A quad-lateral well completed about 15 feet from the oil-water contact was introduced into the model. A valve and sensor were

installed at the junction of each of the laterals and the mainbore. The mainbore was not completed. Constant total fluid rate control was specified for the well with a minimum bottomhole pressure of 1,500 psi. The initial production was specified at a total liquid rate of 10 MSTB/day. The simulation proceeded for 800 days and the valve settings were updated every 200 days. Again, the intent here is to show the adequacy of the history match module.

Table 2.3: Simulation Model Properties for M2

drainage area	4000 x 4000 ft^2
oil thickness	50 ft
gas cap thickness	50 ft
grid dimension	20 x 20 x 6
p_b	4000 psi
R_s at p_b	1.0 MSCF/STB
c_w at p_b	3.04 x 10^{-6} psi^{-1}
k_{ro}	0.80 at S_{wc} = 0.20, S_{gr} = 0.0
k_{rw}	0.40 at S_{or} = 0.30
k_{rg}	0.90 at S_{wc} = 0.20
γ_o at 14.7 psi	0.85
γ_w at 14.7 psi	1.0
γ_g at 14.7 psi	0.71
μ_o *at* p_b	0.42 cp
μ_w *at* p_b	0.30 cp
μ_g *at* p_b	0.025 cp
B_o at p_b	1.55 RB/STB
B_w at p_b	1.02 RB/STB
B_g at p_b	0.71 RB/MSCF
k_v/k_h	0.1

Table 2.4: Model M2 Permeability Statistics

Facies	Average (md)	Standard Deviation (md)
Channel Sand	2478	911
Mudstone	5.32	1.02

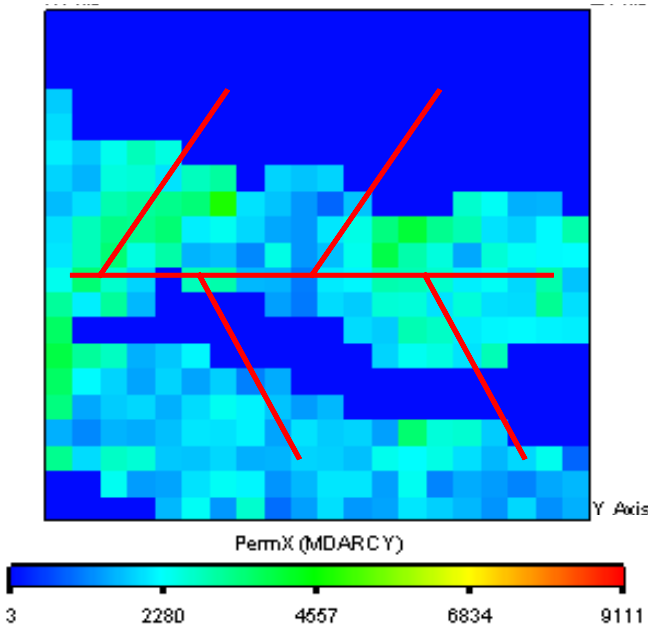

Figure 2.9: Permeability field of channel reservoir M2.

The simulation results are shown in figures 2.10 and 2.11. Figure 2.10 displays the cumulative oil plot of data from M2 compared with that from the initial and final history-matched models. The figure shows that the cumulative oil from the final history-matched model falls on top of the actual data (plotted as circles). In figure 2.11, the water cut of the final HM model tracks the actual data reasonably well. Notice in both figures that the initial HM model did not match the actual data.

In summary, both example cases considered showed reasonable overall agreement between the actual models and the final history-matched models. We take this as verification that the HM procedure is suitable for our work. In the next chapter, we will present the application of the overall procedure to other variogram-based and channel reservoir models.

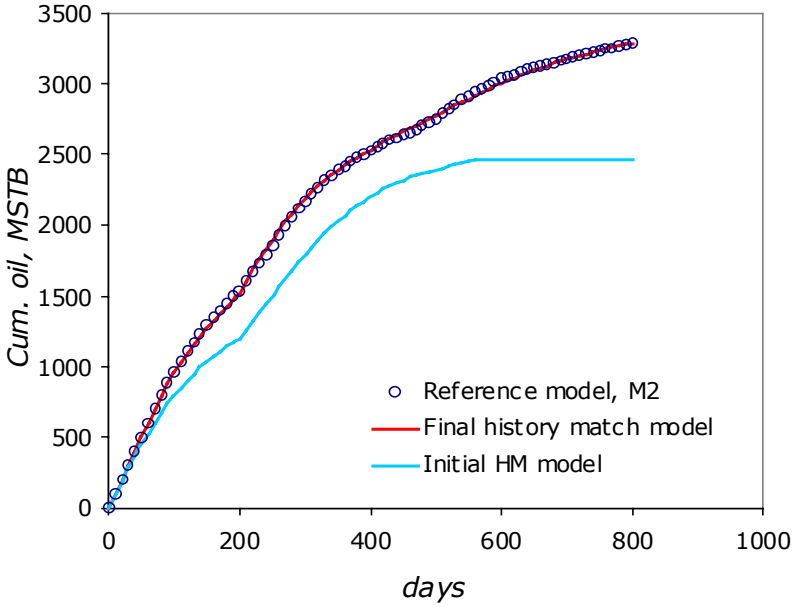

Figure 2.10: History matching plot of cumulative oil data of M2.

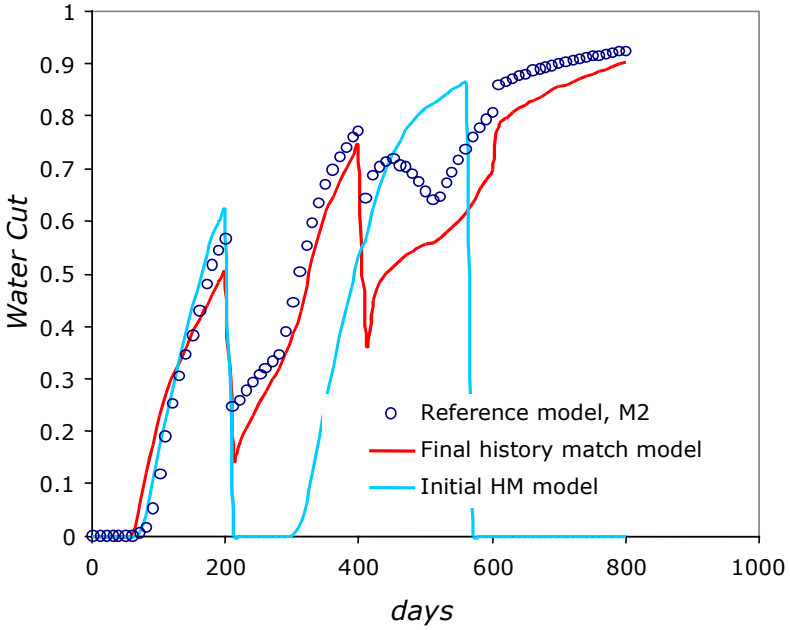

Figure 2.11: History matching plot of water cut data of M2.

Chapter 3

3. Application of the Overall Procedure

We now apply the overall procedure (valve optimization and history matching) to some synthetic models. First, the procedure is applied to geological models that are not conditioned to any well data. Next, we apply it to geological models that are conditioned to hard data (well data) events. The conditioning data was prescribed along the well trajectory (mainbore and laterals of a multilateral well). This is reasonable because the well and its branches are usually logged while drilling with `LWD`. In all of the cases considered, we demonstrate the benefits of instrumentation by comparing the results of the instrumented and the uninstrumented cases for particular geological models.

3.1. Unconditioned Models

We consider an unconditioned variogram based model with a segmented horizontal well and two different unconditioned channelized systems with quad-lateral wells. In cases where multiple models were used, Eq. 2.3 (valve optimization module) was computed with equal weighting ($\lambda_j = 1/N$) of the recovery factors.

3.1.1. Variogram Based Model

This model contains only oil. An analytical aquifer acts at the bottom edge of the model to maintain pressure (two-phase flow). The model parameters are as given in Table 2.1 above. The mobility ratio is 6.7. The facies were unconditionally generated using indicator geostatistics. The permeability within each facies was independently and unconditionally populated with sequential Gaussian simulation, `SGSIM` (Deutsch and Journel, 1998). The top view of the model is shown in figure 3.1. The statistics of the permeability distribution for each facies is the same as was used for the model M1 in the previous chapter (summarized in table 2.2), though here we consider a different realization. The porosity was uniform ($\phi_{ave} = 0.19$). A single horizontal well was

introduced into the model. The well is located 10 ft from the top of the model and has a total length of 2400 ft. The well has five segments and each is instrumented with a valve and sensor (shown in figure 3.1 as red dots).

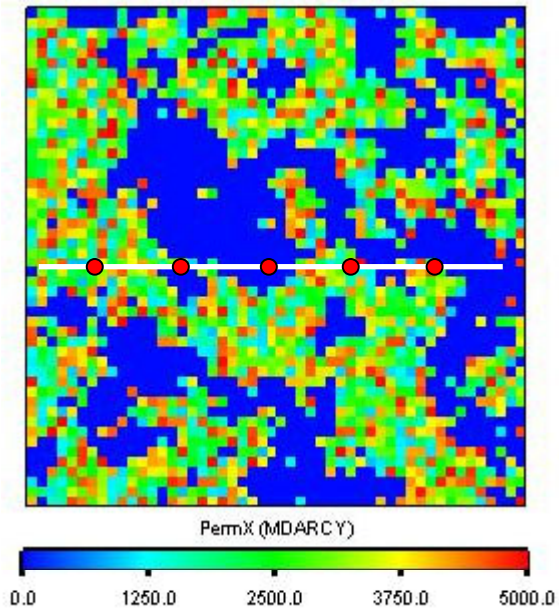

Figure 3.1: Permeability field of variogram based model.

Constant liquid rate control was specified for the well with a minimum bottomhole pressure of 1,500 psi. The initial / maximum production was specified at a total liquid rate of 6 MSTB/day. The simulation proceeded for 720 days and the valve settings were updated every 240 days. This implies that the entire simulation period was divided into three optimization and history matching steps.

Known Geology Case (no history matching):

We first ran this geological model with no instrumentation on the well (base case), which did not require any optimization. Since the model is known, history matching is also not required. For this uninstrumented base case, the horizontal bore was completed with 7-inch diameter slotted casing and was segmented. Next, we considered the well to be a smart well with five segments. In addition to completing the bore as before, 3½-inch

production tubing was installed in the well. Five control devices, one per segment, were introduced on the tubing so that each of the segments can be controlled independently. The uninstrumented and instrumented cases were compared to demonstrate the benefits of the control devices. In the instrumented case, the objective function to be optimized was cumulative oil produced.

Figures 3.2 and 3.3 show the cumulative oil plot and the water cut plot, respectively, for the base (without valves) and optimized (with valves) cases. The water cut plot of figure 3.3 shows that the uninstrumented case experienced very early water breakthrough. The optimization of the valves tends to delay water breakthrough, leading to a substantial increase in oil recovered. Specifically, the percentage gain in cumulative oil recovered over the no-valve base case was approximately 58%.

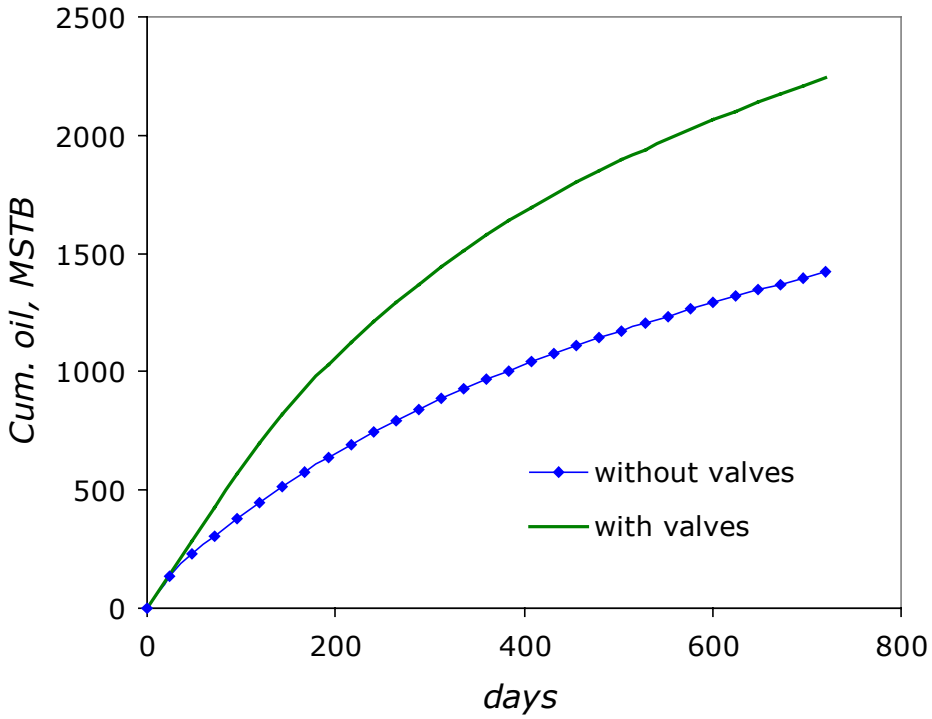

Figure 3.2: Cumulative oil plot of the base (without valves) and the optimized (with valves) cases for known geology.

