GEOLOGICALLY CONSISTENT HISTORY MATCHING WHILE PERTURBING FACIES

A DISSERTATION

SUBMITTED TO THE DEPARTMENT OF PETROLEUM ENGINEERING

AND THE COMMITTEE ON GRADUATE STUDIES

OF STANFORD UNIVERSITY

IN PARTIAL FULFILLMENT OF THE REQUIREMENTS FOR THE DEGREE OF

DOCTOR OF PHILOSOPHY

B. Todd Hoffman

May 2005

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

_____

Jef Caers, Principal Advisor

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

_____

Andre Journel

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

_____

Anthony Kovscek

Approved for the University Committee on Graduate Studies

Abstract

Reservoir modeling is used to better understand the complex flow behavior in the reservoir. History matching is an integral part of any reservoir modeling study where the reservoir has produced for some time. History matching is the process where properties in the reservoir model are changed until the production data from the model matches the field data. Traditionally, this has been a heuristic trial-and-error process, and often other reservoir data, such as geologic continuity, is sacrificed to achieve a match.

This dissertation presents an innovative methodology to match production data and honor other reservoir data including prior geologic information. For most reservoir modeling, geostatistical methods are used to generate properties for the models (e.g. porosity and permeability). By directly perturbing these properties, there is a risk that the intended geological continuity or static data conditioning is destroyed. In this dissertation, a method termed probability perturbation is proposed that relies on perturbing the underlying probabilities used to generate the properties and not the properties directly. It is shown that the geology and all other data in the models can be maintained with this perturbation method. By including all reservoir data, predictions from reservoir models are more accurate than models that match production data but do not honor other reservoir data.

One of the most novel aspects of this history matching method is that large scale parameters, such as facies geometry, can be perturbed, contrary to traditional methods that focus mostly on gridblock permeability. Large scale parameters have a substantial impact on fluid flow in the reservoir, but perturbing them in a history matching application has been notoriously difficult due to their discrete nature. The main focus of this dissertation is to vary facies distributions and proportions. A multi-parameter method is developed that allows different parts of the reservoir model to be perturbed by

different amounts, dramatically increasing the efficiency of the method. Numerous synthetic examples are used to study the effectiveness and robustness of the method.

While the method is shown to be theoretically sound, its greatest strength may be its practicality. A major focus of the work is to create a methodology that can be applied to actual reservoirs; the method is successfully demonstrated on two large real field cases. In both cases, the proposed method is able to significantly improve the match and honor the static information from the reservoir.

Acknowledgments

Rather than thanking my family at the end of the section, as most do, I want to start by thanking my lovely wife, Holly. She has endured the strain of a Ph. D. candidate, and the many late nights without so much as a complaint. She is always smiling, encouraging me, and helping me appreciate events in life besides those related to petroleum engineering and for this I am truly grateful. To my daughters, Dove and Reagan, thanks for your patience with your dad when he had too little sleep and too much stress. Every day you bring me back to earth and make me realize what is truly important. I would like to express my gratitude to all my family for their love and confidence in me, in particular, my father, Roy, for his unconditional support, my mother, Sharon, for her untiring care and watchfulness, and my brothers, R.J. and Kevin, for their affection.

This work would not have been possible without my advisor, Jef Caers. His thoughtful guidance and advice added to the quality of this project as well as my development as a person, and his enthusiasm is something to aspire to. I thank him for this and the time he spent preparing me for an academic position.

I thank Andre Journel for his wisdom and for imparting some of his knowledge to me through our many meetings, discussions and classes. I also appreciate his thorough reading of this work and his even more thorough set of constructive comments clarifying many issues.

I want to thank both Jef and Andre for the hard-working but fun research group they run. I have enjoyed the numerous SCRF outings we have partaken and the exceptional conferences we have attended.

I want to thank Tony Kovscek for being on my reading committee but more importantly for all the guidance he provided my first two years at Stanford as my Master's advisor. I

also appreciate Khalid Aziz for serving on my committee and Peter Kitanidis for chairing my committee. All faculty at Stanford, both in and out of the Petroleum Department, have had a tremendous impact on my life and to all of them, thank you.

To all my friends I have made here at Stanford, I will treasure all the memories we have created. Thank you for making my years at Stanford truly enjoyable.

I appreciate Statoil and ChevornTexaco for providing me with two exception case studies to test my initial algorithm. The experience I gained from working on real fields contributed greatly to the quality of this thesis. In particular I want to thank Martin Springer from Statoil for answering my many questions on the North Sea example and our many discussions on my trip to Norway. I also want to thank Xian-Huan Wen and Sebastien Strebelle for their assistance in my summer job with ChevronTexaco that lead to the West Coast African example. The financial support of all SCRF affiliates is also gratefully acknowledged.

Table of Contents

List of Tables

Chapter 1: Introduction

Chapter 3: Multi-Parameter PPM

Chapter 4: Channel Reservoirs

Chapter 5: Perturbing Proportions

List of Figures

Chapter 1: Introduction

Chapter 2: Probability Perturbation Method

Chapter 3: Multi-Parameter PPM

Chapter 4: Channel Reservoirs

Chapter 5: Perturbing Proportions

"If we knew what it was we were doing, it would not be called research, would it?"

- Albert Einstein

Chapter 1

Introduction

Reservoir simulation is a powerful tool available to petroleum engineers. It is used in many aspects of reservoir engineering to help understand the complex nature and behavior of petroleum reservoirs during production. Reservoir simulation is vital for developing new well and facility plans, for determining the optimal strategy to implement secondary or tertiary recovery, and for estimating production forecasts for economic models. In most reservoir simulation studies, the primary goal is to have an accurate prediction of future reservoir behavior. To obtain more reliable models, this goal should be kept in mind throughout all aspects of building the reservoir model.

A wide variety of data are used to build the reservoir models. Well logs and some core data provide high resolution information about properties near the wells. Seismic data provides information over the entire field, but the resolution is much poorer than well data. A further source of information comes from the geological setting of the reservoir. While this may be difficult to quantify, it is critical for building a reliable reservoir model. Collectively, all of these data are labeled "static" data. Geostatistics is used to

combine these varying sources of information into a consistent reservoir model that honors all of the data. Other engineering data such as relative permeability, capillary pressure and productivity indices are also treated as static data.

An important source of information comes from producing the reservoir, typically in the form of fluid rate and pressure data. This data, often termed "dynamic" data, should also be included in the reservoir model. The process of including historical production data into the reservoir models is called history matching, and it is often the most time consuming and difficult part of building a reservoir model (Mattox & Dalton, 1990). It is challenging for a number of reasons.

(1) Integrating production data into the model requires an iterative process of changing the model and observing how those changes affect the flow of fluids. Thus, multiple flow simulations are required to history match a reservoir, and each flow simulation may take hours to complete. Therefore, a history matching project may take weeks or even months to complete. Limiting the number of flow simulations is vital to developing an efficient and practical history matching methodology.

(2) The relationship between the properties of the reservoir model and the simulated production is multivariate and non-linear (Tan & Kalogerakis, 1991). Therefore, it is difficult to predict or model a priori the changes in the production data due to the changes in the reservoir model. For this reason, production data requires an entirely different treatment than 3D seismic does.

(3) It has long been recognized that history matching is not only a difficult problem but also a non-unique problem (Veatch & Thomas, 1971; Gavalas *et al.*, 1976; Yamada, 2000). Non-uniqueness means that multiple different combinations of model parameters can lead to acceptable matches to the production data. Although the various models produce similar results to the historical reservoir performance, they may produce a wide range of results

> when it comes to the prediction of future performance (Landa & Guyaguler, 2003). Most of the "matched" models will not honor the reservoirs static data; however, a subset of them will match both the static and dynamic data. When history matching, we should be searching for this subset.

Due to the difficulties of history matching, reservoir models traditionally are constructed in two decoupled steps. First the static data is incorporated into a high resolution geocellular model by geoscientists. Then the model is passed to an engineer who adjusts the model to match the dynamic data observed in the field. However, most current techniques for model adjustment are disconnected from the methods to generate the high resolution geologic models. Consequently, the completed reservoir model may honor the dynamic data but may not continue to honor the static data, and in many cases, the model is unacceptable from a geological point of view.

For example, Figure 1.1A displays a channel reservoir model that has been built by geoscientists using available well, seismic, and geologic data. Flow simulation results for two wells from the model are displayed in Figure 1.1A along with the "true" production data. The production data does not match, and the model is passed to the engineer who "history matches" it. The history matched model agrees with the production data quite nicely (Figure 1.1B); however, any concept of a channel reservoir has been destroyed. Throughout the dissertation this type of history matching will be referred to as traditional history matching.

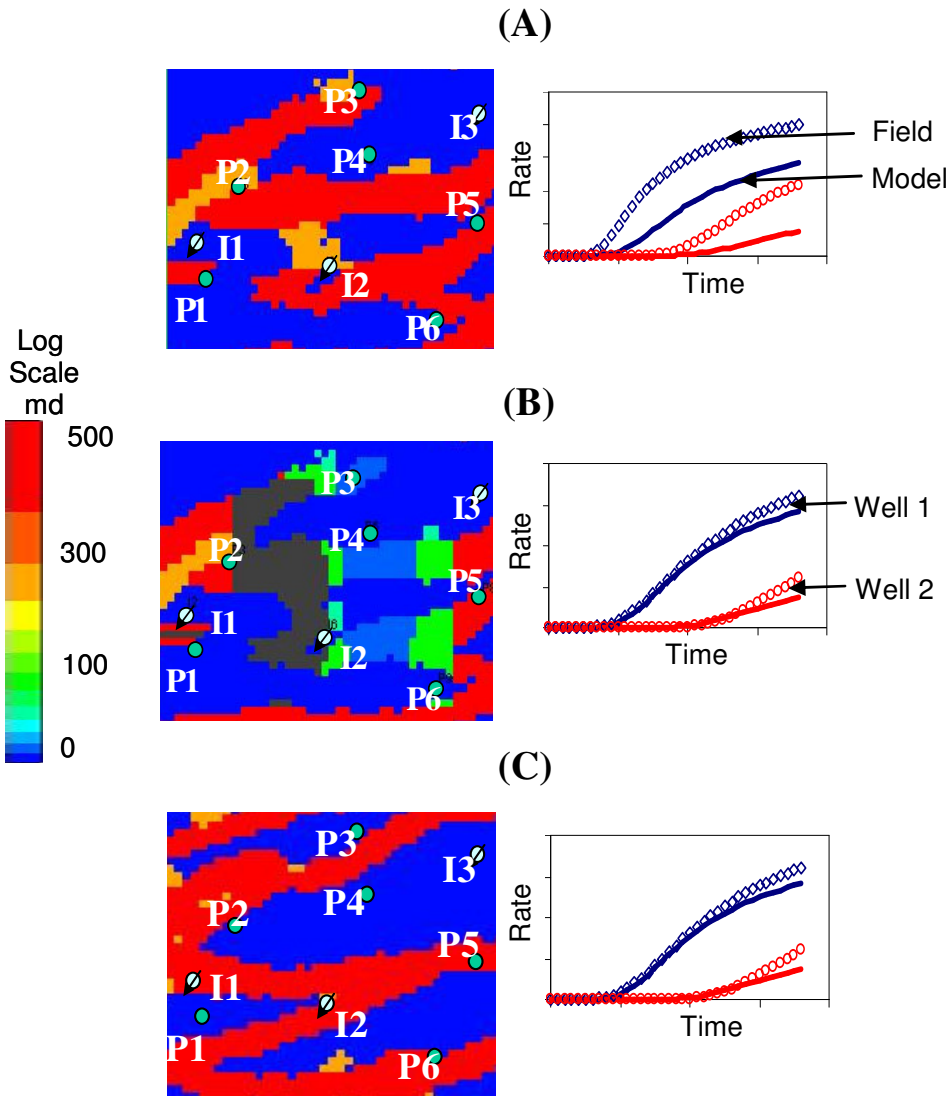

Figure 1.1: History matching example: (A) Initial model that does not match production data, (B) traditional history matching method that matches production data but does not match geology, and (C) method where geology and production data are matched.

The current thesis attempts to better connect building the reservoir model by the geoscientists and matching the production data by the engineer. The premise of the work is to match the production data and honor all other static data including critical geological interpretation. Figure 1.1C displays an example where this has been successful.

1.1 Reservoir Modeling

Sophisticated computer models are built for most petroleum reservoirs. To construct the models, a large amount of information about the reservoir is required. Properties about the rocks and fluids are necessary, as well as the "geology" of the reservoir. The geology of a reservoir encapsulates a wide variety of concepts and model attributes. It usually describes the structure and layering of the reservoir, but it can also include within layer architecture and the petrophysical properties like porosity and permeability. It is sometimes used to describe details at an even smaller scale (e.g. the pore-level scale), but in most cases it refers to the larger scale parameters (layering, facies), and the current work will be restricted to that meaning.

The reservoir geology is measured at a number of different scales. At the finest scale, there are core measurements of porosity and permeability. This is taken on a scale of a few inches and may be the only direct measurements of permeability for the reservoir. The well log data is also at the fine scale. From this data we acquire fluid saturations, facies or sand/shale identification and indirect measurements of porosity. The scale for well log data is normally around one foot.

On the medium scale, we classify properties such as facies distributions and sand/shale ratios. This data is commonly on the order of tens to hundreds of feet. The model properties depend largely on prior geologic information or geological interpretations because few direct measurements exist at this scale. Well measures are much too fine, while seismic is too coarse to resolve variability at this scale. Well test data provides only average properties at this scale. Consequently, these properties typically have a considerable amount of uncertainty; however, they are important to model, for example, inter-well connectivity.

At the largest scale, one considers the major reservoir boundaries, the fluid contacts, and fault blocks. These properties are often well known because of the numerous sources of

data informing these properties. Well and seismic data can provide a wealth of information; and coupling these data with the geologic knowledge of the reservoir usually provides a fairly clear description of the reservoir structure; however, some amount of uncertainty will always exist for any sub-surface reservoir property. The large scale properties range from hundreds to thousands of feet in dimension.

Table 1.1: Scale of simulation properties used to build a reservoir model.

Field (large) Scale
- Horizons
- Faults
- Fluid Contacts
Regional (medium) Scale
- Net to Gross Ratio
- Facies Locations
- Fluid Properties (Bo, Rs, etc.)
- Flow Properties (kr, Pc)
Gridblock (small) Scale
- Porosity
- Permeability
- Initial Saturations

When building a reservoir model, the need for starting with the larger scale properties and subsequently working on smaller scales should be obvious. The locations of the different facies cannot be determined until the overall structural framework of the reservoir model is completed, and the porosity cannot be assigned to a gridblock until the facies that the gridblock belongs to has been determined. While the stages of building a reservoir model are not independent (for example petrophysical property measurements often inform the location of facies), the need for building the model with the appropriate parameter hierarchy is important.

The actual framework for building a reservoir model consists of the following steps:

(1) The boundaries of the reservoir must be defined first. This includes defining the separate horizons, incorporating the water-oil and gas-oil contacts, if appropriate, and determining the location of the major faults. The culmination of this step will be a 3D stratigraphic grid. Seismic data provides long range structure; whereas, well data constrains surfaces locally.

(2) Once the overall structural framework is complete, the separate facies or rock types are added to the model, and the sand/shale proportions are assigned to different regions. This data is usually inferred from well measurements and is extrapolated throughout the model using the geologic understanding of the reservoir. Other parts of the reservoir model (such as fluid properties [e.g. formation volume factors] and flow parameters [e.g. relative permeability]) are typically defined at regional scale, as well.

(3) Finally petrophysical properties such as porosity and permeability are allocated to each gridblock. Porosity is usually directly measured at the well locations and populated throughout the reservoir with a geostatistical technique. The relationship between porosity and permeability is determined from core measurements, and the relationship is used to determine gridblock permeability. Once this step is finished, the high-resolution 3D geocellular model is completed, and it can be used for flow simulation.

1.2 Uncertainty

At the largest scale, reservoir simulation properties have uncertainty. Seismic data usually is available for the entire reservoir but can not be resolved to the scale of the reservoir simulation blocks nor does it predict whether the fault is conducting or sealing. It is also by nature inexact; for example, horizons may be shown in the incorrect locations. Well log information is on a fine scale but does not inform properties between wells. Similarly, uncertainty exists at the medium scale where facies locations and sand proportions are modeled. To determine the spatial properties between wells, there is a large dependence on geologic interpretation of the reservoir. The type of geologic model and especially the exact locations and proportion of the various facies are inherently uncertain. Uncertainty in the parameters at the smallest scale also exists. Localized measurements on a very fine scale (well-logs and cores) inform non-measured locations only qualitatively (e.g. they are used to help generate a geologic map, a variogram or a training image), but porosity and permeability between wells are never known precisely.

These uncertainties are due to limited access and therefore limited measurements of the reservoir. For different reservoirs the uncertainty may be greater for one scale than another, but for all reservoir models there will be some uncertainty at each scale. From a geostatistical point of view, the uncertainty tends to be grouped into two general categories: geologic scenarios and reservoir properties. A further source of uncertainty, while fundamentally different than the other two, is important in reservoir modeling, and it termed random seed uncertainty.

The geologic scenario accounts for the type of reservoir (e.g. channel versus beach sand), the major structure (e.g. the existence of an aquifer) and the type of geostatistical method used. The geologic scenario has a significant impact on flow in a reservoir; however, since the set of scenarios is discrete, it cannot be perturbed in a typical history matching workflow. For this type of uncertainty, we are strongly dependent on the geologist's and geophysicist's ability to describe the nature of the reservoir, the

uncertainty associated with a given geologic scenario, and the possible alternative geologic scenarios.

The next source of uncertainty is associated with the reservoir properties in the model. Large and middle scale parameters, for instance, fault transmissibility, fluid contacts, relative permeability and facies proportions, as well as small scale parameters such as permeability and porosity histograms and variograms are subject to uncertainty. The uncertainty at the property level is usually very large and also important to flow simulation. Unlike geologic scenarios, reservoir properties are often easily parameterized in history matching, so they can be perturbed within their range of uncertainty.

The other uncertainty is associated with what is known in geostatistics as "multiple realizations." Once various input parameters such as facies proportions, histograms, and variograms are defined, geostatistics is used to generate a "realization" of a particular property be it facies or petrophysical. Actually, multiple equal-probable realizations can be generated by changing the random seed in the geostatistical algorithm. All realizations honor the same input parameters; therefore, this type of uncertainty is only due to the random seed. This uncertainty is fundamentally different than the first two because it is due to the algorithm while the other two come from physical reality. Depending on the type of geostatistics and the particular case, this uncertainty can be substantial, but it is less that the previous two types of uncertainty (Caumon *et al.*, 2004).

Uncertainty provides both the need and the means for history matching. Because the reservoir properties are uncertain, the static reservoir model (built without incorporating production data) will output production that does not match the field-measured production data. Since uncertainty is known to exist in all reservoir model properties, these properties can be changed, within their range of uncertainty, until the simulated production is equal to the observed production. However, determining which properties to change, how to change them and by how much (a vital part of history matching) is not a simple task.

1.3 History Matching

In most reservoir simulation studies, production data from the reservoir is available. The production data is comprised of oil, gas and water rates and wellbore and reservoir pressures. Once the reservoir model is completed, flow simulation is run on the model and the resulting rates and pressure are compared with field data. If the two are different (they almost always are), then changes are made to the reservoir model until both data sets match. The process of iteratively comparing data and changing the reservoir model is termed history matching.

History matching has been applied to petroleum reservoir models from the time reservoir models were first being built. Actually before numerical simulators were available, reservoir engineers performed "history matching" by adjusting resistors and capacitances in electrical models (Wahl *et al.*, 1962). Traditionally, history matching has been a heuristic trial and error process to validate the model. After each flow simulation, the engineer must compare the output of the simulator with previous results and change the model parameters accordingly. This type of history matching is often termed "manual". The success or failure of the history matching endeavor largely depends upon the experience, fortitude and sometimes good fortune of the engineer.

Procedures have been developed to assist manual history matching by attempting to systemize the trial and error process (Willians & Keating, 1997; Mattax & Dalton, 1990). These approaches provide workflows to follow when history matching and include guidelines for what to change for each type of data (rate and pressure) that is not matching; however, the evaluation of the results and the adjustment of the model is still manual.

From nearly the beginning of history matching, the desire to computerize at least some aspects of the history matching process has been present (Jacquard & Jain, 1965; Coats *et al.*, 1970). Since this time, numerous techniques have been introduced, and more

recently, the research effort has been to automate the history matching process. The different types of automatic history matching methods are divided generally into the following two categories:

(1) gradient methods, and

(2) stochastic methods.

The idea for the gradient methods comes from optimal control theory. For these types of algorithms (e.g Quasi-Newton, Conjugate Gradient, and Levenberg-Marquardt), derivatives are required (Gill *et al.*, 1986). In history matching context, the derivatives of the production response with respect to changes in the reservoir parameters are needed. Derivatives (gradients) are used to determine the direction and the amount to change the parameters. Calculating the derivatives has been and still is a major topic of research. In the 1970s two distinct methods to efficiently calculate gradients were developed. Works by Chen *et al.* (1973) and Chavent *et al.* (1973) provide methods to generate the necessary gradients based on the adjoint equation. At the same time, Carter *et al.* (1974) and Hirasaki (1975) developed methods for calculating gradients based on sensitivity coefficients. The efficiency and practically of the sensitivity coefficient methods for multi-phase petroleum reservoirs was more fully established by Anterion *et al.* (1989) and Bissel *et al.* (1992).

One of the first attempts to include geological (geostatistical) information in the petroleum inversion techniques was reported by Fasanino *et al.* (1986). However, previous methods to honor geology while inverting for dynamic data were developed in the hydrology field (Delhomme, 1979). These methods combine the optimal control theory with the concept of kriging and pilot points in the reservoir. It is too expensive (and redundant) to calculate gradients with respect to properties in every gridbock. Thus the derivatives are calculated at only a few locations (pilot points) in the reservoir, and the results are interpolated to the remaining locations with an interpolator like kriging.

Wen *et al.* (1998) further developed this method to allow more geologic realism at a finer scale in a method termed Self Sequential Calibration (SSC).

Landa (2001) extended the method of Anterion *et al*. (1989) to account for geostatistical parameters such as variogram range, direction of maximum continuity, and global mean. Landa *et al*. (1996) and Phan & Horne (2002) used object based geostatistical modeling to try to perturb large scale geological continuity such as channels. This was done by working with the properties of the object model like channel sinuosity and thickness.

Oliver (1994) proposed an approach to compute permeability and porosity from pressure transient data and prior information. He was able to include the geostatistical information in the least-square formulation of the objective function following a Bayesian method developed in geophysical inverse theory (Tarantola, 1987). This work has been extended to account for 3D three-phase reservoirs, and it can match water-oil ratios and gas-oil ratios as well as pressure (Li *et al*., 2003). Another method based on the Bayesian framework has been proposed to integrate historical production data into reservoir models by Tjelmeland & Omre (1997) and Omre (2000).

The other main class of history matching generally is called stochastic search, and the primary advantage of these methods is the relatively simple integration of static data. In these methods, gradients are not used, and only the evaluation of the forward model (i.e. flow simulation) is required. The other main attraction of these methods is their theoretical ability to converge to the global minimum, whereas gradient methods may become trapped in local minima.