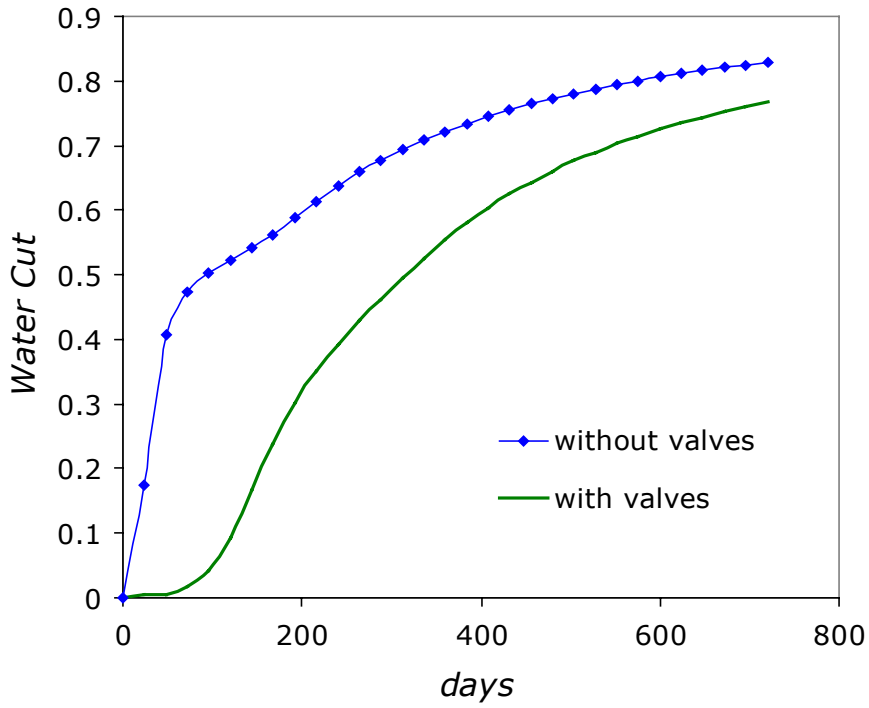

Figure 3.3: Water cut plot of the base (without valves) and the optimized (with valves) cases for known geology.

Valve Optimization with History Matching (VOHM):

Since the geology is not known in real applications, history matching (or some type of model updating) is necessary for valve optimization to be beneficial in practice. We therefore applied the overall VOHM procedure, described in section 2.4.4, to this case. We generated an initial geologic model that was not conditioned to any data and used it to determine optimized valve settings. This model was specified to have the correct variogram model and the correct global proportions of the facies. The valve settings were applied to the well in the reference model after which flow simulation was performed to generate production data for history matching. The history-matched models were subsequently used during the other optimization and HM steps.

Figure 3.4 is a plot of the HM cumulative recovery together with the cumulative recovery curve of the known geology base case and the optimized (known geology) case. The HM recovery curve (red curve) lies between the base case and the optimized case curves and

is quite close to the known geology optimized case. It should be noted that the known geology optimized case is presumably the best that can be obtained, while the base case is presumably the least that would be achieved. Figure 3.5 displays the water cut responses for the three cases. A sharp discontinuity occurs at about 240 days on the HM water cut curve. This occurs because the initial model used in the first step has little resemblance to the actual model and was therefore unable to significantly delay water breakthrough. At the second step, the history-matched model was used to update the valve settings. This caused some of the valves to close, resulting in the discontinuity observed. The percentage gain in cumulative oil recovered over the no-valve base case due to VOHM was approximately 54%. This translates to about 93% of the gain attainable if the geology were known, which, for this case, is 58%.

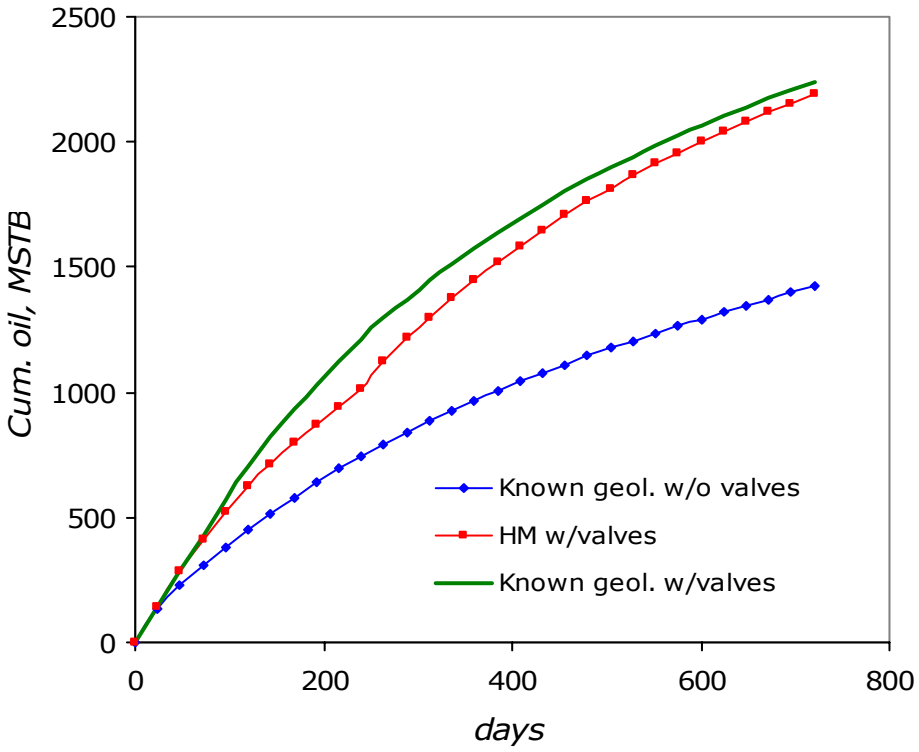

Figure 3.4: Cumulative oil plot from the overall procedure (VOHM) and the known geology cases.

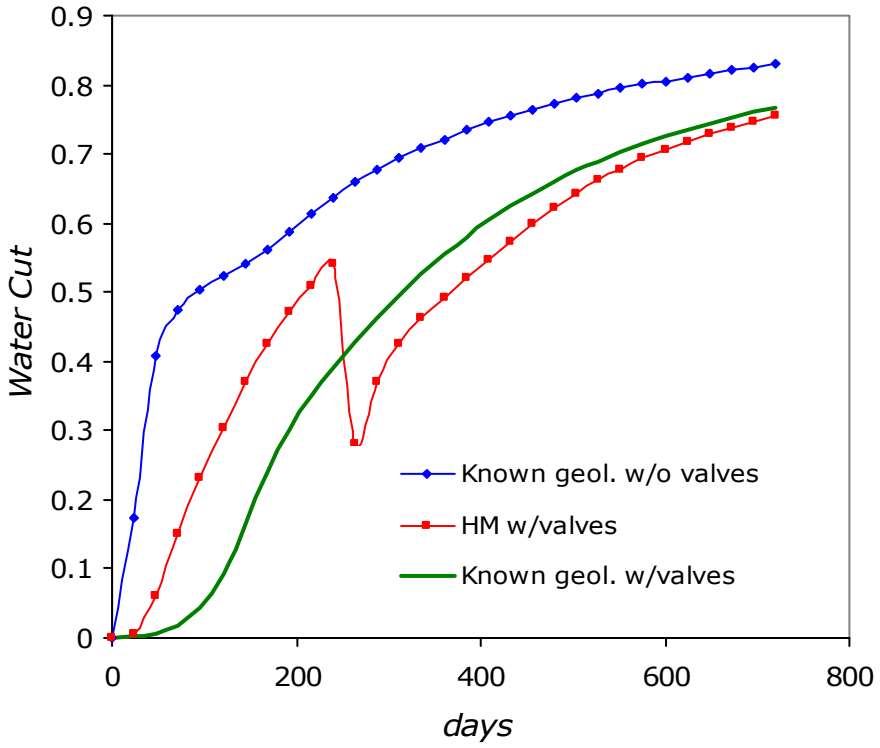

Figure 3.5: Water cut plot from the overall procedure (VOHM) and the known geology cases.

3.1.2. Channel System

We now consider a three-dimensional channelized model (`fluva`) in three-phase flow. The model parameters are as presented earlier in table 2.3. This model has a 50-ft gas cap with an analytical aquifer acting at its bottom edge to maintain pressure. The mobility ratio is less than 1.0. The `SNESIM` algorithm was used to unconditionally generate the facies from a specified training image (this training image is different from that used to generate model M2 in the previous chapter). Independent and unconditional population of the permeability and porosity within each facies was then performed with `SGSIM`. Permeability and porosity histograms are shown in figure 3.6 while the statistics are presented in table 3.1. The bottom view of the permeability field is shown in figure 3.7.

A quad-lateral well located 20 ft from the oil-water contact was introduced into the model. The mainbore is of a total length of 3200 ft. Each of the laterals has a length of

approximately 2200 ft. Only the laterals were opened to production; the mainbore was not perforated.

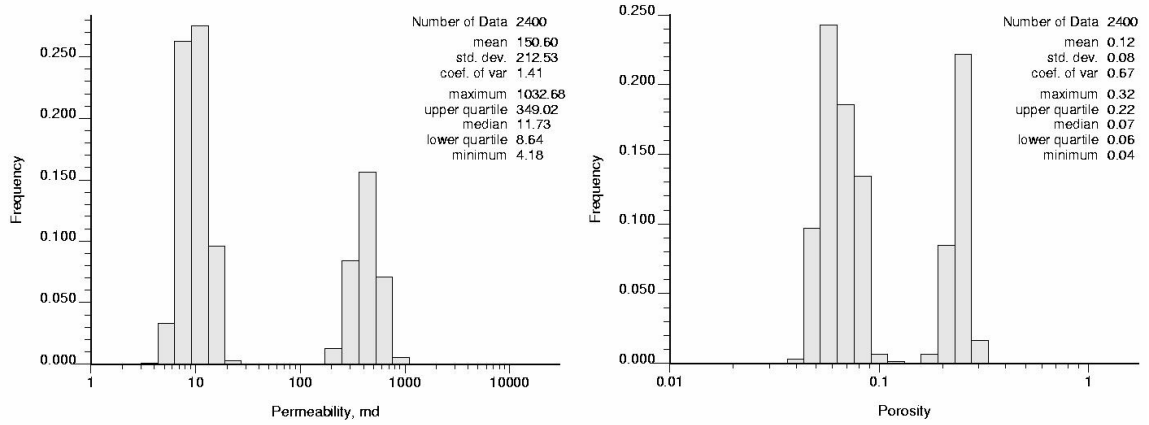

Figure 3.6: Permeability and porosity histograms for `fluva` channel reservoir.

Table 3.1: `fluva` Permeability and Porosity Statistics

Facies	Permeability		Porosity	
	Average (md)	Standard Deviation (md)	Average	Standard Deviation
Channel Sand	436	123	0.24	0.02
Mudstone	10.1	2.85	0.07	0.01

Constant total fluid rate control was specified for the well with a minimum bottomhole pressure of 1,500 psi. Initial / maximum production was specified at a total liquid rate of 10 MSTB/day. The simulation proceeded for 800 days and the valve settings were updated every 200 days; i.e., the simulation period was divided into four optimization and history matching steps. The entire well was shut-in if the producing water cut exceeded 80% or if the oil rate fell below the economic rate of 200 STB/day.

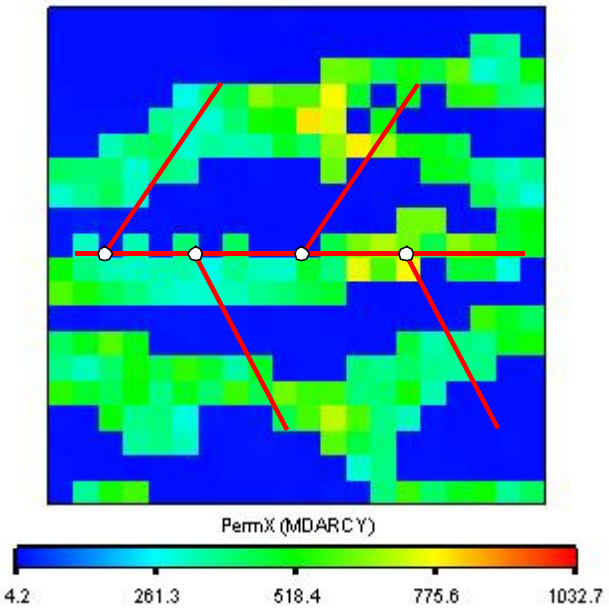

Figure 3.7: Permeability field of `fluva` channel reservoir.

Known Geology Case (no history matching)

As with the variogram based model, we first simulated the channelized geological model with no instrumentation on the well (base case). Since the model is known, history matching is not required. For this uninstrumented base case, the mainbore was completed with 7-inch diameter casing and each of the laterals was completed as 5-inch diameter open holes. For the smart well simulation (optimized case), 3½-inch production tubing was run inside the 7-inch casing. Four control devices (valves and sensors) were installed on the tubing. Each device was placed close to the junction of the branch and the mainbore. Lateral completions were the same as in the base case. We again compare the uninstrumented and instrumented cases to demonstrate the benefits of the control devices.

Figure 3.8 shows the cumulative oil produced for both the base and optimized cases. The no-valve base case stops producing at 420 days because the well attained the water cut constraint of 80% and was therefore shut-in. This is apparent in figure 3.9, which presents the water cut plots of both the base and optimized cases. Breakthrough occurred at about 100 days for both cases but the base case water cut rose quickly to 80% (at which

point the well was shut-in). Notice the sharp discontinuities, which are the effect of valve optimization, in the water cut curve of the optimized case. The valves shut off branches with high water production and/or low oil rate. As a result of valve optimization, the percentage increase in cumulative oil over the base case for this channel reservoir is approximately 40%. This increase is presumably close to the best that can be obtained since we are optimizing the actual model.