Traditional stochastic optimization methods such as simulated annealing (Aarts & Korst, 1989) and genetic algorithms (Goldberg, 1989) have been applied on petroleum engineering problems. Sultan *et al*. (1994) attempted to use parallel computing to speed up the simulated annealing routine of Ouenes *et al*. (1993). Datta-Gupta *et al.* (1994) also applied simulated annealing, and later Datta-Gupta *et al*. (1995) improved the efficiency of their algorithm by using a streamline simulator; nonetheless, over 5000

streamline simulations were required for a 500 cell model. Sen *et al.* (1992) and recently Carter & Ballester (2004) presented cases where genetic algorithms are used in reservoir modeling.

The next algorithm is still classified as a stochastic search method, but it is implemented in a manner that improves the efficiency. It is termed gradual deformation and was introduced by Roggero & Hu in 1998. In this method two Gaussian realizations are combined to create a realization that matches the production data better that the initial two realizations. The combinations are parameterized with a single parameter, consequently the large high dimensional optimization is reduced to a much simpler 1D optimization. It is a stochastic search routine because the newly created realization can be further combined with another equiprobable realization to find an even better match, and the process is repeated in a stochastic manner until the match is satisfactory.

Almost all reservoirs have multiple wells and usually some wells are matched easier than others. Only areas around the non-matching wells will need to be modified. In this case the single parameter combination will not suffice because it modifies the model by the same amount everywhere and may destroy the match already achieved at some wells. Thus, the gradual deformation method has been extended to allow different amounts of deformation in different sections of the reservoir model and have multiple parameters, one for each region of the reservoir (Hu *et al*., 2001).

Streamline-based methods for history matching are appealing for two reasons. First, the forward simulation is potentially much faster than conventional simulation methods for displacement-type problems. Second, time-of-flight information along the streamlines can be used to find sensitivity coefficients in an efficient and elegant manner (Wang & Kovscek, 2000; Agarwal & Blunt, 2001). However, the spatial variability of the permeability and porosity was not respected, and often the streamline locations were apparent in the history matched models. Caers et al. (2002) overcame this problem by updating the model in a geostatistical framework.

Simulating flow data is usually the most time consuming aspect of any method, so very fast approximate simulators (e.g. streamlines simulators or very coarsely gridded simulators) provide the opportunity to speed up this part of the process by several orders of magnitude. We already mentioned work by Datta-Gupta *et al.* (1995) to try to speed up simulated annealing. Another popular use of "proxy" simulators is to rank a large number of geostatistical realizations (Wang & Kovscek, 2002; Lodoen & Omre, 2004), and then choose the models that best match the production data from the proxy simulators.

Automatic history matching has been around for almost 40 years, yet its use has mainly been restricted to university research and industry research centers, and it is primarily demonstrated on synthetic examples. Real field applications have been limited to Bissel *et al.* (1994) and a few very recent papers (Feraille *et al.*, 2003; Qassab *et al.*, 2003; Gross *et al.*, 2004). In field-operating companies, history matching continues to be done almost exclusively with a manual trail and error method.

1.4 Limitations of the Available Methods

Manual history matching takes a significant amount of human time to interpret the results, observe where and what to change, and make the changes in the model. Additionally, it is very difficult to honor geologic information when these types of changes are implemented. Often in manual history matching methods, petrophysical properties in the non-matching regions are simply multiplied by a constant until a match is achieved (Milliken *et al.*, 2001). While an impressive match can be found, any previous knowledge of geology continuity has been destroyed, and the previously stated goal of accurate future predictions will not likely be accomplished. Despite its limitations, even today, the trial and error method is the most common form of history matching used for real field applications.

Gradient-based methods tend to converge quickly but are susceptible to local minimum; thus, the initial guess needs to be close to the optimum. Additionally, it is often time consuming to calculate gradients, and for many discrete parameters such as channel facies or fault locations, gradients are simply inappropriate. Furthermore, most gradient-based methods do not rigorously ensure that the prior geologic information is maintained.

Full thoroughness requires that the derivative of every response with respect to every parameter be calculated. In reservoir flow simulators there can be thousands of responses and millions of parameters making it impossible to determine every derivative. Currently there is a large amount of work in simplifying assumptions to get the number of derivatives to a reasonable amount. The simplifications have limited the advantage of using derivatives in the first place, and these methods have shown little success in real field applications.

Bayesian methods use gradients, thus they suffer from the same problems as gradient methods; however, Bayesian methods have additional issues as well. For mathematical tractability, the parameters in the Bayesian methods are often assumed to be multivariate

Gaussian, and measurement errors must be known and are usually assumed to be independent. These limiting conditions mean that Bayesian methods can only be applied in a limited number of very specific cases, many of them academic.

In all of the simulated annealing and genetic algorithm methods presented, too many flow simulations are required to be practical for real reservoir studies. Gradual deformation methods improve on the efficiency, but they have only been applied in a history matching context for continuous parameters such as permeability. Le Revalec-Dupin and Hu (2004) state that discrete objects such as nodules or channels cannot be history matched with the gradual deformation method.

Streamlined-based methods are fast but are limited to the robustness of streamlines (i.e. they work well for waterfloods), and they have not been applied to many discrete parameters such as channel facies or fault locations. Other approximate reservoir flow simulators (proxies) have been used extensively in history matching applications. Proxies can only be applied in limited situations, but they are only "approximate" to finite difference simulators, so they will not replace or dramatically alter traditional history matching. However, their use can aid some aspects of history matching, and they are used to some extent in the current work.

Issues with history matching can be summarized by the following three points:

- History matching is often performed without any regard to geologic data. This limits the predictive capabilities of the "matched" models.
- It is difficult to incorporate large scale parameters in an efficient manner. However, the large scale parameters often have the most significant impact on production data.
- Most current automatic history matching methods are not practical for real field applications. They require too many flow simulations or can not handle the geologic complexity found in most reservoir models.

1.5 Proposed Method

The primary objective of this work is to develop a practical history matching method that will generate reservoir models constrained to all types of information. The reservoir models will match the historical production data as well as honor seismic data, well log and core data, and especially large scale geologic information. Results from a model that honors all available data will undoubtedly have a better chance at predicting reservoir behavior than models where only part of the information is maintained.

Large scale properties do have uncertainty, and perturbing them often has a greater effect on production data than the fine-scale properties (Selari *et al.*, 1992). This work will primarily focus on perturbing large scale parameters such as facies distributions and proportions in an efficient manner using a history matching technique called the probability perturbation method (PPM). One aspect of the PPM is that the model perturbations are done not directly on the model parameters but on the underlying probabilities used to generate the model. This allows the geology of the reservoir to be honored in all realizations.

The current method is not presented as a method to solve all history matching problems. Although we examine some simple cases where fractures or continuous properties (permeability) are matched, other works have much more detail for faults (Suzuki & Caers, 2004), fractures (Voelker *et al.*, 2003; Suzuki *et al.*, 2005), and continuous parameters (Caers, 2004), and there are different history matching techniques that can honor geology and work well for certain reservoir model parameters such as gridblock permeability (Gross *et al.*, 2004). The current work is mainly aimed at reservoirs where uncertainty associated with *facies distributions* dominates. This work will primarily examine the uncertainty associated with random seeds and reservoir properties; the geologic scenario uncertainty will only briefly be discussed (i.e. the geologic scenario will assumed to be known). Furthermore, for this work, the geocellular models are built with the same grid used for flow simulation. However, this history matching technique

has been used when the geocellular model is on a much finer grid than the flow simulation model (Tureyen, 2005).

Applying the method to real world history matching problems will not be overlooked in this work. On the contrary, this work will focus on creating a method that can successfully be applied to actual field cases. The method will be demonstrated on two large real field case studies (with real production data). It is the author's opinion that showing the method can be applied to real field applications is as important as showing the method is theoretically well founded.

1.6 Thesis Outline

The current Ph.D. dissertation is comprised of six chapters that fully describe a history matching method designated the Multi-Parameter Probability Perturbation Method (MP-PPM). All aspects of the method are developed; the theoretical foundation for the method is established, the algorithm is described in detail and backed up with numerous demonstrative synthetic examples. Two case studies demonstrate the practicality of the method by showing how it can be applied to real field examples.

The probability perturbation method (PPM) is described in Chapter 2. It starts with a review of the geostatistical techniques that are required for the method. Then a complete description of how the method works including the algorithm is given. The theoretical ties of the probability perturbation method with Bayesian inversion techniques and Markov chain Monte Carlo methods are presented. The chapter concludes by discussing some of the limitations of the PPM and the need to extend the method so it can be applied to realistic field applications.

The extensions of the method are described in Chapter 3 along with applications that exemplify the more robust multi-parameter method. The chapter begins by describing the multi-parameter method and its advantages and by showing that the method does not lose any of the advantages of the previous method. A simple 2D example is used to explore numerous aspects of the method and to compare its behavior to other methods. Additional 2D and 3D examples demonstrate the flexibility of the method; it can work with other data and can be parameterized by more than one approach.

Chapter 4 validates the method on two 3D examples with channel reservoirs. The first is a synthetic example called Stanford V that is based on a North Sea channel reservoir. The MP-PPM is able to match the Stanford V data quite nicely in relatively few flow simulations. Furthermore, a comparison with a traditional history matching technique illustrates the advantage (better prediction) of having the geology honored in the history

matched model. The second example is a reservoir model that is currently being used by a major oil company for an offshore African field. This example demonstrates how the method matches real reservoir models that contain noisy field data. Many of the practical aspects of history matching are illuminated through this example. Additionally, we show how seismic data (along with well and geologic data) is included in the workflow.

The multi-parameter probability perturbation method is generalized further in Chapter 5. The chapter opens by discussing that perturbing only the random seed in a geostatistical routine often is not sufficient to achieve a history match, and in many cases, the properties of the geostatistical routine must be perturbed also. One particular property, facies proportion, is presented as an important parameter that can be perturbed in this manner, and the idea is confirmed with a 2D synthetic example. A real case study from the North Sea shows how this would work in practice by perturbing the proportion of calcite nodules in the model to match the production.

The final chapter discusses the major findings of this work, and their potential impact on history matching and petroleum engineering in general. Additionally Chapter 6 puts forward some ideas about the future of history matching, what the main goals of upcoming history matching research should be (i.e. the holy grail of history matching) and some possible paths to attain them.

"Just as our eyes need light in order to see,
our minds need ideas in order to create."

- Napoleon Hill

Chapter 2

Probability Perturbation Method

The probability perturbation method is a new approach to build reservoir models constrained to both static and dynamic data. The method is able to perturb large scale structures such as facies or faults as well as continuous properties such as porosity or permeability, and no discontinuities are created during the perturbations. A very simple example explores the fundamentals of the method, and a more complex example demonstrates the applications of the method. Additionally, the main purpose of this chapter is to develop formally the theoretical foundations of the method based on Bayesian inversion and Markov chains. The chapter begins with a small background of geostatistics and describes the main geostatistical tools used in the dissertation.

2.1 Geostatistics

Geostatistics as applied to reservoir modeling is focused on integrating data from various sources into a consistent model that is used to study the behavior of the reservoir. Traditionally, geostatistics is thought of as a way to combine the static data but not the dynamic data. The current work maintains that distinction in terminology (although it will not in practice). When static data is being combined, the discussion is of geostatistics and geostatistical algorithms, and when the dynamic data is being incorporated, history matching terminology and language is used. As stated in the introduction, however, the main point of this work is to combine these two seemingly diverse aspects into one consistent method.

Geostatistics is a relatively young field, only about half a century old. It started in the mining industry (Krige, 1951; Matheron, 1963; Journel & Huijbregts, 1978), and it later spread to the petroleum industry (Haas, 1976). Kriging was developed as an alternative to hand contouring or inverse distance weighting to generate maps of reservoir properties. However, it was recognized early that kriging maps were often too smooth and did not account for the proper connectivity for petroleum reservoirs. Consequently, sequential simulation techniques were developed (Journel & Alabert, 1988) to introduce more realistic correlation structures into reservoir models.