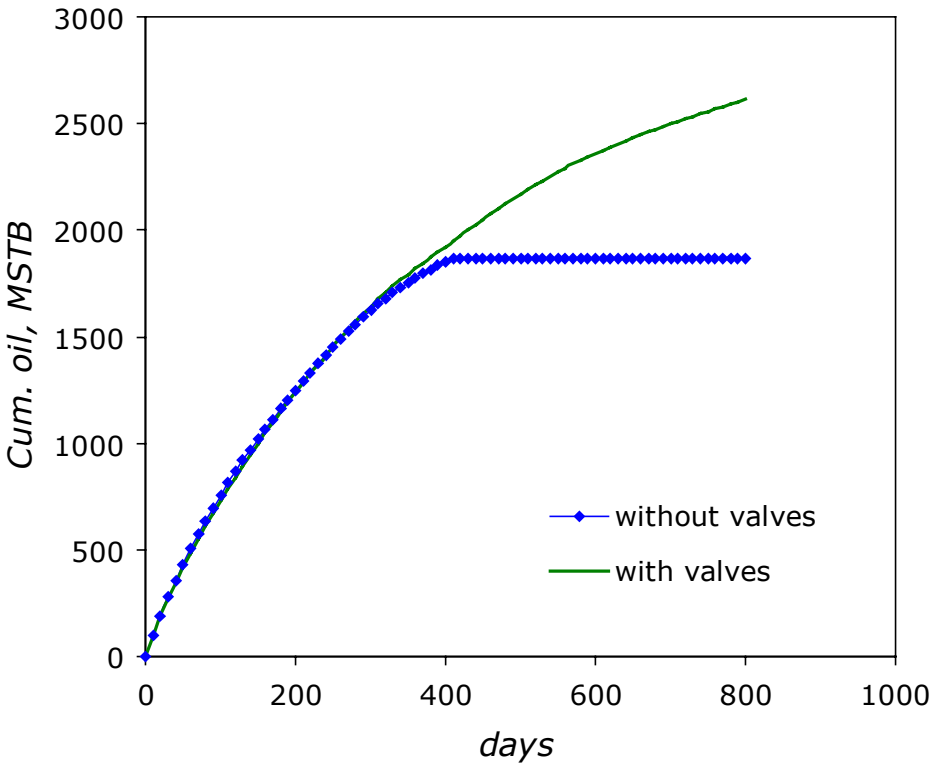

Figure 3.8: Cumulative oil plot of the base (without valves) and the optimized (with valves) cases for known geology (`fluva` channel reservoir).

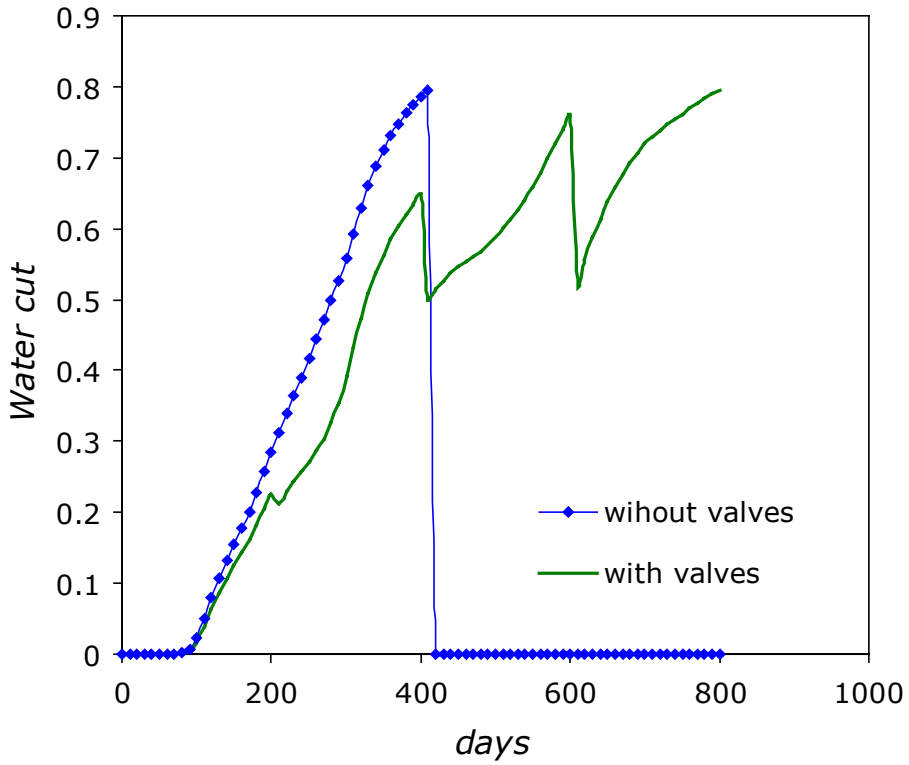

Figure 3.9: Water cut plot of the base (without valves) and the optimized (with valves) cases for known geology (`fluva` channel reservoir).

Valve Optimization with History Matching (VOHM):

We now apply the overall VOHM procedure to this channel reservoir. The initial geological model was based on the same training image used to generate the actual model. Figure 3.10 is a plot of the HM cumulative recovery together with the cumulative recovery curve of the known geology uninstrumented base case and the optimized case. The HM recovery curve (red curve) is very close to the known geology optimized case, as was observed with the variogram case considered above. In figure 3.11, which is the plot of the water cut responses, discontinuities are observed in the HM water cut curve. These are due to the closure of branches not contributing to oil recovery; typically branches producing at high water cut. The percentage gain in cumulative oil recovered over the no-valve base case due to VOHM was approximately 38%. This translates to about 94% of the gain attainable if the geology were known, which, for this case, was 40%.

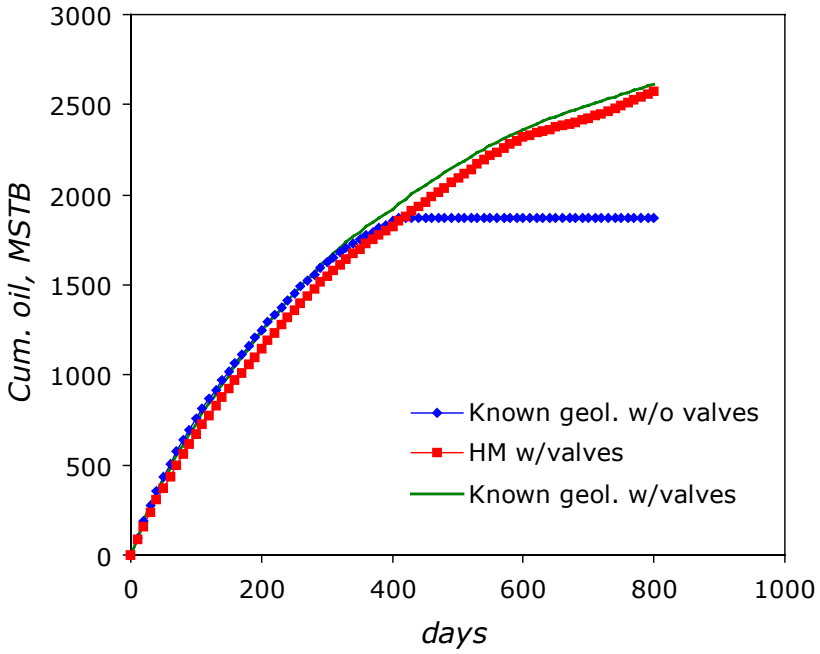

Figure 3.10: Cumulative oil plot from the overall procedure (VOHM) and the known geology cases.

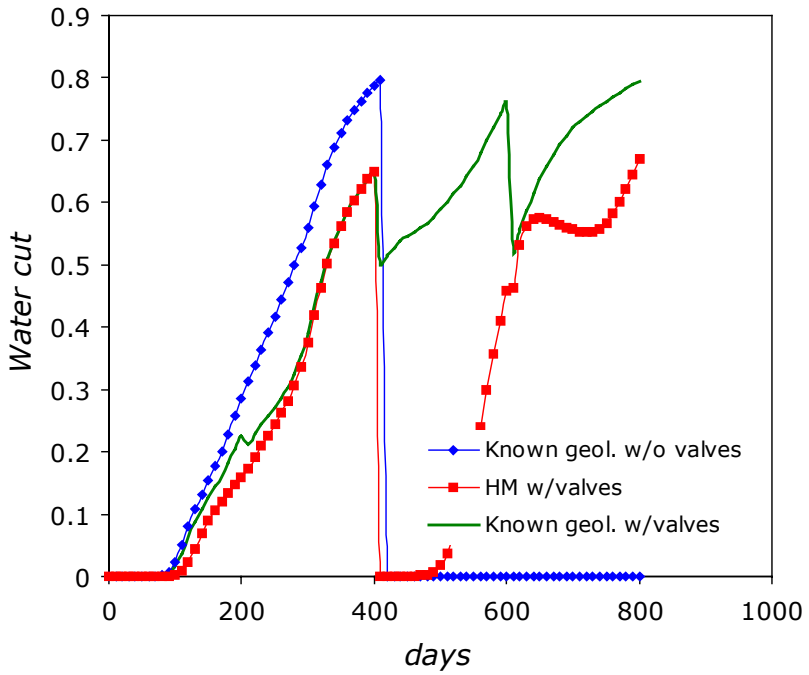

Figure 3.11: Water cut plot from the overall procedure (VOHM) and the known geology cases.

3.1.3. Second Channelized Case

We consider another two-facies channelized reservoir model (`fluvb`). The model is similar to the `fluva` channelized model except that we generated `fluvb` with a different training image. The permeability and porosity within each facies were unconditionally populated and uniformly distributed. Permeability and porosity histograms are shown in figure 3.12 while the statistics are presented in table 3.2. The quad-lateral well and the top view of the permeability field are shown in figure 3.13. The well is in the same location as in the previous example (`fluva`). We also specify the same initial production rate, well and production constraints as were used in the previous channelized reservoir case.

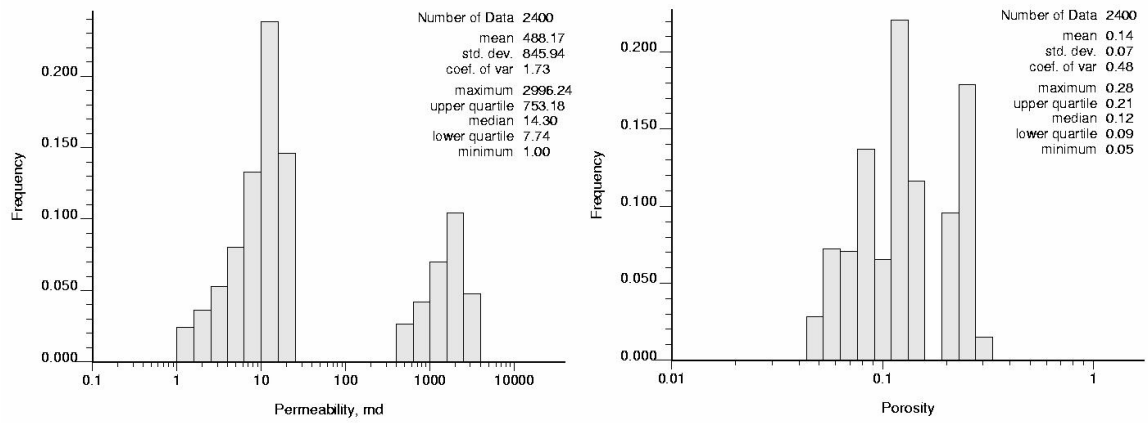

Figure 3.12: Permeability and porosity histograms for `fluvb` channel reservoir

Table 3.2: `fluvb` Permeability and Porosity Statistics

Facies	Permeability		Porosity	
	Average (md)	Standard Deviation (md)	Average	Standard Deviation
Channel Sand	1658	736	0.24	0.02
Mudstone	10.5	5.5	0.10	0.03

Assuming `fluvb` reservoir geology is known, we first ran the model with no instrumentation on the well (base case). Next, we simulated the quad-lateral smart well and ran the model (optimized case). We compared both of these cases to again demonstrate the benefits of the control devices and for baseline purposes.

Figure 3.14 shows the cumulative oil plot for both the base and optimized cases while figure 3.15 presents the corresponding water cut plot. The cumulative oil for the no-valve base case reached a plateau because the well reached the water cut constraint and was therefore shut-in. This is revealed in the water cut plot (figure 3.15). The gain in cumulative oil over the base case due to valve optimization is about 28%.

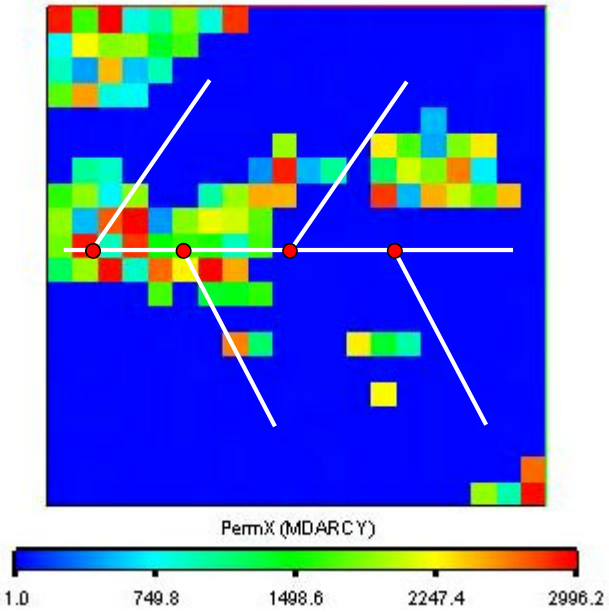

Figure 3.13: Permeability field of `fluvb` channel reservoir.