Sequential simulation accounts for most of the reservoir characterization that is completed in the petroleum industry today; moreover, it will be used exclusively for the current work. Therefore, the generalities of sequential simulation are described in greater detail, and a few specific algorithms are developed more completely.

All types of geostatistical algorithms incorporate the same idea: given a set of known data, $Z(\mathbf{u}_\alpha)$, a grid of unknowns, $Z(\mathbf{u})$, is populated. The term $\mathbf{u} = (x, y, z)$ denotes the location of a grid node not necessarily in a regular grid. These unknowns could be porosity values, facies indicators or any other geologic reservoir property to be

determined. The paradigm of sequential simulation is *how* the unknowns are estimated. Each node that does not have a property value is visited randomly, and a conditional probability, P(A|B), is estimated for that node. P(A|B) is the probability of the unknown property, A, occurring given some other information B. For example, A could stand for "channel occurs" or "permeability is less than 100 millidarcy (md)" and B could be well-log data and/or geologic knowledge of the reservoir. Then, one randomly draws from that probability distribution to get the property value for that node. Once a node has been simulated, its value is used with the known data to create the conditional probability for the next node. This is continued until properties fill the entire model. Because the nodes simulated previously are retained for simulating all subsequent nodes, the spatial relations of the unknown values (up to the statistics of the algorithm) are honored. Three desirable properties of sequential simulation are:

(1) it reproduces the desired statistics,

(2) multiple realizations are easily created by changing the random path, and

(3) integration of different types of data is straightforward.

Numerous sequential simulation methods have been developed over the last 15 years; however, the only difference between most of them is how P(A|B) is determined. Three of the most common sequential simulation algorithms include:

- sequential Gaussian simulation, sgsim, (Gomez & Journel, 1989), after transferring the variable to a Gaussian space, P(A|B) becomes a Gaussian distribution with mean equal to kriging mean and variance equal to the kriging variance. This technique is mostly used for generating petrophysical properties within facies.

- sequential indicator simulation, sisim, (Journel & Isaaks, 1984), perform indicator kriging at various thresholds at each unknown location to determine P(A|B). This technique is mostly used to generate facies realizations.

- single normal equation simulation, snesim, (Strebelle, 2002), P(A|B) is built by scanning a training image for equivalent data events A|B. This method allows generation of a wide variety of geological scenarios and goes beyond the variogram approaches of sgsim and sisim.

snesim is the latest development in sequential simulation. Traditional sequential simulation methods rely on two-point statistics (variograms) to determine the degree of correlation in the model (e.g. sgsim). However, variograms can not reproduce well-connected curvilinear features such as channels because they only use the correlation of two points in space. Therefore, new methods have been developed that use multiple-point statistics and training images to define correlations that are geologically more complex. Training images are non-conditional and purely conceptual depictions of the geological patterns deemed relevant for a particular subsurface. The training image is used to extract geological patterns, in terms of multiple-point statistics, not limiting oneself to a variogram (two-point statistic). These statistics, just as with the variogram, are anchored to the subsurface well and seismic data. However, contrary to variogram and kriging-based methods, the probability model, P(A|B), is directly inferred by scanning the training image for equivalent data events A|B. Subsequent examples and case studies of the proposed method utilize multiple-point statistics, but the method is not limited to multiple-point statistics or even sequential simulation. Any technique that uses conditional probabilities can be history matched using the probability perturbation method.

One of the desirable properties of sequential simulation is its ability to incorporate different types of data. Journel (2002) has demonstrated an elegant yet efficient method to integrate soft or secondary data, C, with the previous probability model, P(A|B) to get an updated conditional probability model, P(A|B,C). Each simulated value should depend not only on the primary data, B, but also on the secondary data, C. Before incorporating the secondary data into the model, it needs to be calibrated with the

primary data; such calibration results in a probability model, P(A|C). For example if the secondary data is seismic amplitude of a channel reservoir, P(A|C) would be the probability of A = "channel occurring" at a location given some local secondary information, C. To combine P(A|B) and P(A|C) into P(A|B,C), Journel uses the following expression based upon an improved form of conditional independence.

$$\frac{x}{a} = \left(\frac{b}{a}\right)^{\tau_1} \left(\frac{c}{a}\right)^{\tau_2} \tag{2.1}$$

where:

$$x = \frac{1 - P(A \mid B, C)}{P(A \mid B, C)}, \quad b = \frac{1 - P(A \mid B)}{P(A \mid B)}, \quad c = \frac{1 - P(A \mid C)}{P(A \mid C)}, \text{ and } \quad a = \frac{1 - P(A)}{P(A)}.$$

P(A) is the global proportion of A occurring; therefore, a can be interpreted as the distance to the event A occurring, prior to knowing the information carried by B or C. If P(A) = 1, then the distance $a = 0$ and A is certain to occur. Likewise the values b and c state the uncertainty about the occurrence of A given information B and C, respectively, and x is the uncertainty when both B and C are known. The parameters τ_1 and τ_2 account for redundancy between the data B and C, but in this work, they are chosen as $\tau_1 = \tau_2 = 1$, which amounts to a form of standardized conditional independence. The combined probability model, P(A|B,C), for unit tau values is equal to the following expression.

$$P(A \mid B, C) = \frac{1}{1 + x} = \frac{a}{a + bc} \tag{2.2}$$

2.2 Description of the Method

For demonstration purposes, we consider only the case of a binary spatial variable described by an indicator random function model

$$I(\mathbf{u}) = \begin{cases} 1 & \text{if a given facies occurs at } \mathbf{u} \\ 0 & \text{else} \end{cases}$$

where $\mathbf{u} = (x, y, z) \in$ reservoir model, R, is the spatial location of the node, and $I(\mathbf{u}) = 1$ means, for example, that channel occurs at location $\mathbf{u}$, while $I(\mathbf{u}) = 0$ indicates non-channel occurrence. Given a conditional probability, P(A|B), at all grid nodes, $\mathbf{u}$, an initial realization is generated using sequential simulation and is termed $i^{(0)}(\mathbf{u})$. This realization is conditioned to the static data, but it does not honor the dynamic data. In all forms of history matching, an initial realization is created and subsequently changed until a satisfactory match with the production data is achieved, and in this respect the probability perturbation method is no different. The unique aspect is how the initial realization is changed (perturbed). Rather than perturbing the initial realization directly, we propose to perturb the probability model, P(A|B), used to generate the initial realization. This is done by introducing another probability model, P(A|D), that depends on dynamic data, D. The perturbation of P(A|B) by P(A|D) is achieved by combining both conditional probabilities using Journel's method (Eq. 2.2). The resulting new probability model, P(A|B,D), is used to populate the next realization. P(A|D) is defined as follows:

$$P(A|D) = (1-r_D)\, i^{(0)}(\mathbf{u}) + r_D P(A) \in [0,1] \qquad (2.3)$$

where r_D is a parameter between [0,1] that controls how much the model is perturbed. To better understand the relationship between r_D and P(A|D) consider the two limiting cases when r_D=1 and r_D=0. When r_D=0, P(A|D) = $i^{(0)}(\mathbf{u})$ and the initial realization, $i^{(0)}(\mathbf{u})$, is retained in its entirety, and when r_D=1, P(A|D) = P(A) and a new equiprobable realization, $i^{(1)}(\mathbf{u})$, that is equally probable as $i^{(0)}(\mathbf{u})$. The parameter r_D, therefore, defines

a perturbation of an initial realization to another equiprobable realization. To illustrate the role of the free parameter, r_D, Figure 2.1 shows six different realizations, $i^{(r_D)}(\mathbf{u})$, for different r_D values. For $r_D = 0$, the realization is exactly the same as the initial realization, $i^{(0)}(\mathbf{u})$, and for small r_D values the models are very similar. When $r_D = 1$ the new realization appears very different from the initial realization, and when $r_D = 0.5$ the new realization appears as a blend of the initial realization and the equiprobable realization.

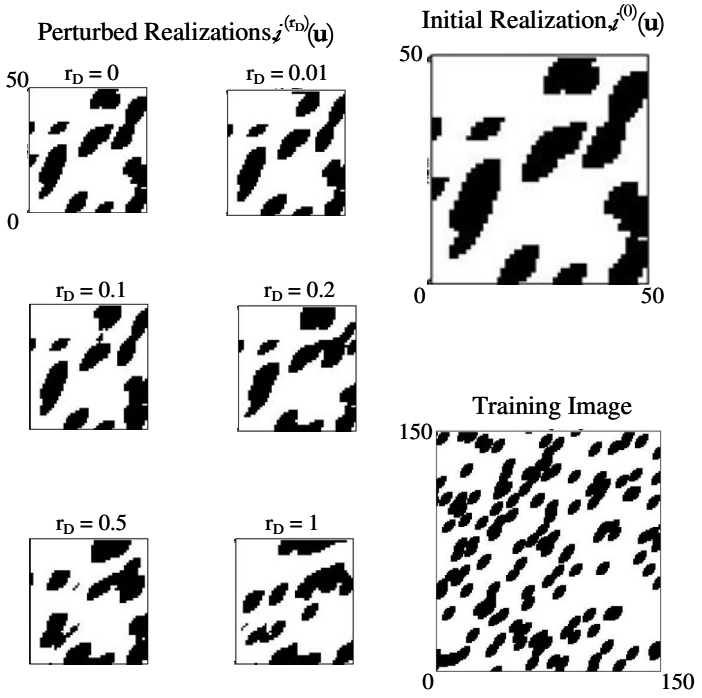

Figure 2.1: The initial realization and six perturbed realizations with perturbation parameter, r_D, values ranging from 0.0 to 1.0 are displayed. All realizations are generated from the current training image.

There may exist a value of r_D, such that $i^{(r_D)}(\mathbf{u})$, matches the production data better than the initial realization. Finding the optimum realization, $i^{(r_{Dopt})}(\mathbf{u})$, is a problem parameterized by only one free parameter, r_D; therefore, finding the optimum realization is equivalent to finding the optimum r_D value.

$$\min_{r_D}\{O(r_D) = \| D^S(r_D) - D \|\} \tag{2.4}$$

where $O(r_D)$ is the objective function and is defined as some measure of difference between the simulated production data, $D^S(r_D)$ and the observed field data, D. Note that the simulated production data is different for each geostatistical realization, $i^{(r_D)}(\mathbf{u})$, The value of r_{Dopt} and consequentially the optimum realization, $i^{(r_{Dopt})}(\mathbf{u})$, can be selected from any one-dimensional optimization routine (e.g. the Brent method (Press *et al.*, 1989)).

The best realization between two equiprobable realizations may not yet represent an acceptable match to the production data. The search space is too limited because there are only two realizations. Thus, a two-loop optimization routine is necessary to introduce new details into the history matching process. The previous optimum realization, $i^{(r_{Dopt})}(\mathbf{u})$, is used as the initial realization in the next step by replacing $i^{(0)}(\mathbf{u})$ in Eq. (2.3). Also during each outer loop, the random seed is changed. By doing so, a new equiprobable realization is generated when $r_D = 1$. This allows the method to again search between two equiprobable realizations in the next inner iteration. The outer iteration is stopped when a satisfactory history match is achieved (i.e. the error is below a tolerance ε). In summary, there is an inner iteration that finds the optimum realization between two equiprobable realizations and an outer iteration that converges on a history match by continually adding new information.

Algorithm

A summary of the single parameter probability perturbation method SP-PPM is provided.