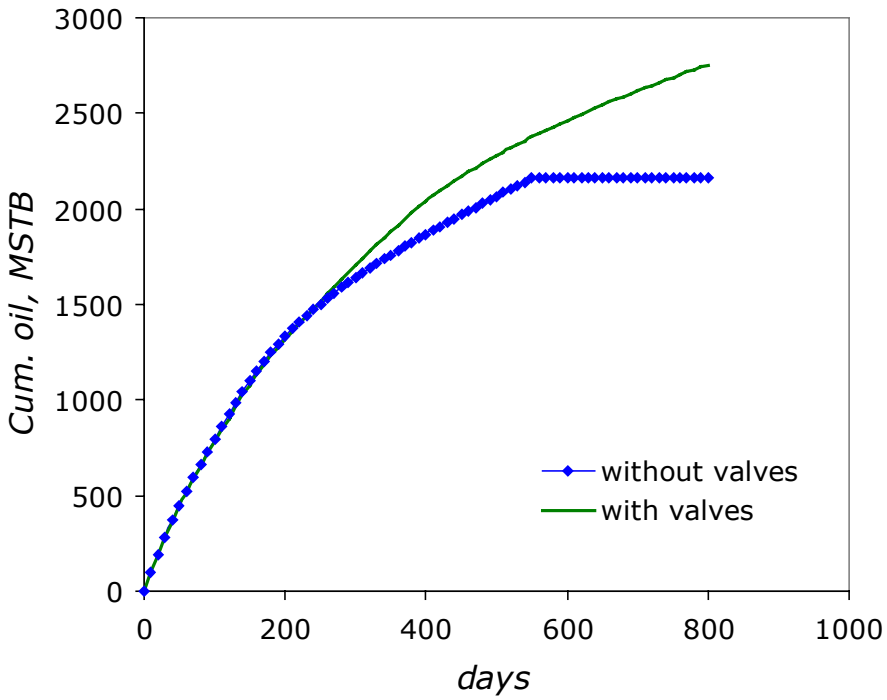

Figure 3.14: Comparison of cumulative oil for known geology case for `fluvb` model.

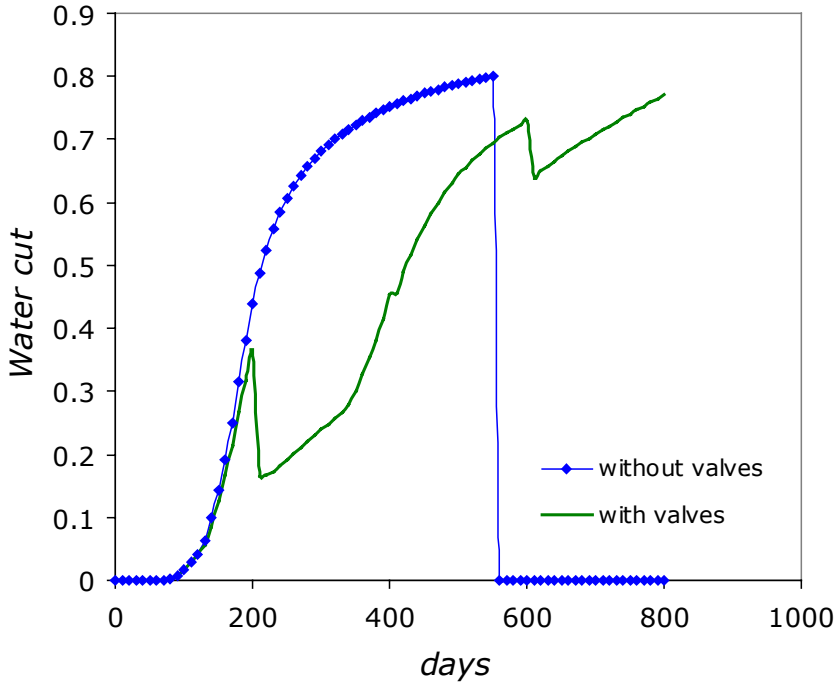

Figure 3.15: Comparison of water cut curves for known geology case for `fluvb` model.

We again apply the overall valve optimization and history matching (VOHM) procedure to this channel reservoir. Figure 3.16 displays the cumulative recoveries for the HM model, the known geology base case, and the optimized (known geology) case. Unlike the `fluva` case, the HM recovery curve is not close to the known geology optimized case. The final recovery was however higher than the known geology base case. The percentage gain in cumulative oil recovered over the no-valve base case due to VOHM was approximately 11%. This is about 41% of the gain attainable if the geology were known, which, for this case, is 28%. Notice how the recovery curve is consistently below the uninstrumented base case until about 600 days. Figure 3.17 shows the water cut responses. Here we see that water cuts in the HM case are quite high compared to the optimized case with known geology.

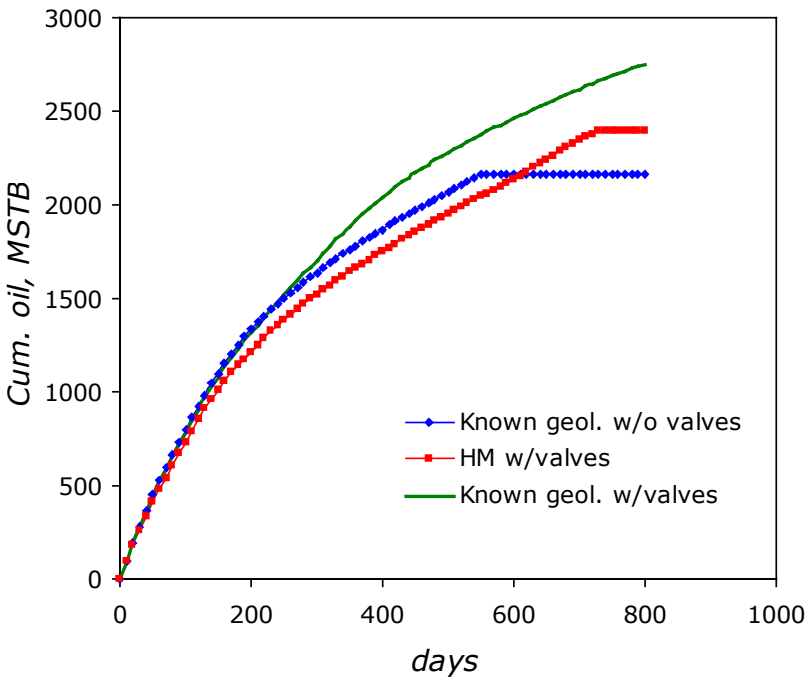

Figure 3.16: Cumulative oil plot from the overall procedure (VOHM) and the known geology cases for `fluvb` model.

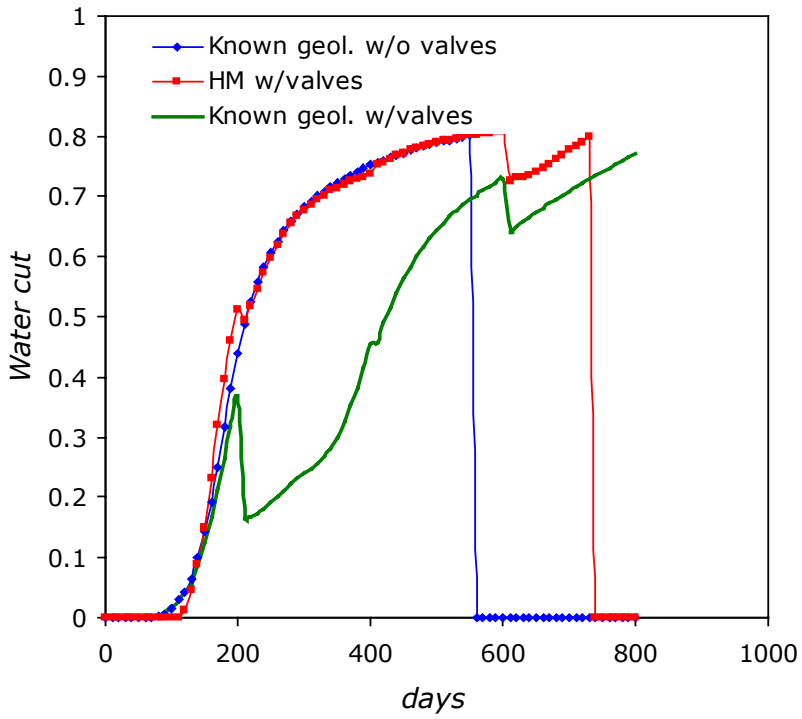

Figure 3.17: Water cut plot from overall procedure and the known geology cases for `fluvb` model.

These results are clearly inferior to those obtained for the previous cases. The problem is likely due to the fact that the history matched model does not adequately represent the reservoir geology. In order to investigate this issue, several (six) different single HM realizations were used (one at a time). There was in general little improvement when compared to the case of just using the initial geologic model without HM. On average, about 44% (± 27%) of the gain attainable with known geology was realized with single HM models. Using the initial geologic model only, the gain obtained was about 39% (± 51%) of that with the actual (known) geology.

This unimpressive result is due, at least in part, to the nonuniqueness of the history matched models. The inaccuracy in any of the individual history-matched models renders them incapable of providing optimization results comparable to those achieved for known geology. We address this problem by optimizing the valve settings over multiple updated models. We consider three and five history-matched models to gauge the sensitivity of

the results. For comparison purposes, we also optimize the valve settings using the initially specified geologic models only (i.e., optimize over multiple models without HM).

In order to quantify the improvement in recovery relative to the gain with known geology, we define a dimensionless cumulative oil difference, ΔN, as:

$$\Delta N = \frac{N_{p_{target\ model,\,w/valve}} - N_{p_{known\ geology,\ no\ valve}}}{N_{p_{known\ geology,\ w/valve}} - N_{p_{known\ geology,\ no\ valve}}} \tag{3.1}$$

where N_p is cumulative recovery. If $\Delta N = 0$, the cumulative oil recovered from the target model will be the same as that of the known geology no-valve base case. On the other hand, if $\Delta N = 1$, the cumulative oil recovered from the target model will be the same as that of the known geology optimized case (with valves).

Table 3.3 shows the result for one, three, and five HM models. The mean and standard deviation are also reported. These were computed by considering many groups (e.g., 21 for the case without HM) of one, three or five models (with HM, we used six, five and two groups, respectively). Note that we do not report standard deviations for the case of five history-matched models because only two such groups were considered, due to the computational costs of these runs. For the case of optimization without HM, ΔN remained at approximately 0.4 for the single and multiple realizations. All of the cases gave high standard deviations, indicating considerable variation between realizations.

Table 3.3: Result of sensitivity runs on `fluvb` for different number of models

Number of realizations	without HM		with HM	
	ΔN	Number of groups	ΔN	Number of groups
1	0.393 ± 0.508	24	0.438 ± 0.273	6
3	0.417 ± 0.372	21	0.852 ± 0.165	5
5	0.358 ± 0.410	18	0.844	2

Using multiple history matched models, ΔN improved significantly, from about 0.44 (with one model) to about 0.85. This clearly demonstrates the potential benefit of optimizing over multiple HM models. The standard deviation was reduced in this case (for one and three models) by about 50% compared to the case of optimizing without HM. Both the increase in ΔN and the reduction in standard deviation can be attributed to the use of multiple realizations and history matching.

A similar sensitivity was performed on the `fluva` channel model (figure 3.7) considered previously. The results are presented in table 3.4. For the case of optimizing without HM, ΔN ranged between 0.5 and 0.6 depending on the number of models considered.

Table 3.4: Result of sensitivity runs on `fluva` for different number of models

Number of realizations	without HM		with HM	
	ΔN	Number of groups	ΔN	Number of groups
1	0.519 ± 0.255	23	0.903 ± 0.179	4
3	0.552 ± 0.252	20	0.939 ± 0.043	4
5	0.587 ± 0.234	19	0.928 ± 0.003	3

In the case of optimizing with HM, for the single realization case, ΔN increased significantly to about 0.9 (as is apparent from figure 3.10). Using multiple realizations, ΔN increased slightly relative to the result with a single HM model (to 0.93 or 0.94), while the variance decreased considerably. This reduction in variance is a more subtle but very useful effect, as it renders the results relatively insensitive to the choice of the initial geologic models.

The combined effect of multiple realizations and HM appears to generally improve ΔN and reduce its variance. As observed in the `fluvb` case, HM does not always significantly improve ΔN when a single model is used but it does reduce the variance.

3.2. Conditioned Models

We now consider conditioned models of `fluva` and `fluvb`. The models are conditioned using facies data along the trajectory of the mainbore and laterals. `LWD` provides the ability to obtain these data in real applications. Porosity and permeability fields were generated unconditionally with `SGSIM`. Well location, production and well constraints, and reservoir properties remain the same as used in previous examples. It should be noted that the actual models (`fluva` and `fluvb`) were treated as unknowns. For multiple models, Eq. 2.3 (valve optimization module) was computed with the weights, λ_j, taken as the rank transform of the inverse objective function value from the HM module.

Sensitivity runs were again made with and without HM (note that the models within these two cases were both conditioned to the facies data). As in the previous case of unconditioned models, groups (e.g., eight) of one, three and five history matched models were used and their mean ΔN and standard deviation are reported.

3.2.1. First Channel Reservoir (`fluva`)

Table 3.5 shows the result of the sensitivity runs with `fluva`. Without HM, ΔN ranged between 0.55 and 0.65 for all of the realizations. The standard deviation was fairly constant. Compared to the similar case in table 3.4 where ΔN ranged between 0.5 and 0.6, well data conditioning of the models generally provided a slight improvement in ΔN. A reduction in the standard deviation was also observed.