- Define a training image depicting the desired geological continuity
- Generate an initial realization, $i^{(0)}(\mathbf{u})$ for $\mathbf{u} \in R$
- Change random seed
- Outer loop: until history match, $O(r_D) < \varepsilon$:

- Inner loop: perform a 1D optimization to get the r_D that results in the best match to the production data, D. The objective function is evaluated as follows:
  - Calculate $P(A|D) = (1-r_D)i^{(0)}(\mathbf{u}) - r_D P(A)$ and combine it with P(A|B) to get the probability model, P(A|B,D) using Eq. 2.2.
  - Generate new model $i^{(r_D)}(\mathbf{u})$ and run flow simulation to get $D^S(r_D)$
  - Calculate objective function, $O(r_D) = \| D^S(r_D) - D \|$
- If $O(r_D) < \varepsilon$ – end; else – change random seed and set $i^{(0)}(\mathbf{u}) = i^{(r_{Dopt})}(\mathbf{u})$

Multiple Facies and Continuous Properties

Although the PPM has only been discussed in terms of two facies, the method is applicable to multi-facies and continuous variables with little change. The extension to the multi-facies case is straightforward. Rather than use Eq. (2.3) to define P(A|D), the probability is defined for each facies with the following equation:

$$P(A_m|D) = (1-r_D)\, i^{(0)}(\mathbf{u},s_m) + r_D P(A_m) \in [0,1] \tag{2.5}$$

where the indicator denotes that

$$I(\mathbf{u},s_m) = \begin{cases} 1 \text{ if a certain facies, } s_m, \text{ occurs at } \mathbf{u} \\ 0 \text{ for all other facies} \end{cases}$$

and A_m denotes that $\{I(\mathbf{u},s_m) = 1\}$, and m is the facies: 1, 2 … M where M is the total number of facies. When generating a multiple facies realization, at every gridblock there is a probability of each facies occurring. The PPM applies a perturbation to all of these facies probabilities. If the perturbation parameter, r_D, is small, all facies are perturbed by a small amount, and if r_D is large, all facies are perturbed by a large amount.

The continuous case while similar to the discrete case requires a little more alteration. The similarity is that production data is used to create soft data that conditions the next

realization. Instead of calculating P(A|D) and combining it with prior data using Eq. (2.2), collocated co-simulation is performed using a correlation coefficient, ρ. In this instance the soft data is simply the previous "best" realization. Similar to the perturbation parameter, the correlation coefficient quantifies the amount by which this current best realization is perturbed. Actually, the correlation coefficient and the perturbation parameter are related through $\rho = 1 - r_D$. When $\rho = 1$ ($r_D = 0$), the soft data is honored exactly, and the current best realization is reproduced. Conversely when $\rho = 0$ ($r_D = 1$), the soft data provides no information, and the realization is maximally perturbed into a new equiprobable realization. The continuous property case is a probability perturbations method because the soft data "perturbs" the probability distributions (i.e. the ccdfs) used to calculate the gridblock properties. A more detailed explanation and derivation of the probability perturbation method applied to continuous cases is provided in Suzuki *et al.* (2005).

A Simple Illustrative Example

A very simple example is presented to clarify the approach and to illustrate various properties of the method. The model consists of a grid with only three nodes, $\mathbf{u}_1$, $\mathbf{u}_2$ and $\mathbf{u}_3$. Each node can be either black, $I(\mathbf{u})=1$ or white, $I(\mathbf{u})=0$. The model is therefore simply $\{I(\mathbf{u}_1), I(\mathbf{u}_2), I(\mathbf{u}_3)\}$. There are a number of data in the example, and the first data is a point measurement, namely, $i(\mathbf{u}_2)=1$ (a black pixel in the middle). This is equivalent to a hard data measurement from a well in a petroleum reservoir. The spatial dependency of this 1D model is described by the training image shown in Figure 2.2. The training image and the data point comprise the "static" data for the model (i.e. they are the B data).

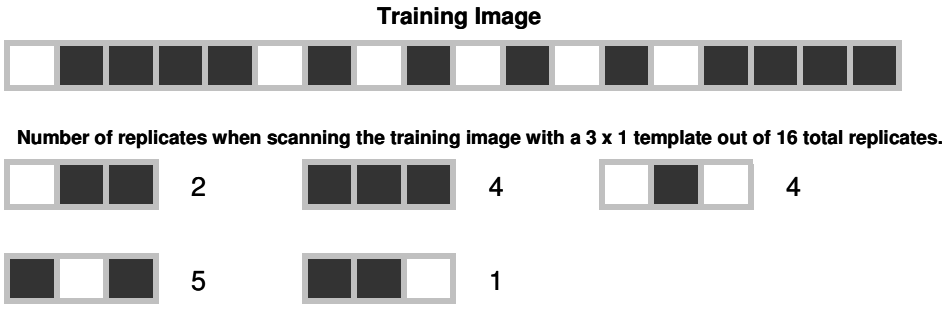

Figure 2.2: Training image and number of 3 x 1 replicates derived from the training image.

The other piece of information is $I(\mathbf{u}_1) + I(\mathbf{u}_2) + I(\mathbf{u}_3) = 2$. This obviously is not dynamic data in the sense of a petroleum reservoir; however, for this illustrative example, it will correspond to the "dynamic" data in the method (i.e. this is the D data). The objective function is provided by the following equation.

$$O = \left| \sum_{k=1}^{3} i(\mathbf{u}_k) \; - \; 2 \right| \tag{2.6}$$

Since the value of the middle node is always one, the summation can only equal one, two or three; thus, the objective will be either zero or one. If the objective equals zero, the model will be matched, and the routine will be stopped; otherwise, the routine will continue until the objective is satisfied.

The first part of the PPM requires an initial model to be generated and sequential simulation is used to do this. A random seed is set, and all hard conditioning data are included. The hard data is put directly into the model and fixed throughout the simulation (e.g. $i(\mathbf{u}_2)=1$). All conditional distributions of the type Prob(A|B) are determined by scanning the training image with a 3 x 1 template. Throughout this example, the term 'A' will stand for "a black pixel occurring." There are only a few possible alternatives, so the entire sequential simulation process is reviewed. To begin any sequential simulation, first, the random path must be determined based on the given random seed; for this example, let's say the left node, $i(\mathbf{u}_1)$, is visited first.

One wants to know the probability of having a black pixel in the left node given that a black pixel occurs in the center node and given the information from the training image. This can be stated in probability notation as Prob($I(\mathbf{u}_1)$=1| $i(\mathbf{u}_2)$=1, TI). There are eleven replicates that have a black pixel in the center (Figure 2.2), and of those eleven, five have a black pixel in the left node:

$$\text{Prob}(I(\mathbf{u}_1)=1|\ i(\mathbf{u}_2)=1, \text{TI}) = 5/11 = 0.45$$

We randomly draw from this probability to determine if the left node will be black or white. If it is black, we calculate the probability that the right node is black given the left and center nodes are black, (Prob($I(\mathbf{u}_3)$=1| $i(\mathbf{u}_2)$=1, $i(\mathbf{u}_1)$=1, TI)). Of the five nodes that satisfy the condition, four have a black node at location $\mathbf{u}_3$:

$$\text{Prob}(I(\mathbf{u}_3)=1|\ i(\mathbf{u}_2)=1,\ i(\mathbf{u}_1)=1, \text{TI}) = 4/5 = 0.80$$

However, if the left pixel is white ($i(\mathbf{u}_1)$=0), there are six replicates and two of them have black pixels for their right nodes:

$$\text{Prob}(I(\mathbf{u}_3)=1|\ i(\mathbf{u}_2)=1,\ i(\mathbf{u}_1)=0, \text{TI}) = 1/3 = 0.33$$

The same analysis can be done starting with the right node ($i(\mathbf{u}_3)$), however, the final probabilities for the four possible models are the same. Figure 2.3 demonstrates all the different ways that the initial model can be built and the probabilities of each model occurring.

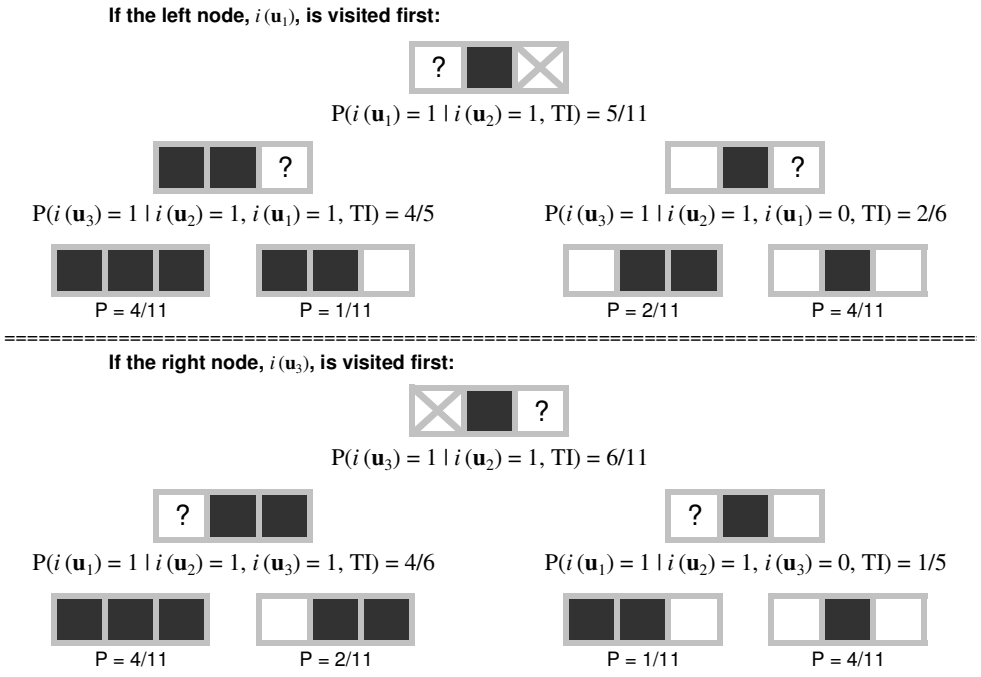

Figure 2.3: Building the initial model using sequential simulation.

Up to this point, only "static" data has been used to build the model. Next, the "dynamic" data is incorporated using the probability perturbation method. Since the D data has a complex and non-linear relationship with the model, an iterative technique is required to find the r_D and P(A|D) that create a realization that best matches the data.

An initial guess is created. Suppose the guess does not have two black pixels. Then the random seed is changed and a value for r_D is chosen between zero and one. The two conditional probabilities based on the D data (Prob($I(\mathbf{u}_1)$=1|D) and Prob($I(\mathbf{u}_3)$=1|D)) are now be calculated from Eq. 2.3. Sequential simulation will again be performed, but this time the conditional distributions will be of the type P(A|B,D). They are determined by combining Prob(A|B) and Prob(A|D) using Eq. 2.2, where Prob(A|B) is the same one used to generate the initial realization. Table 2.1 shows the probability values for four different possible r_D values for the 1D example.