Table 3.5: Result of sensitivity runs on `fluva` for conditioned models

Number of realizations	without HM		with HM	
	ΔN	Number of groups	ΔN	Number of groups
1	0.581 ± 0.174	16	0.881 ± 0.062	8
3	0.548 ± 0.146	16	0.885 ± 0.091	8
5	0.643 ± 0.195	16	0.906 ± 0.093	8

With HM, ΔN from the single realization and the multiple realizations was approximately 0.9. This is consistent with the results of the unconditioned case in table 3.4.

3.2.2. Second Channel Reservoir (`fluvb`)

Table 3.6 presents the result of the sensitivity runs with `fluvb`. Without HM, ΔN was approximately 0.54 for the single realization case. For the multiple realization case, ΔN was approximately 0.6 (here we only consider three realizations). Compared to the similar case in table 3.3, where $\Delta N \sim 0.4$, conditioning the realizations to well data improved ΔN and also reduced the standard deviation.

Table 3.6: Result of sensitivity runs on `fluvb` for conditioned models

Number of realizations	without HM		with HM	
	ΔN	Number of groups	ΔN	Number of groups
1	0.543 ± 0.270	16	0.645 ± 0.173	8
3	0.604 ± 0.238	12	0.826 ± 0.096	6

With HM, ΔN from single realizations was about 0.65. Here, we observe that conditioning and history matching provides improved results with single HM realizations relative to the case with no history matching or conditioning for the `fluvb` model (ΔN = 0.645 compared to 0.393). Using three HM realizations, ΔN was approximately 0.83, which is quite similar to what was obtained using history matching without conditioning (table 3.3). It is possible that conditioning does not have more of an impact on ΔN in these cases because there is a degree of redundancy in the facies and production data (so the benefit of conditioning is not that great). This issue requires further investigation.

3.3. Computational Requirements

The valve optimization module takes about one hour to optimize a single model on PetEng-NTTS11 (Dual 2.8 GHz Intel® Xeon™ CPUs, 4.0 Gb RAM). For three and five realizations, this takes about three and five hours, respectively. The number of function evaluations is of *O*(100). The history matching module performs about 30 function evaluations (~19 evaluations in the inner loop and ~11 evaluations in the outer loop) per iteration. Each HM iteration (inner and outer loop) takes approximately 24 minutes (0.4 hour) to complete. The number of HM iterations needed to generate reasonable one, three, or five history matched models is 10, 15 or 18, respectively (though these can be reduced if the convergence criterion is relaxed). There is some linkage between the various models which acts to reduce the number of iterations per model when multiple models are used.

Let us define number of optimization and HM steps as *STEPS*, optimization time per run as *OPTtime*, HM time per iteration as *HMITERtime*, and the number of HM iterations as *numHMITER*. We can compute the maximum time required to complete the overall procedure (*VOHMtime*) with the equation:

$$VOHMtime = (STEPS \times OPTtime) + \left[(STEPS - 1) \times HMITERtime \times numHMITER\right] \quad (3.2)$$

Therefore, it will take a maximum time of about 16 hours to perform the overall VOHM procedure with a single HM model using four optimization and HM steps. Using three history matched models, this will take about 30 hours.

With multiple models, both the valve optimization and the HM modules will benefit immensely from parallelization. For the valve optimization module, this can be accomplished by running each model on different processors. With the HM module, parallelization can be accomplished by running the module on different processors to generate independent single history matched models.

Chapter 4

4. Alternative Optimization Procedures and Techniques

In previous chapters, a particular valve optimization procedure (closely related to that of Yeten, 2003) was proposed and utilized. Here, we introduce two new optimization procedures and compare them with the current approach. Furthermore, we compare the Levenberg-Marquardt (LM) optimization technique with the current conjugate gradient (CG) technique for use in the overall valve optimization and history matching (VOHM) procedure.

4.1. Alternative Optimization Procedures

The current procedure, which we will call the *global procedure*, optimizes the valve settings for each optimization step such that the objective function is maximized over the entire remaining simulation period. The disadvantage of this is that the optimization module becomes very time consuming, particularly for multiple HM models. To alleviate this problem, we propose the *local r_0 procedure* and the *r_1 procedure* (or *extended r_0 procedure*). Rather than maximize the objective function at each optimization step over the entire remaining simulation period, the local r_0 procedure maximizes the objective function at each step by considering that step only. It disregards what may happen at later times and focuses only on what occurs at the current step. The r_1 procedure is an extension of the local r_0 procedure. It maximizes the objective function by looking one step ahead. This means that it considers the current optimization step plus the next optimization step when determining the valve settings for the current step. Figure 4.1 shows the schematic representation of each of the procedures, similar to that presented in figure 2.5 above.

In order to more comprehensively compare the procedures, we introduce two new models: the layered system presented by Yeten et al. (2002) and a third channel reservoir which we refer to as `fluvc` (figure 4.2). A water cut constraint of 90% was defined for `fluvc`. The simulation was run for 1200 days with an optimization step of 240 days.

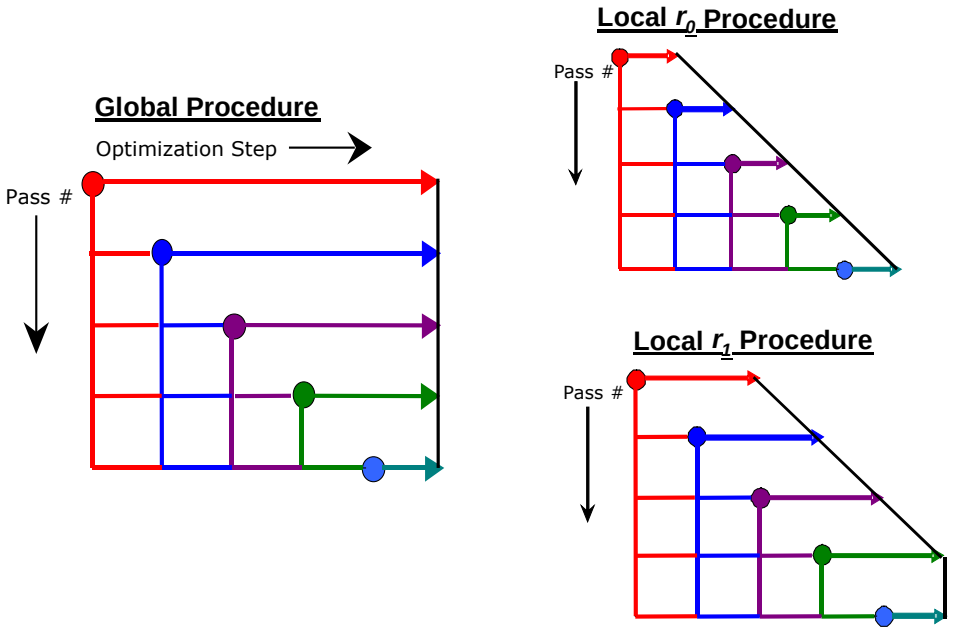

Figure 4.1: Different valve optimization procedures.

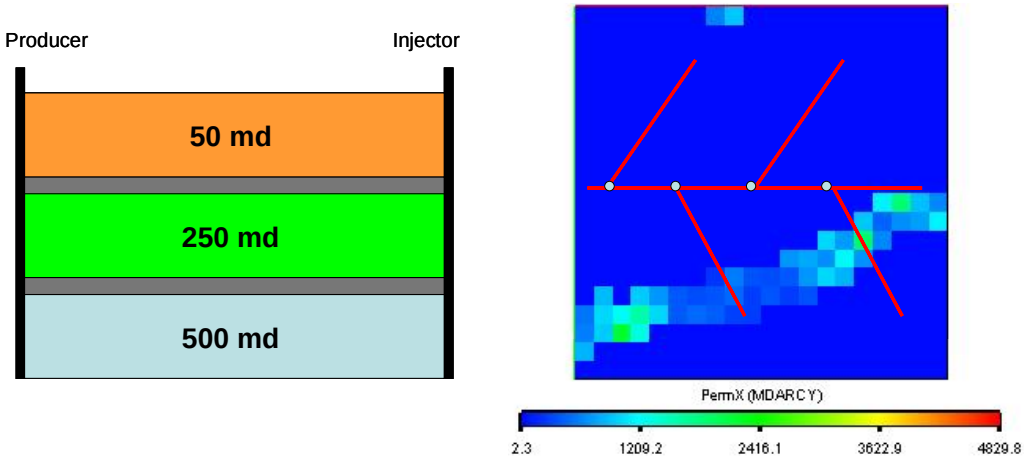

Figure 4.2: Layered model (left) and `fluvc` channel reservoir (right)

We apply these procedures to the layered model, the variogram-based model of figure 3.1, `fluva`, `fluvb` and `fluvc`. The objective function is to maximize recovery. Figures 4.3, 4.4, 4.5, 4.6 and 4.7 show the cumulative oil and water cut result for each of the models. Also included in each figure is the result for the uninstrumented case.

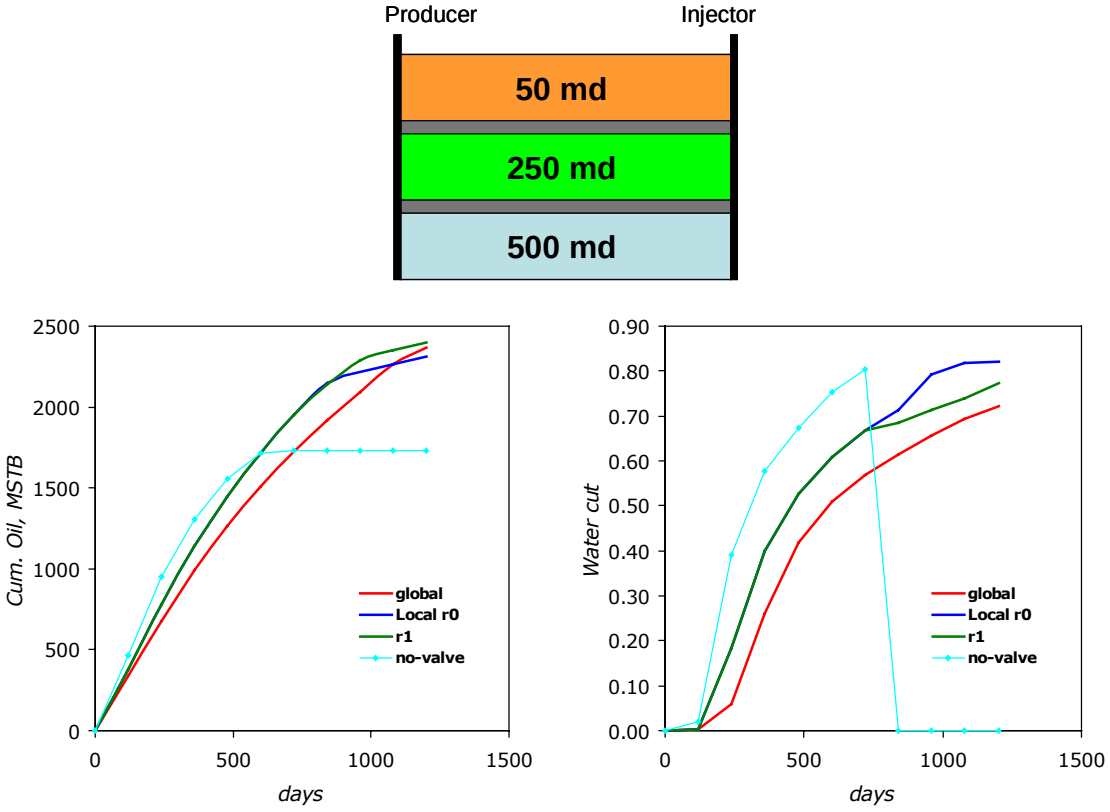

Figure 4.3: Comparison of valve optimization procedures for the layered model.

We observe in figure 4.3 that the water cut from the global procedure is lower than the two new procedures (though water cut does not appear directly in the objective function). However, the oil recovery was improved by using either the local r_0 or the r_1 procedure. In terms of the final oil recovery, the r_1 procedure performed the best of the three procedures in this case.

Figure 4.4 presents the comparison of the valve optimization procedures for the variogram based model. All of the procedures gave essentially the same results for this case.

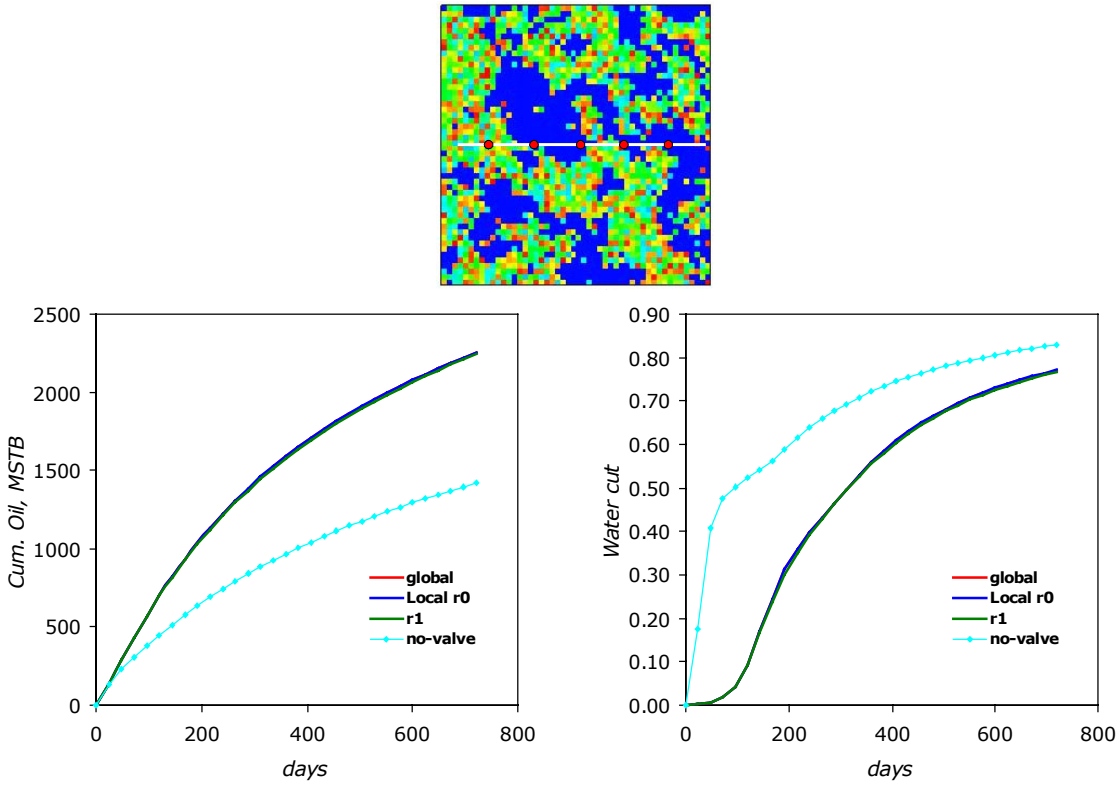

Figure 4.4: Comparison of valve optimization procedures for the variogram based model.

The comparison of the valve optimization procedures for the `fluva` channel reservoir is shown in figure 4.5. The local r_0 procedure slightly extended the life of the well in this case and provided the best result in terms of cumulative recovery. The well reached the water constraint earlier with the r_1 procedure. Note also the substantial differences in the water cut curves in this case. This indicates that the valve settings can differ significantly for the different optimization procedures.

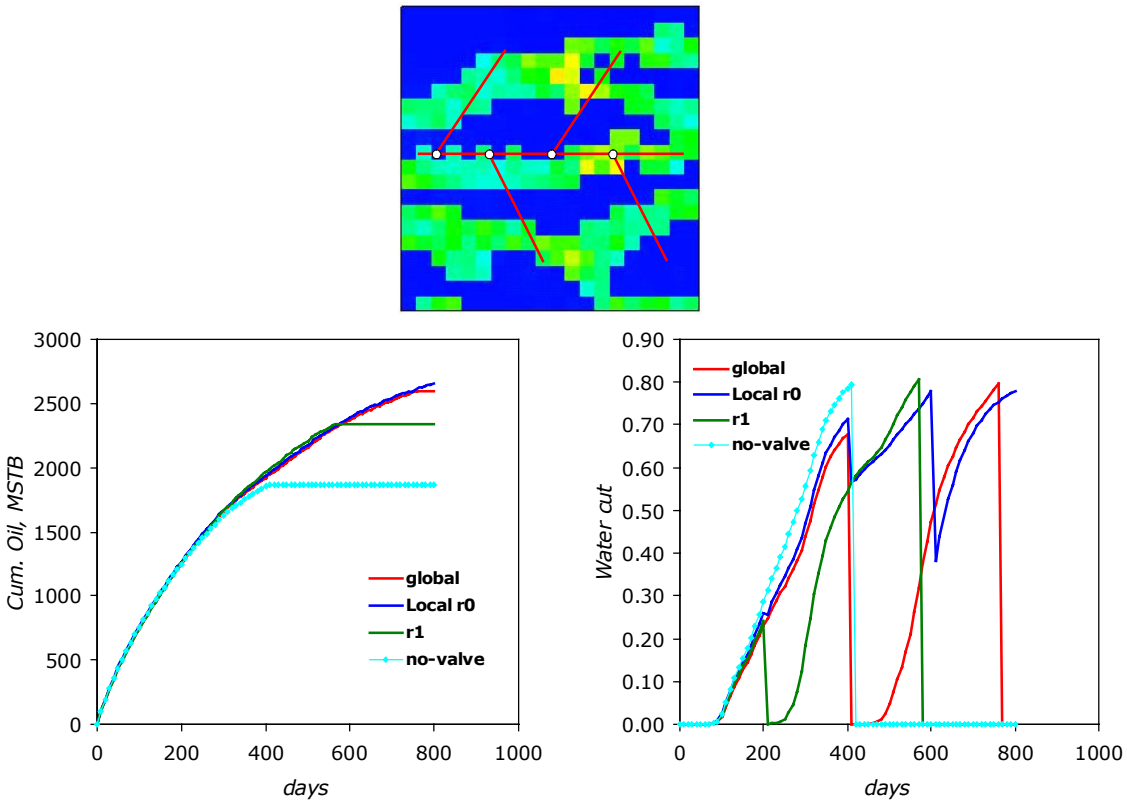

Figure 4.5: Comparison of valve optimization procedures for `fluva` channel reservoir.

Figure 4.6 shows the comparison of the procedures for the `fluvb` channel reservoir. All of the procedures gave comparable results in terms of both the water cut and cumulative oil recovery. However, the local r_0 procedure provided the best result by a slight margin.

With the `fluvc` channel reservoir, the local r_0 procedure clearly gave the best result both in terms of cumulative recovery and water cut. This is shown in figure 4.7 below. Again, the water cut varies considerably between the various procedures.

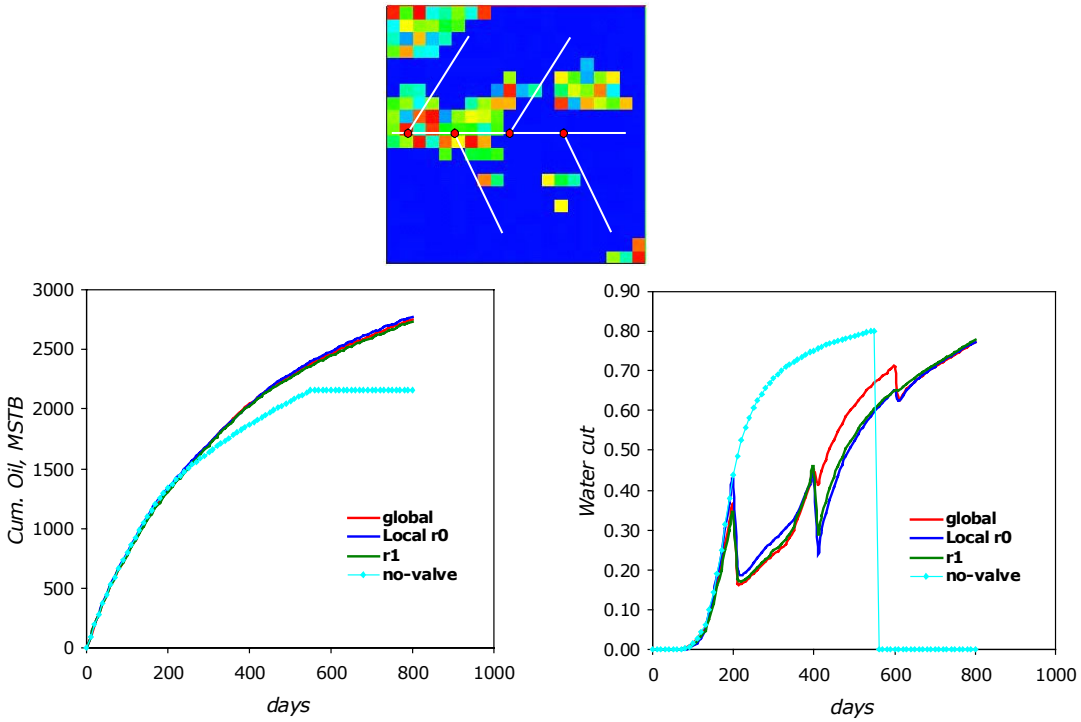

Figure 4.6: Comparison of valve optimization procedures for `fluvb` channel reservoir.

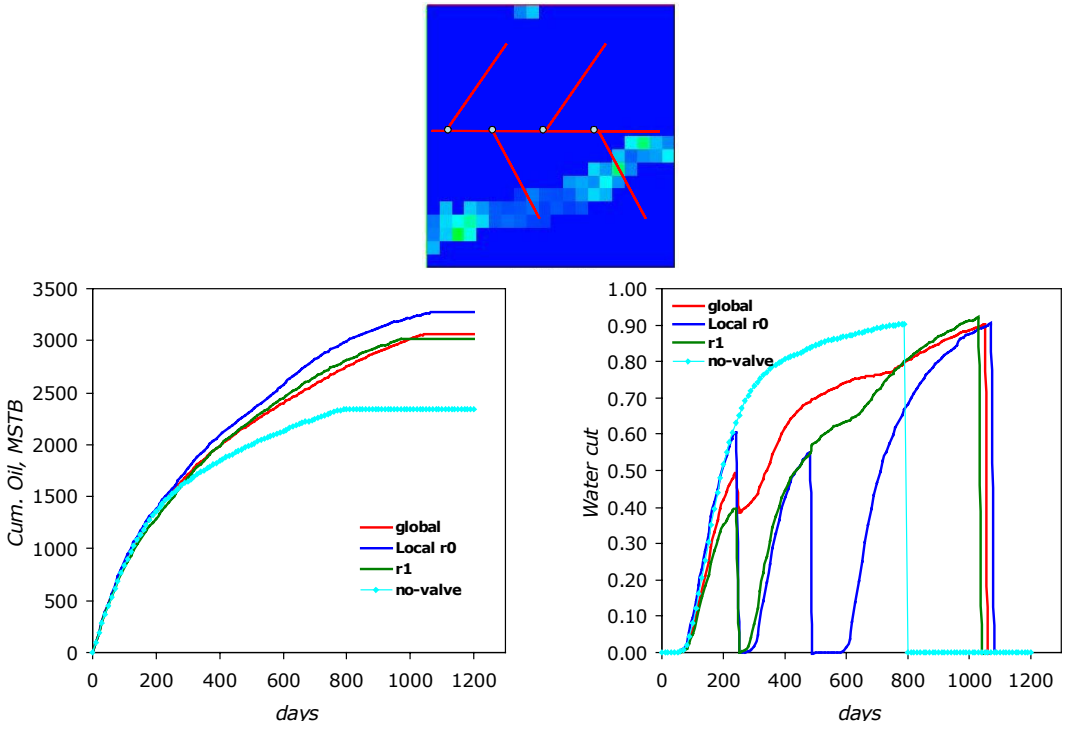

Figure 4.7: Comparison of valve optimization procedures for `fluvc` channel reservoir.

Table 4.1: Summary of the comparisons of the various procedures

Model Types	Cumulative oil, MSTB			
	Global	r_0	r_1	*No-valve*
Layered	2365	2309	2395	1733
Variogram	2252	2254	2243	1422
Fluva	2596	2659	2462	1870
Fluvb	2751	2775	2732	2159
Fluvc	3066	3273	3010	2344

The above results are summarized in table 4.1. The best result for each case is shown in red. The extended r_1 procedure appears to be the best for the layered system, while any of the procedures may be used for the variogram based model. For the channel systems, the local r_0 procedure appears to be the preferred option. Therefore, rather than optimize the valve settings over the entire simulation period at each optimization step, we can optimize only over the current optimization step, leading to computational savings.

4.2. Optimization Techniques

The optimization engine used for the VOHM results was a conjugate gradient (CG) procedure. We now consider the use of Levenberg-Marquardt (LM) instead of CG. The advantage of LM is that it is more computationally efficient than CG. Eq. 2.3 cannot be used in its current form if we apply LM. This is because LM requires second derivative information of the objective function, usually expressed as a first order quadratic term. In effect, second derivative terms are neglected based on the assumption that they are small relative to the first derivative terms (Press *et al.*, 2002). Moreover, if the second derivative terms are retained, the system will be more computationally expensive, which is undesirable.

We restate Eq. 2.3 here for convenience:

$$\underset{x_i \in [0,1]}{\text{maximize}}\, F(\mathbf{x}) \;=\; r\sigma \;+\; \sum_{j=1}^{N} \lambda_j \, f_j(\mathbf{x}) \tag{4.1}$$

Utilizing this equation as it is in an LM procedure will yield second derivative terms only and will lead to Newton's method of optimization. We therefore define the modified objective function $V(\mathbf{x})$ as the square of $F(\mathbf{x})$:

$$\underset{x_i \in [0,1]}{\text{maximize}}\, V(\mathbf{x}) \;=\; F^2(\mathbf{x}) \tag{4.2}$$

Differentiating $V(\mathbf{x})$ twice gives:

$$\begin{aligned} \mathbf{G}(\mathbf{x}) &= \frac{1}{2} V'(\mathbf{x}) = \nabla_{\mathbf{x}} V = F(\mathbf{x}) F'(\mathbf{x}) \\ \mathbf{H}(\mathbf{x}) &= \frac{1}{2} V''(\mathbf{x}) = F F'' + F' F'^{T} \end{aligned} \tag{4.3}$$

where $\mathbf{G}(\mathbf{x})$ is the gradient vector and $\mathbf{H}(\mathbf{x})$ is the Hessian matrix. $F \in [0, 1]$ is a function of the recovery factor and it multiplies the second derivative term of $\mathbf{H}(\mathbf{x})$. In the LM procedure, we neglect $F F''$ relative to $F' F'^{T}$. Eq. 4.3 is now written as:

$$\mathbf{H}(\mathbf{x}) \approx F' F'^{T} \tag{4.4}$$

The LM transformation also introduces a stronger diagonal:

$$\mathbf{H}_{\mathbf{LM}} \;=\; \mu \mathbf{I} \;+\; \mathbf{H} \tag{4.5}$$

where $\mathbf{I}$ is the identity matrix. The purpose of the μ in Eq. 4.5 is to combine both steepest descent and Gauss-Newton methods. If μ is large, the matrix becomes diagonally dominant and the LM method is similar to a steepest descent type algorithm. The LM algorithm used follows that of Press et al. (2002):

a) Set μ to 0.001 initially.

b) Compute the gradient and the Hessian matrix and solve the matrix equation $\mathbf{H}_{LM}\ \mathbf{x} = -\mathbf{G}$.

c) If $F(\mathbf{x})$ increases (maximized), reduce μ by 10; otherwise, increase it by 10 and repeat until convergence.