Table 2.1: Conditional probability values for four different r_D values.

	r_D=0.05		r_D=0.25		r_D=0.50		r_D=0.95	
P(A)=0.67	P^0(A|D) = 0.033 P^1(A|D) = 0.983		P^0(A|D) = 0.167 P^1(A|D) = 0.917		P^0(A|D) = 0.333 P^1(A|D) = 0.833		P^0(A|D) = 0.633 P^1(A|D) = 0.683	
P(A|B)	P^0(A|B,D)	P^1(A|B,D)	P^0(A|B,D)	P^1(A|B,D)	P^0(A|B,D)	P^1(A|B,D)	P^0(A|B,D)	P^1(A|B,D)
0.45	0.01	0.96	0.08	0.82	0.17	0.68	0.42	0.47
0.80	0.06	0.99	0.29	0.96	0.50	0.91	0.78	0.81
0.33	0.01	0.94	0.05	0.73	0.11	0.56	0.30	0.35
0.55	0.02	0.97	0.11	0.87	0.23	0.75	0.51	0.56
0.67	0.03	0.98	0.17	0.92	0.33	0.83	0.63	0.68
0.20	0.00	0.88	0.02	0.58	0.06	0.38	0.18	0.21

*Superscripts 0 or 1 indicate whether initial value at the node is black (1) or white (0)

The manner in which P(A|B,D) is calculated is discussed for one r_D value, 0.25. P(A|D) is first calculated by substituting 0.25 for r_D in Eq. 2.3 to get P(A|D) = (1-0.25)$i^{(0)}$(**u**) + (0.25)P(A). The prior information about A occurring, P(A), is 2/3 for this example. This information is known before the model is constructed, and for consistency the training image in this example also has the same proportion. The value of the initial node (either zero or one) will determine P(A|D). If the initial value was zero, then P(A|D) is 0.167, and if the initial value was one, then P(A|D) is 0.917. The P(A|B) values were calculated while building the initial model (Figure 2.3), and they are listed in the left most column of Table 2.1. The P(A|D) values and the P(A|B) values are combined using Eq. 2.2 with $\tau = 1$, which be can expanded to:

$$\text{Prob}(A \mid B, D) = \frac{\left(\frac{1-P(A)}{P(A)}\right)}{\left(\frac{1-P(A)}{P(A)}\right)+\left(\frac{1-P(A \mid B)}{P(A \mid B)}\right)\left(\frac{1-P(A \mid D)}{P(A \mid D)}\right)}$$

For example, if the initial value is black (P(A|D) = 0.917) and the left node is visited first (PA|B) = 0.45), the probability that the node is black given both the B and D data is 0.82. If it is originally black, there is a high probability that it will stay black; similarly, if it was originally white, there is a low probability it will change to black (0.08). For each r_D

value, Table 2.1 lists the P(A|B,D) for the six P(A|B) situations and for the two possible initial values (0 or 1). If r_D is high, P(A|D) is close to the prior probability and P(A|B,D) is similar to its corresponding P(A|B) values; therefore, if the random seed is changed, the new realizations are much different than the initial. If r_D is low, both P(A|D) and P(A|B,D) are close to their indicator values (0 or 1), so the model is not expected to change much.

We draw randomly from the new probability distribution, P(A|B,D), to generate the next realization. If the sum $i(\mathbf{u}_1) + i(\mathbf{u}_2) + i(\mathbf{u}_3)$ equals two, we are done. If the sum doesn't equal two, we select a new r_D and repeat the process until the data matches. The entire process from generating the initial realization to matching the "dynamic" D data is displayed in Figure 2.4.

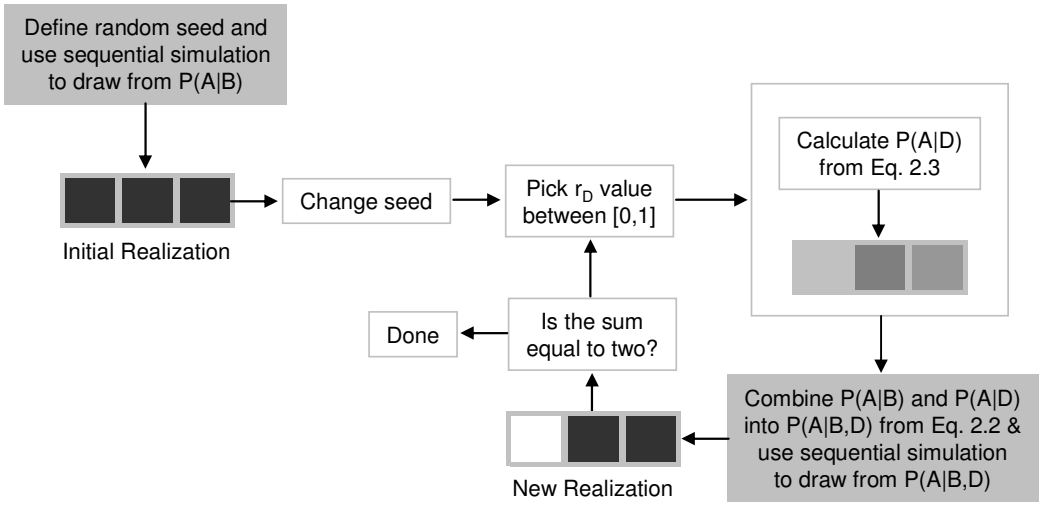

Figure 2.4: Probability perturbation flowchart for 1D example.

The only difference between this example and the way history matching is done with the PPM is how r_D is treated. In this example we choose randomly an r_D value; however, when matching history with the PPM, a 1D optimization method is used to find an r_D that generates a model that best reproduces production data. A history matching example is covered in the next example.

Before we go to the next example, however, one final aspect from the 1D example is considered. Given that the two conditions are satisfied, only two different realizations are generated: one where node 1 is black and node 2 is white and one where node 1 is white and node 2 is black. The question is posed: when the conditions are satisfied, what is the probability that node 1 is black? In notation, what is $P(I(\mathbf{u}_1) = 1 \mid i(\mathbf{u}_2) = 1, I(\mathbf{u}_1) + I(\mathbf{u}_2) + I(\mathbf{u}_3) = 2)$ or simply what is P(A|B,D)?

To determine the true answer (given the prior geologic model or training image is correct), a large number of models are matched to the two conditioning data with the PPM. Then we calculate how many of the matched models have a black pixel for node 1, and the fraction with a black pixel in node 1 is 0.36. The prior information in Figure 2.2 contains only three models that honor both conditioning data, and of those three only one has a black pixel for node 1. Therefore, the true probability is equal to 0.33. The small discrepancy is likely due to the assumption of conditional independence between the two data used in PPM. However, nothing in the PPM requires the assumption of conditional independence. If the dependence among the data is known, it can be accounted for in Eq. (2.2) through non-unity tau values (Journel, 2002; Krishnan, 2004). A more detailed discussion of this example can be found in Caers & Hoffman (2005).

Synthetic 2D Example

The first flow model example is a simple two facies model with 25 gridblocks in the x and y directions; the "true" reference model is shown in Figure 2.5A. A high permeability facies has a permeability of 750 md, and a lower permeability facies is 150 md. There are three wells, one injector in the center of the model (I1) and two producers at the corners of the model (P1 and P2). The majority of the high permeability facies is located between the injector and well P2; consequently, the water-cut profiles for the two production wells are significantly different (Figure 2.5B).

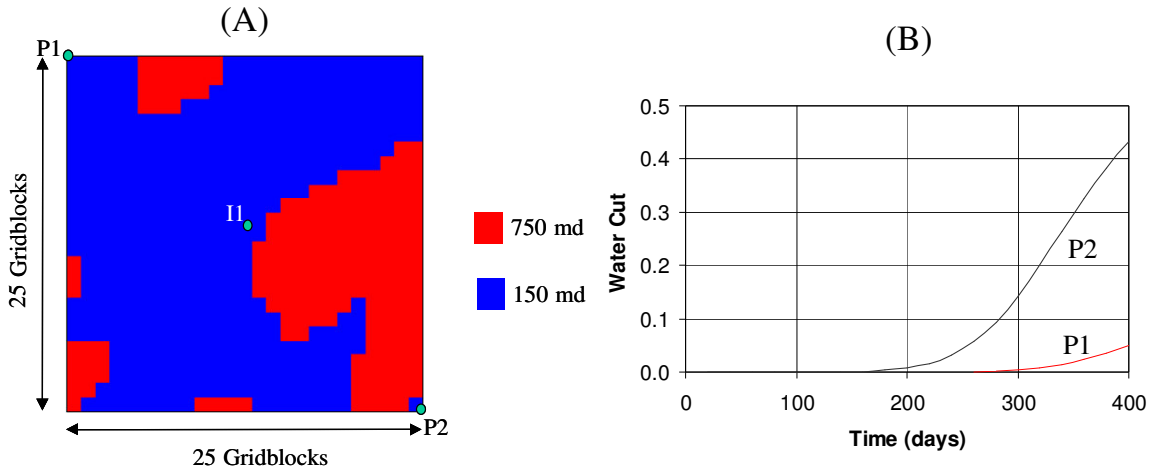

Figure 2.5: (A) The "true" reference model for the 2D example, and (B) the corresponding water cut curves for the two production wells.

The porosity in the model is constant for the two facies at 20%. The model is a two phase oil-water system. The viscosity of the oil is 1.0 centipoise (cp), and the water is 0.325 cp. The density of the oil is 49.94 lb/ft^3 and the water density equals 62.48 lb/ft^3. The oil-water relative permeability curves are provided in Figure 2.6.

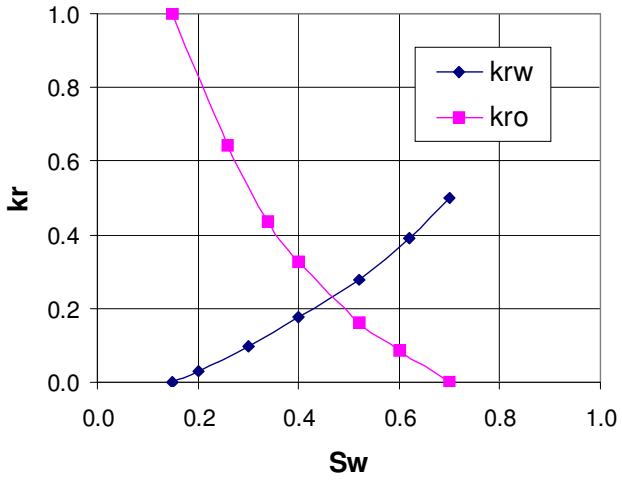

Figure 2.6: Relative permeability for simple 2D example.

The same technique used for the 1D model will be employed here as well. An initial realization is generated that is constrained only to the static data. For this example, the static data consists of the facies measurements from the three well locations (they are all low permeability facies) and the geological information provided by the training image (Figure 2.7).

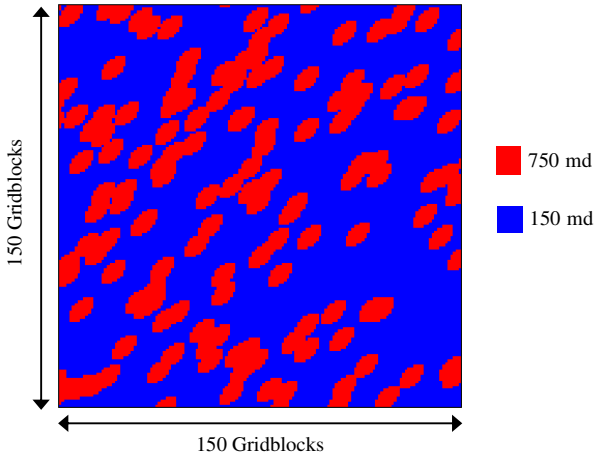

Figure 2.7: Training image for 2D example.

The dimensions of the training image in Figure 2.7 is 150 by 150, while the model is only 25 by 25, so the high permeability bodies appear larger in the model than the training image, but in fact the bodies have the same actual dimensions.

Once the initial model is constructed, its facies distribution is modified using the PPM. The other model properties (e.g. porosity, relative permeability etc.) remain fixed at their initial values. The perturbations continue until the model output matches the dynamic data. For this example, the water cut at the producers is the dynamic data. The standardized least square error between the reference water cut and the simulated water cut is calculated. Once the error is below 0.001, the procedure is stopped, and a model that honors both static and dynamic data is present.

All steps are essentially the same as the 1D example except an optimization technique updates the r_D values. For this example, the Brent method (Press, et al., 1989) is used. The advantage of Brent is that it does not require derivatives, and it is efficient and robust. Brent is a combination of a bisection method (Golden Search) and a parabolic interpolation method; the robustness comes from the bisection method, and the efficiency comes from the parabolic interpolation method.