We demonstrate the VOHM procedure with the LM algorithm by applying it to the `fluva` channel reservoir. We use the global procedure and consider one conditioned HM model to determine the valve settings. For comparison purposes, we also show results with the CG algorithm. All well and production constraints are the same as specified in Chapter 3.

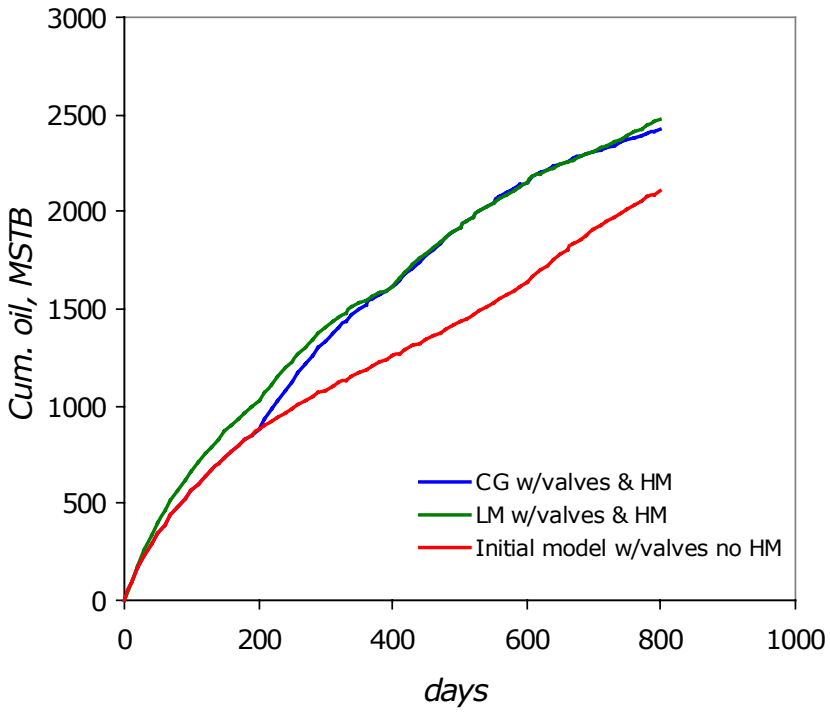

Figure 4.8: Cumulative oil comparison of LM and CG for `fluva` model.

Figure 4.8 shows the HM cumulative oil recovery with the LM and CG algorithms. Also included in the figure is the cumulative recovery from the initially specified geologic model without HM (CG is applied for the optimization of this model). Both the CG and LM curves track at later times and produced approximately the same final recovery. The water cut results of the two algorithms, shown in figure 4.9, differ considerably. This indicates that the two algorithms can produce valve settings that are significantly different.

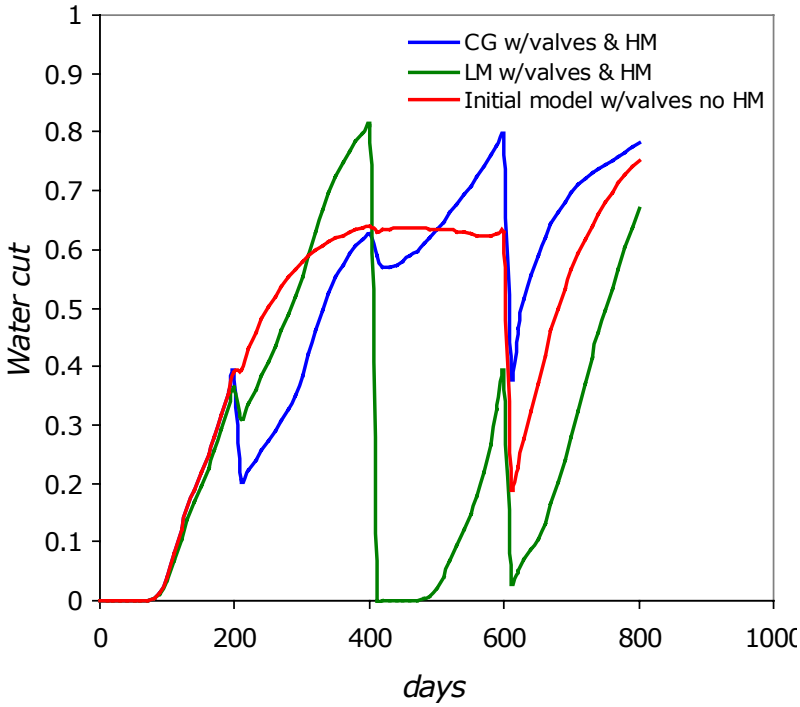

Figure 4.9: Water cut comparison of LM and CG for `fluva` model.

We next apply the VOHM procedure with the LM algorithm to the `fluvb` channel reservoir. Figure 4.10 presents the cumulative oil plot of the two algorithms and that of the initial geologic model. The cumulative oil from the CG and LM algorithms tracks reasonably well and produced similar final recoveries. The water cut responses of the

two algorithms, shown in figure 4.11, look very much alike, indicating that the valve settings from both algorithms are similar in this case.

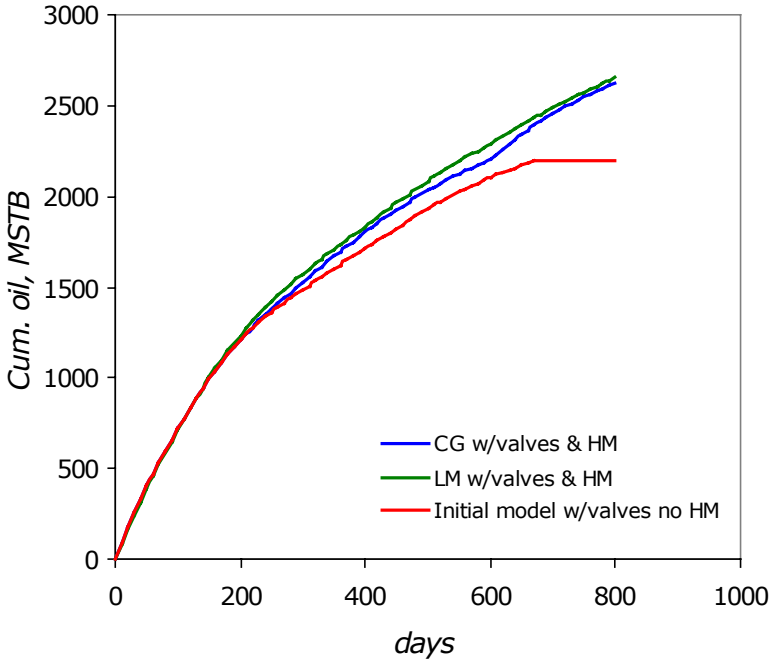

Figure 4.10: Cumulative oil comparison of LM and CG for `fluvb` model.

The LM method was three to four times faster than the CG method. This becomes important when multiple models are considered. The procedure will run even faster when the *local* r_0 procedure is used in conjunction with LM.

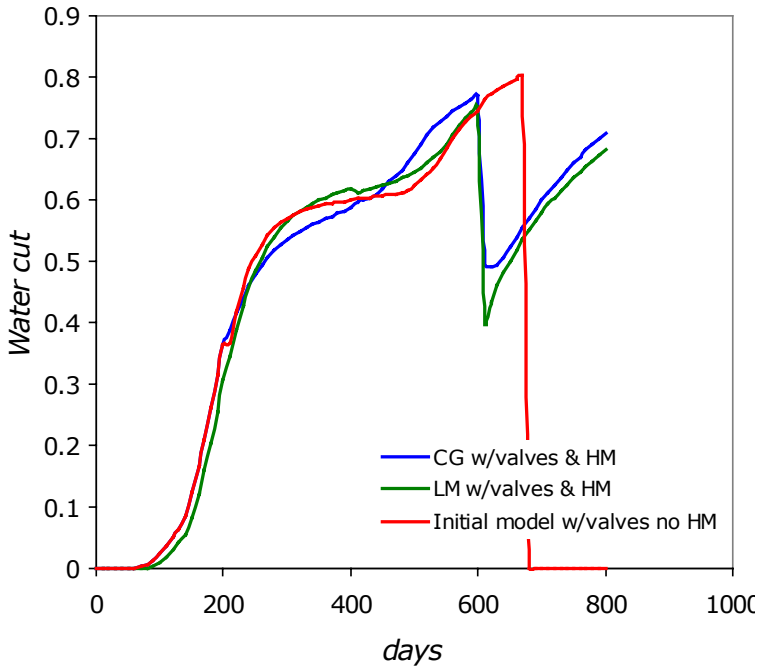

Figure 4.11: Water cut comparison of LM and CG for `fluvb` model.

Chapter 5

5. Conclusions and Future Work

5.1. Conclusions

A procedure that combines smart well valve optimization and history matching (VOHM) was developed. The procedure utilized a previous smart well valve optimization method that was based on a fixed geological model and a defensive control strategy. Valve settings were determined via optimization using a nonlinear conjugate gradient (CG) algorithm, by considering the current optimization step and the entire remaining simulation period (global procedure). In this work, geologic models were updated with an automatic history matching (HM) technique based on the probability-perturbation principle.

The benefit of the VOHM procedure was investigated. It was applied to variogram-based and channelized models using unconditioned and conditioned HM realizations. With this procedure, the improvement using HM ranged approximately between 80% and 90% of the gain attainable with known geology for both the unconditioned and conditioned cases. In some of the cases considered, the use of single HM models did not improve the result over the no-HM case. This is due to the nonuniqueness of HM solutions and was addressed with the use of multiple HM models. Using multiple HM models, the result improved significantly over the no-HM case. For one of the conditioned HM cases considered, results were similar with both the single and the multiple HM model options. In the other conditioned HM case, however, the result with single HM models was less than that with three HM models. Conditioning, did, however, act to improve the performance with single HM models for this difficult case.

Although the overall procedure was effective, it was very time consuming. Efficient alternative optimization procedures - *local* r_0 *and* r_1 - were therefore developed and

investigated. The *local* r_0 procedure was shown to provide the best overall results and was the most efficient. We also investigated the effectiveness of the Levenberg-Marquardt (LM) optimization algorithm by comparing its results with those of the CG algorithm. Both optimization algorithms produced similar cumulative recovery results. LM was found to be three to four times faster than CG, leading to significant computational savings.

5.2. Future Work

We suggest the following areas for future investigations:

- Introduce noise effects in the sensor data.
- Investigate the effectiveness of the LM method in the facies optimization part of the HM procedure, instead of the Brent method currently used. This may improve the speed of this module through a reduction in function evaluations.
- Develop a parallel implementation of the overall procedure.
- Consider the use of other data (e.g., 4D seismic) in the history matching procedure.
- Apply the overall procedure to suitable real field applications.

Nomenclature

A	binary random function event
B	training image event or formation volume factor (rb/stb or rb/mcf)
$\mathbf{b}$	arbitrary vector
c	compressibility, psi^{-1}
D	production data event
$\mathbf{d}$	optimization search direction
$\mathbf{e}$	set of unit vectors (identity matrix)
F	objective function
f	recovery factor
$\mathbf{g}, \mathbf{G}$	gradient vectors
$\mathbf{H}$	Hessian matrix
h	objective function derivative step size
I	binary random function
i	binary data / outcome
K	maximum number of iterations
k	permeability, md
k_r	relative permeability
N	total number of realizations (or models)
N_p	cumulative oil production, MSTB
ΔN	dimensionless cumulative oil difference
P	probability
p	pressure, psi
R	gas-oil ratio, scf/STB
r	risk factor
$\mathbf{r}$	residual vector
r_D	free parameter used in history matching
s	facies
$\mathbf{u}$	grid location
V	modified objective function
WOR	water-oil ratio
$\mathbf{x}$	scaled vector of valve settings

Supercripts

k	iteration level
m	segment number
T	transpose

Subscripts

ave	average
b	bubble point
f	friction
g	gas
h	horizontal
i	intial
l	facies number
min	minimum
o	oil
obs	observed
s	solution
v	vertical
w	water

Greek Symbols

α	step size in search direction
$\boldsymbol{\alpha}$	vector of scaled multivariate parameters
β	Gram-Schmidt constant
ε	convergence factor
ϕ	porosity
γ	specific gravity, dimensionless
λ	weights
μ	viscosity, cP or Levenberg-Marquardt transformation factor
σ	standard deviation, data units

Abbreviations

CG	conjugate gradient
HM	history matching
LM	Levenberg-Marquardt
NCW	nonconventional well
`SGSIM`	sequential Gaussian simulation
`SNESIM`	single normal equation simulation
VOHM	valve optimization and history matching

References

Arenas, E. and Dolle, N., 2003, "Smart Waterflooding Tight Fractured Reservoirs Using Inflow Control Valves", paper SPE 84193 presented at the SPE Annual Technical Conference and Exhibition, Denver, Colorado, 5-8 October.

Al-Khodhori, S. M., 2003, "Smart Well Technologies Implementation in PDO for Production & Reservoir Management & Control", paper SPE 81486 presented at the SPE Middle East Oil Show and Conference, Bahrain, 5-8 April.

Bittencourt, A. C. and Horne, R. N., 1997, "Reservoir Development and Design Optimization", paper SPE 38895 presented at the SPE Annual Technical Conference and Exhibition, San Antonio, Texas, 5-8 October.

Brouwer, D. R., Jansen, J. D., Van Der Starre, S., Van Kruijsdijk, C. P. J. W., and Berentsen, C. W. J., 2001, "Recovery Increase through Water Flooding with Smart Well Technology", paper SPE 68979 presented at the SPE European Formation Damage Conference, The Hague, Netherlands, 21-22 May.

Brouwer, D. R. and Jansen, J. D., 2002, "Dynamic Optimization of Water Flooding with Smart Wells Using Optimal Control Theory", paper SPE 78278 presented at the SPE European Petroleum Conference, Aberdeen, UK, 29-31 October.

Badru, O. and Kabir, C. S., 2003, "Well Placement Optimization in Field Development", paper SPE 84191 presented at the SPE Annual Technical Conference and Exhibition, Denver, Colorado, 5-8 October.

Caers, J., 2003, "History Matching Under Training-Image-Based Geological Model Constraints", *SPE Journal*, September, pp 218-226.

Deutsch, C. V. and Journel, A. G., 1998, *GSLIB User's Manual*. Applied Geostatistics Series, Oxford University Press, New York, second edition.

Dolle, N., Brouwer, D. R., and Jansen, J. D., 2002, "Dynamic Optimization of Water Flooding with Multiple Injectors and Producers Using Optimal Control Theory", paper presented at the XV International Conference on Computational Methods in Water Resources, Delft, Netherlands, 23-28 June.

Gai, H., 2001, "Downhole Flow Control Optimization in the World's 1st Extended Reach Multilateral Well at Wytch Farm", paper SPE 67728 presented at the SPE/IADC Drilling Conference, Amsterdam, Netherlands, 27 February-1 March.

GeoQuest Schlumberger, 2001a, *ECLIPSE Reference Manual 2001A*.

GeoQuest Schlumberger, 2001b, *ECLIPSE Technical Description 2001A*.

Guyaguler, B., Horne, R. N., Rogers, L., and Rosenzweig, J. J., 2002, "Optimization of Well Placement in a Gulf of Mexico Waterflooding Project", *SPE Reservoir Evaluation & Engineering*, June, pp 229–236.

Hoffman, T. B. and Caers, J., 2003, "History Matching Using the Regional Probability Perturbation Method", Stanford SCRF Group Report 16.

Holmes, J. A., Barkve, T., and Lund, Ø., 1998, "Application of a Multi-segment Well Model to Simulate Flow in Advanced Wells", paper SPE 50646 presented at the SPE European Petroleum Conference, The Hague, Netherlands, 20-22 October.

Kabir, C. S., Chien, M. C. H., and Landa, J. L., 2003, "Experiences With Automated History Matching", paper SPE 79670 presented at the SPE Reservoir Simulation Symposium, Houston, Texas, 3-5 February.

Khargoria, A., Zhang, F., Li, R., and Jalali, Y., 2002, "Application of Distributed Electrical Measurements and Inflow Control in Horizontal Wells under Bottom-Water Drive", paper SPE 78275 presented at the SPE European Petroleum Conference, Aberdeen, UK, 29-31 October.

Press, W. H., Teukolsky, S. A., Vetterling, W. T., and Flannery, B. P., 2002, *Numerical Recipes in C++: The Art of Scientific Computing*. Cambridge University Press, Cambridge, UK, second edition.

Snaith, N., Chia, R., Narayasamy, D., and Schrader, K., 2003, "Experience with Operation of Smart Wells to Maximize Oil Recovery from Complex Reservoirs", paper SPE 84855 presented at the SPE Asia Pacific International Improved Oil Recovery Conference, Kuala Lumpur, Malaysia, 20-21 October.

Strebelle, S., 2000, *Sequential Simulation: Drawing Structures from Training Images*, Ph.D dissertation, Department of Petroleum Engineering, Stanford University, California.

Suzuki, S., 2003, *Determining Petrophysical Properties of Facies Using a Hierachical History Matching Method*, Master's Report, Department of Petroleum Engineering, Stanford University, California.

Yeten, B., Durlofsky, L. J., and Aziz, K., 2002, "Optimization of Smart Well Control", paper SPE 79031 presented at the SPE International Thermal Operations and Heavy Oil Symposium and International Horizontal Well Technology Conference, Calgary, Canada, 4-7 November.

Yeten, B., Durlofsky, L. J., and Aziz, K., 2003, "Optimization of Nonconventional Well Type, Location and Trajectory", *SPE Journal*, September, pp 200-210.

Yeten, B., 2003, *Optimum Deployment of Nonconventional Wells*, Ph.D dissertation, Department of Petroleum Engineering, Stanford University, California.

Yeten, B., Brouwer, D. R., Durlofsky, L. J., and Aziz, K., 2004, "Decision Analysis Under Uncertainty for Smart Well Deployment," *Journal of Petroleum Science and Engineering* (to appear).

Appendix

Flow chart of overall valve optimization and history matching procedure

The programme is written in both C++ and FORTRAN. The valve optimization module is in C++ while the HM module is written in FORTRAN. Both are linked through the *SmartOptHm* C++ code. The overall structure of *SmartOptHm* is shown in the flowchart below.

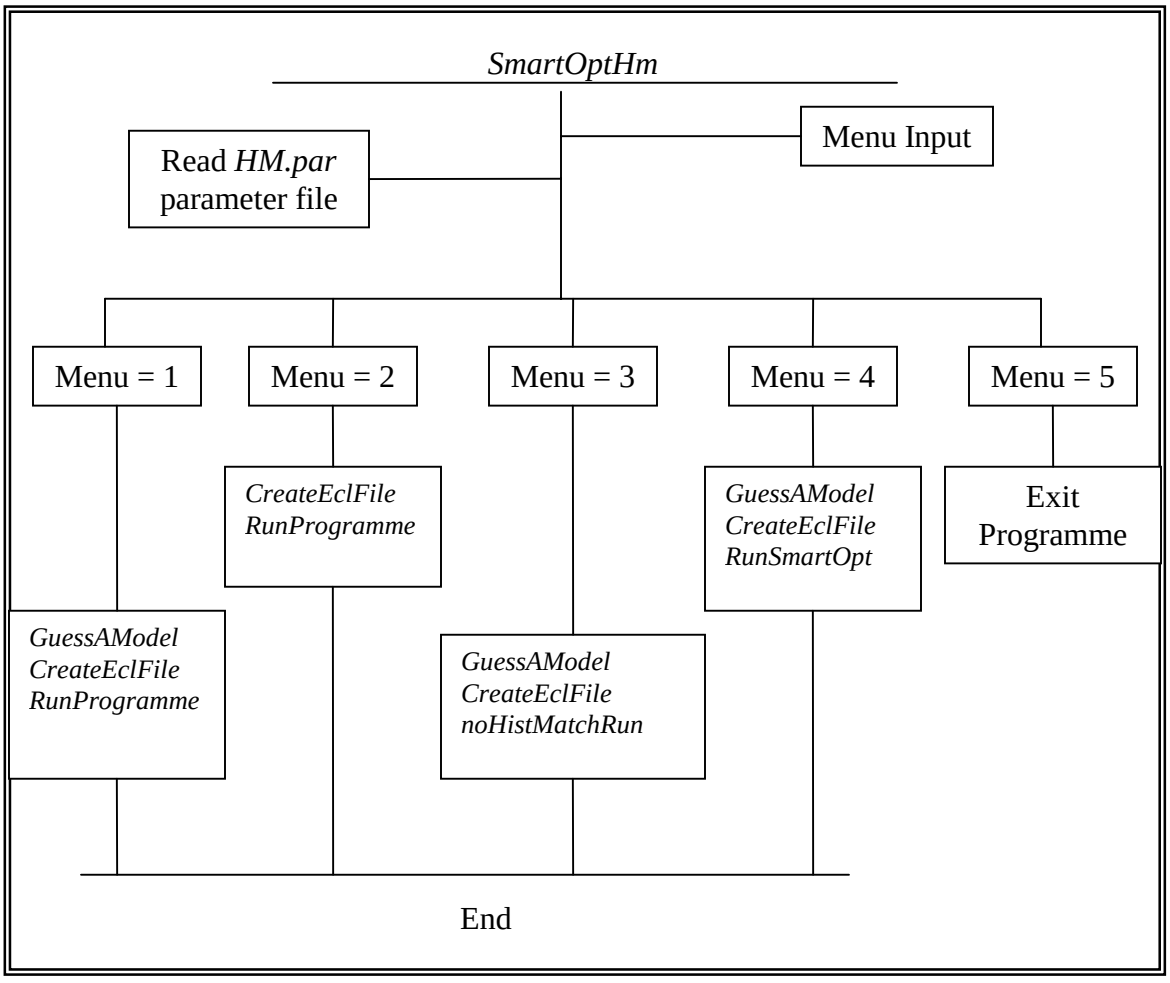

Figure A.1: Overall structure of *SmartOptHm* code.

The structure of each of the subroutines in *SmartOptHm* code is shown in figure A.2. The helper functions are called in the order shown in the figure. Also shown in figure A.2 are the history matching and valve optimization procedures.

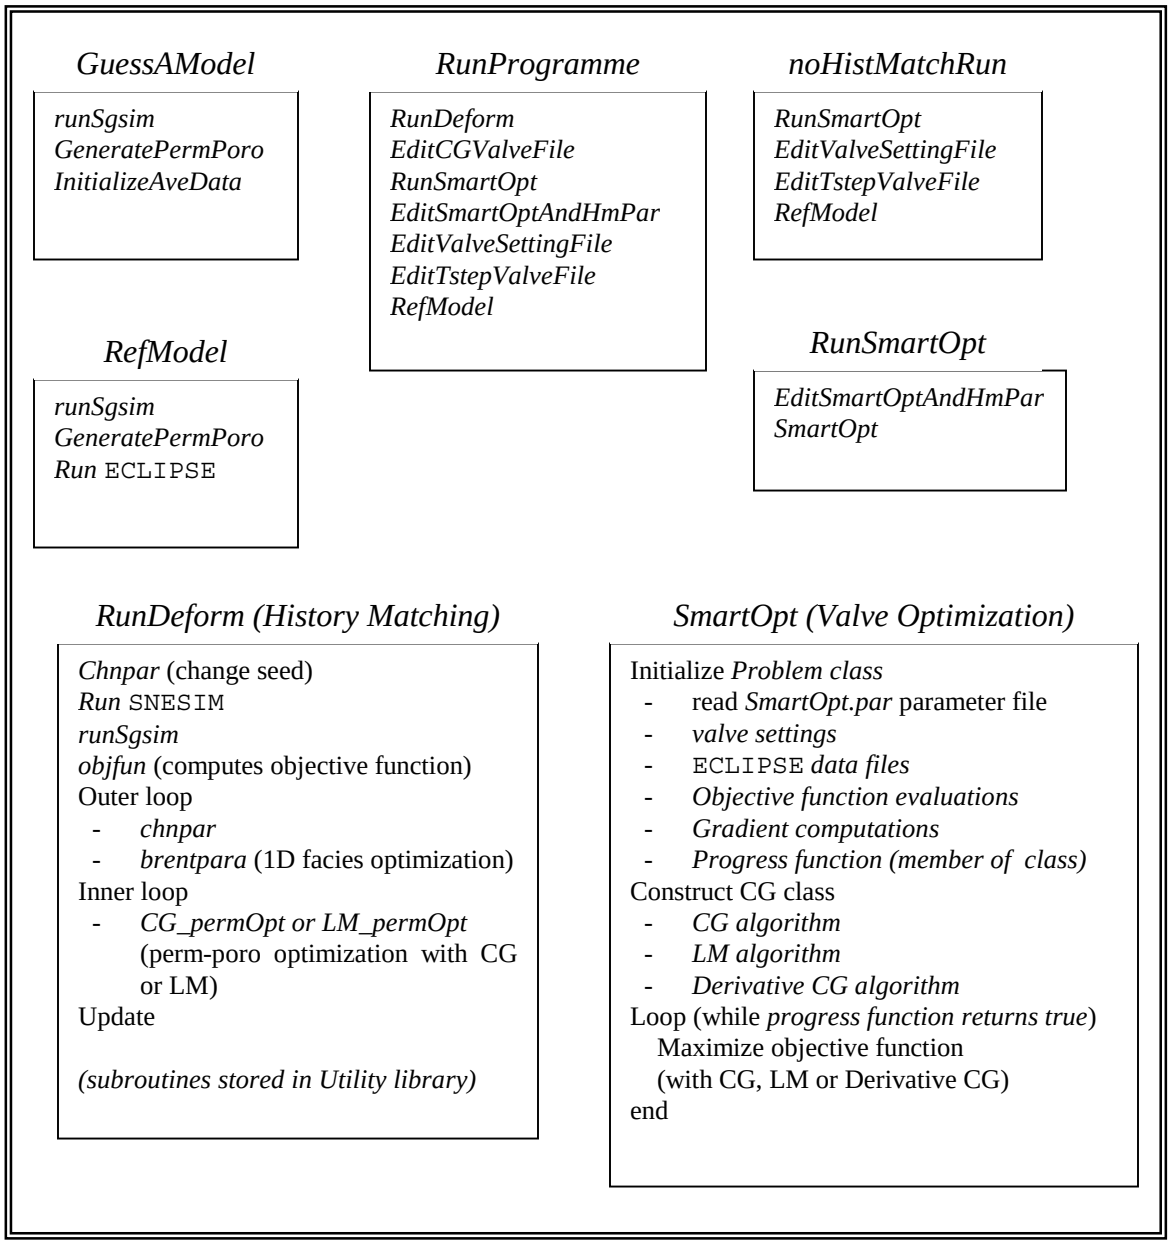

Figure A.2: Structure of subroutines, HM and valve optimization procedures in *SmartOptHm* code