The algorithm starts with an initial guess for r_D, and for this example, 0.50 is used since it is the midpoint of the range of r_D, [0,1]. An objective function value (mismatch) for this r_D value is required. A realization is generated for this r_D, and flow simulation is completed on the realization. The flow simulator, eclipse (Schlumberger, 2003), is used for this example. The mismatch is then calculated by comparing the simulated and reference water cuts. The Brent algorithm chooses the next r_D value by bisection, and the same process is used to determine its objective function value. The third point (r_D value) is also found by bisection, and its mismatch is calculated. Now, however, there are enough points to do parabolic interpolation. The remaining r_D values are determined by parabolic interpolation if it is well behaved. If the parabolic interpolation does not produce consistent results (e.g. a point lies outside the bounds), Brent then reverts back to bisection for that step. Once the Brent method has converged, the r_D value, the mismatch value, and the realization that best match the data are returned to the main program. (A general form of the computer code is provided in Appendix A.)

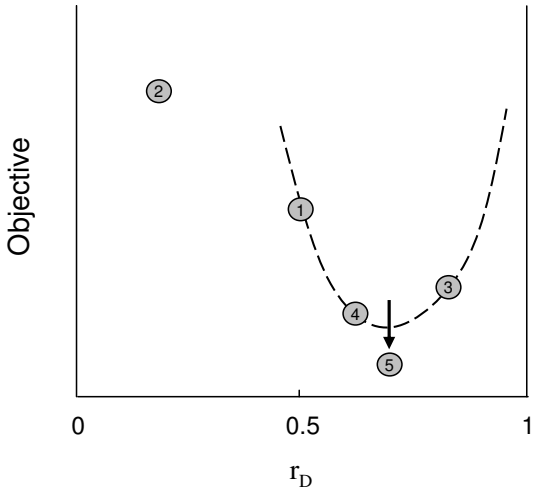

Figure 2.8: Convergence to a minimum using Brent. Bisection is used to determine points 2, 3 and 4 and parabolic interpolation through points 1, 4 and 3 is used to find the fifth point.

Figure 2.8 shows how the Brent algorithm might proceed to find the optimal r_D value for one inner iteration. The initial guess for r_D is provided at 0.5, and the objective for that point is calculated. Points 2 and 3 are obtained by bisection. Since the objective value at point 2 is higher than point 1, the third point explores the other half of the space. (If point 2 was less than point 1, the third point would be located between points 1 and 2.) If a parabola is generated through points 1, 2 and 3, at its minimum, r_D is greater than 1.0 (non-consistent behavior); therefore, the fourth point is also determined by bisection (between points 1 and 3). Now a parabola can be constructed through points 1, 3 and 4. At the minimum of the parabola, a new r_D is selected and used to calculate a new objective (generate a realization and run flow simulation), and this is the value at point 5.

If the lowest mismatch value is less than the tolerance (for this example, 0.001), the PPM ends and the realization is the history matched realization. However, in many instances, this realization will not provide a sufficient match because this part of the PPM only searches between two different realizations (i.e. the search space is small). In these instances, the current best realization becomes the initial realization in Eq. 2.3. The

random seed is changed, so we are again searching between two equiprobable realizations, and the process is repeated until an acceptable match is found.

A PPM flowchart for history matching the 2D example is provided in Figure 2.9.

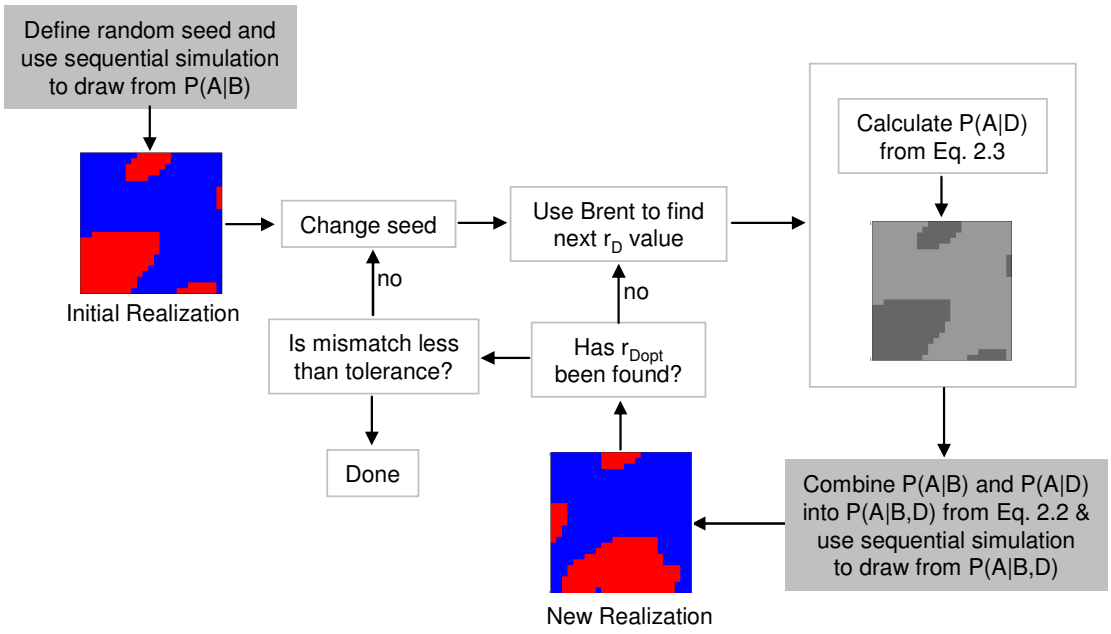

Figure 2.9: Probability perturbation flowchart for 2D example.

The probability perturbation method is successfully employed to history match the synthetic example. The reference water cut is reproduced almost exactly, and the matched realization honors the geology that is provided by the training image and hard data. Figure 2.10 has the initial realization (A) and the matched realization (B) along with the water cut data (C) for the three models (reference, initial and matched).

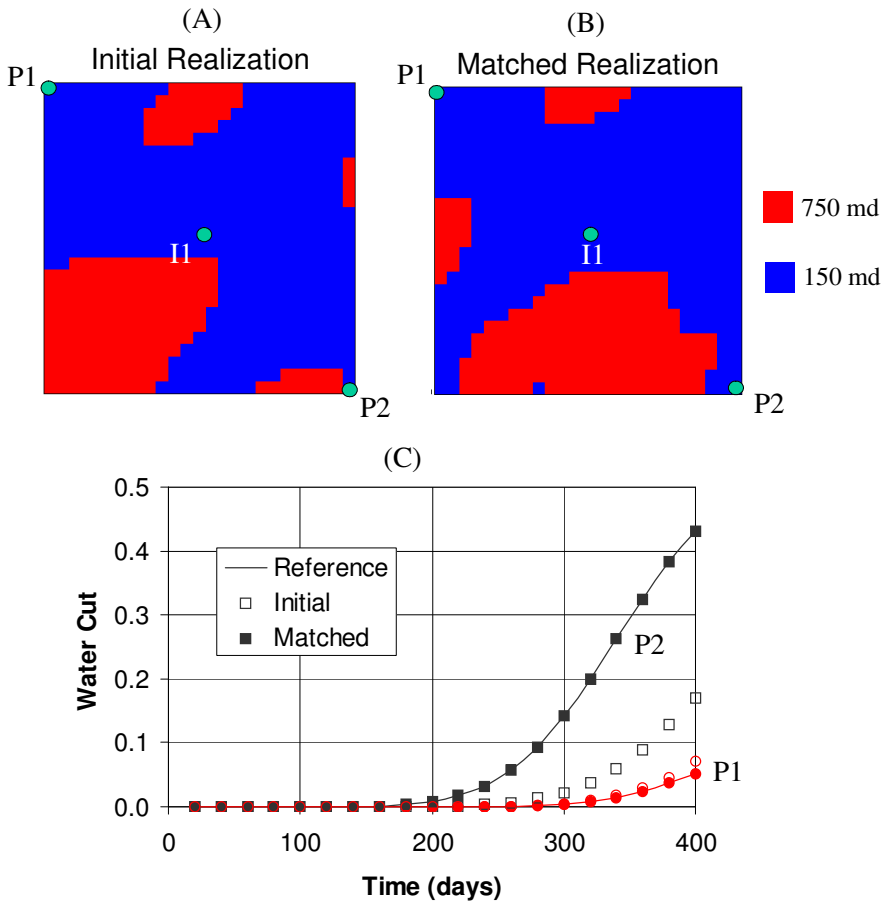

Figure 2.10: An initial and matched realization for 2D example, and water cut data for the reference, and the initial and matched realizations.

For well P2, the water cut data for the initial realization (open squares in Figure 2.10C) is much lower than for the reference. This is because the initial realization has very little high permeability facies between the injector and the P2 producer. Similar to the reference, however, the matched realization has much more high permeability facies between the injector and the P2 well.

The history match required 57 outer iterations and around 400 flow simulations to complete. Figure 2.11 shows the normalized objective function value for the outer iteration and for two inner iterations. For the outer loop, the objective function values

decrease rapidly for the first few iterations. However, it also has many long periods where the model does not improve (i.e. the best r_D is $r_D = 0$), and it takes a large number of iterations to converge the last small amount.

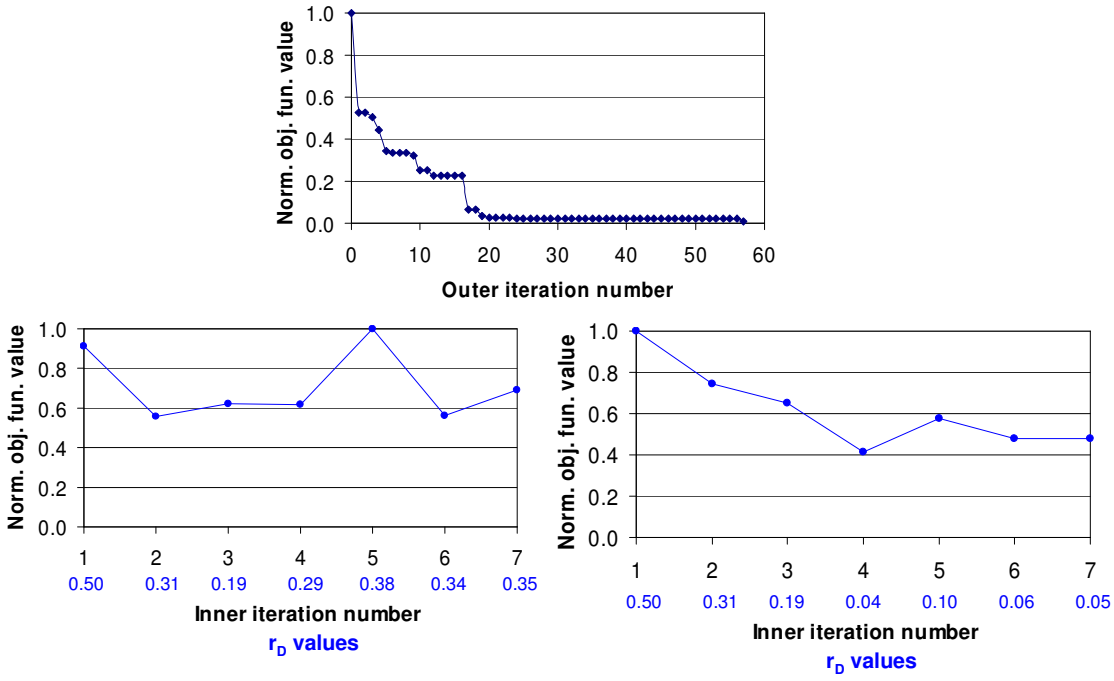

Figure 2.11: Objective function behavior for 2 realizations.

The objective function values for both inner iterations do not monotonically decrease. This is caused by the non-continuous nature of the bodies that are being perturbed. As the realization is perturbed, the bodies can appear, disappear, or be moved by a discrete amount. This can cause large changes in the production data and, consequently, in the objective function. However, it does not affect the overall efficiency of the method because of the two-loop routine. In the inner loop the last iteration is not retained; rather, the realization with the lowest objective function is retained and passed to the outer loop.

Both inner loop plots in Figure 2.11 have seven iterations; this is no coincidence. A sensitivity study was completed to determine a suitable number of flow simulations for the inner loop. In the study, all inner loops were allowed to run until they reached the

stopping criteria set in Brent. Most runs (63%) finished in 10 or less and the longest run required 15. However, the last iteration does not always have the lowest objective function value, so the inner loops were checked for the iteration step that reaches the minimum. Around 85% of the runs found their lowest value in 7 iterations or less; therefore, the maximum number of iteration is set at seven.

2.3 Theoretical Background

A description of how the method works has been provided; however, the theoretical underpinnings have yet to be discussed. Providing a solid theoretical background for the method gives it the consistency and strength needed to be accepted as a valid method. Additionally, the advantages and disadvantages of the method are better understood once its theory has been established. This section describes the PPM in terms of two well established theories, Bayesian inversion and Markov chains. Bayesian methods have been used extensively to solve inverse problems, and the first sub-section shows how PPM is equivalent to the Bayesian inversion methodology. This is followed by a short description of how PPM can be cast as a Markov chain.

Bayesian Inversion

Bayesian inversion is a theoretically well established inversion technique, and this section shows that the probability perturbation method fits into the framework of this method. Bayesian inversion recognizes that solutions to inverse problems are not unique and therefore seeks to sample realizations from the posterior distribution model. The PPM uses a sequential decomposition of the prior *and* posterior distributions. This allows solutions of PPM to be considered as realizations directly sampled from a posterior distributions, although this posterior is never explicitly stated. Before the details are explained, however, Bayesian inversion notation is defined.

All measurements from the reservoir constitute the data and will be termed **d**. This includes any "hard" data typically well log and core data and any secondary or "soft" data, usually seismic data. It also includes any production data such as fluid rate or pressure information from the reservoir. In the current text, this information is referred to as dynamic data. The reservoir will be described by a mathematical model with a set of parameters **m**. The aim of Bayesian inverse methods is to obtain samples of a posterior distribution of the model parameters given the data, $f(\mathbf{m}|\mathbf{d})$. The posterior can be redefined using Bayes relation.

$$f(\mathbf{m} \mid \mathbf{d}) = \frac{f(\mathbf{d} \mid \mathbf{m}) f(\mathbf{m})}{f(\mathbf{d})} \tag{2.7}$$

The prior density $f(\mathbf{m})$ describes the dependency among the model parameters and, therefore, constrains the set of possible solutions. In a spatial context such dependency refers to the spatial structure of $\mathbf{m}$, be it a covariance, Boolean or training image-defined dependency. The likelihood density $f(\mathbf{d}|\mathbf{m})$ models the stochastic relationship between the observed data $\mathbf{d}$ and each particular model $\mathbf{m}$ retained. This likelihood would account for model and measurement errors. Void of any such errors, the data and model parameters are related through a forward model g,

$$\mathbf{d} = g(\mathbf{m}) \tag{2.8}$$

The density $f(\mathbf{d})$ depends on the density $f(\mathbf{m})$ and the forward model g, but its specification can often be avoided in most samplers of $f(\mathbf{m}|\mathbf{d})$.

In this dissertation, we also treat the more general case where two types of dataset $\mathbf{d}_1$ and $\mathbf{d}_2$ are available. In such cases, the Bayesian decomposition of (2.7) is written as

$$f(\mathbf{m} \mid \mathbf{d}_1, \mathbf{d}_2) = \frac{f(\mathbf{d}_1, \mathbf{d}_2 \mid \mathbf{m}) f(\mathbf{m})}{f(\mathbf{d}_1, \mathbf{d}_2)} = \frac{f(\mathbf{d}_1 \mid \mathbf{m}) f(\mathbf{d}_2 \mid \mathbf{m}) f(\mathbf{m})}{f(\mathbf{d}_1, \mathbf{d}_2)} \tag{2.9}$$

where the second expression relies on a conditional independence hypothesis between the two data types. The data $\mathbf{d}_1$ is the static data whose relationship to the model $\mathbf{m}$ is linear or pseudo-linear in nature. The data $\mathbf{d}_2$ is the dynamic data that exhibits a complex multi-point and non-linear relationship with the model $\mathbf{m}$.

For the current discussion, we will again consider binary spatial variables. In this case, the model parameters are given by a set of indicator variables at each grid node location:

$$\mathbf{m} = \{ I(\mathbf{u}_1), I(\mathbf{u}_2), \ldots, I(\mathbf{u}_N) \}$$

where N is the total number of nodes (gridblocks) in the model, and the grid need not be rectangular. The prior distribution is then simply the joint distribution of all indicator random variables at all grid node locations

$$f(\mathbf{m}) = \text{Prob}\{\ I(\mathbf{u}_1) = i(\mathbf{u}_1),\ I(\mathbf{u}_2) = i(\mathbf{u}_2),\ \ldots,\ I(\mathbf{u}_N) = i(\mathbf{u}_N)\ \}$$

A fast and general method for sampling a prior distribution is sequential simulation. Sequential simulation is non-iterative and is completed after a single pass over all grid node locations. Sequential simulation relies on the following decomposition of the prior:

$$\begin{aligned} f(\mathbf{m}) = {} & \text{Prob}\{I(\mathbf{u}_1) = 1\} \times \\ & \text{Prob}\{I(\mathbf{u}_2) = 1 \mid i(\mathbf{u}_1)\} \times \ldots \times \\ & \text{Prob}\{I(\mathbf{u}_N) = 1 \mid i(\mathbf{u}_1), \ldots, i(\mathbf{u}_{N-1})\} \end{aligned} \tag{2.10}$$

This decomposition can be interpreted in two ways:

(1) Sampling a given joint distribution $f(\mathbf{m})$ is equivalent to sampling sequentially a series of univariate conditional distributions. If the joint distribution is known analytically, it determines the shape of each conditional distribution. The only joint distribution where each conditional is known is the multi-variate Gaussian. For a joint multi-Gaussian case, each conditional is also Gaussian with the mean and variance determined by a kriging system. Sequential Gaussian simulation (sgsim) is one of the most efficient algorithms for drawing samples from the multi-Gaussian distribution.

(2) Sequential drawing from a series of given univariate conditional distributions provides a realization (sample) of joint distribution which need not be explicitly stated. The only restriction on each such univariate distribution is that it needs to be conditional on all previously simulated parameters of the vector **m**. The aim in this approach is to generate realizations of **m** exhibiting certain (univariate, bivariate or multi-variate) statistics. The set of realizations generated in this manner represent a sample from some, not explicitly known, joint distribution.

The second approach recognizes that one is not interested in $f(\mathbf{m})$ itself but in the realizations generated from $f(\mathbf{m})$. In actual field cases, prior information on $\mathbf{m}$ comes in the form of limited statistics (e.g. a spatial covariance). The type of multi-variate density f cannot be determined from data and always needs to be assumed. The decision of distribution type is not made on the joint distribution, but on the conditional distributions. An example of such approach is direct sequential simulation, dssim (Journel, 1993), where the conditional distribution can be of any type, as long as they have mean and variance provided by a simple kriging system. The second interpretation also has led to an algorithm that samples from a large number of prior models, snesim (Strebelle, 2002) as discussed before in Section 2.1.

Priors need to be conditioned to data $\mathbf{d}_1$ and $\mathbf{d}_2$, resulting in a posterior distribution of the model parameters given the data, as formulated in Bayes rule, Eqns. (2.7) or (2.9). In determining the posterior, we will follow the same philosophy as in determining the prior: the posterior will not be explicitly stated, only conditionals will be modeled. We consider the case where the data $\mathbf{d}_1$ constitutes direct observations (hard data) of the model parameters at a set of $n << \mathrm{N}$ spatial locations. In the binary random function case,

$$\mathbf{d}_1 = \{i(\mathbf{u}_\alpha), \alpha = 1, \ldots, n\}$$

and $\mathbf{d}_2$ contains data that has a non-linear relationship with the model parameters (e.g. dynamic data)

$$\mathbf{d}_2 = g(\mathbf{m}) = \mathrm{g}(I(\mathbf{u}_1), I(\mathbf{u}_2), \ldots, I(\mathbf{u}_\mathrm{N}))$$

with g as the flow simulator. We will assume for now that there are no data or model errors. One is interested in generating samples from a posterior distribution:

$$f(\mathbf{m} \mid \mathbf{d}_1,\mathbf{d}_2) = \mathrm{Prob}\{ I(\mathbf{u}_1) = i(\mathbf{u}_1), I(\mathbf{u}_2) = i(\mathbf{u}_2) ,\ldots, I(\mathbf{u}_\mathrm{N}) = i(\mathbf{u}_\mathrm{N}) \mid i(\mathbf{u}_\alpha), \alpha = 1, \ldots, n ,\mathbf{d}_2\}$$

To sample such posterior, we do not follow the traditional Bayesian route of likelihood and prior. Instead, we rely on a sequential decomposition similar to the prior case:

$$
\begin{aligned}
f(\mathbf{m}\,|\,\mathbf{d}_1,\mathbf{d}_2) = \text{Prob}\{I(\mathbf{u}_1)=1\,|\,\{i(\mathbf{u}_\alpha),\alpha=1,\ldots,n\},\mathbf{d}_2\}\times \\
\text{Prob}\{I(\mathbf{u}_2)=1\,|\,i(\mathbf{u}_1),\{i(\mathbf{u}_\alpha),\alpha=1,\ldots,n\},\mathbf{d}_2\}\times\ldots\times \\
\text{Prob}\{I(\mathbf{u}_N)=1\,|\,i(\mathbf{u}_1),\ldots,i(\mathbf{u}_{N-1}),\{i(\mathbf{u}_\alpha),\alpha=1,\ldots,n\},\mathbf{d}_2\}
\end{aligned}
\tag{2.11}
$$

Sampling from the complex (joint) posterior f is equivalent to sequential sampling from a series of univariate conditionals of the type:

$$
\begin{aligned}
&\text{Prob}\{I(\mathbf{u}) = 1 \mid i(\mathbf{u}_1),\ldots,i(\mathbf{u}_{j\text{-}1}),\ \{i(\mathbf{u}_\alpha),\ \alpha = 1,\ldots,n\},\ \mathbf{d}_2\} = \text{Prob(A|B,D)} \\
&\quad \text{with A} = \{I(\mathbf{u}) = 1) \\
&\qquad \text{B} = \{\ i(\mathbf{u}_1),\ldots,i(\mathbf{u}_{j\text{-}1}),\ \{i(\mathbf{u}_\alpha),\ \alpha = 1,\ldots,n\}\} \\
&\qquad \text{D} = \mathbf{d}_2
\end{aligned}
$$

where we have re-introduced our original notation. By now we know that the conditionals, Prob(A|B,D), are obtain explicitly using Journel's method (2002) and the two conditional distributions, Prob(A|B) and Prob(A|D).

Working with conditional distributions will lead to an approach that is different from classical Bayesian inversion that involve the likelihoods Prob(B|A) and Prob(D|A); however, the end result, drawing from a posterior, is the same. Some practical limitations of traditional Bayesian inversion techniques are discussed, and some distinctions between the two methods are highlighted.

(1) Bayesian inversion methods involve iterative samplers that may require many thousands of evaluations of the forward model g. In history matching cases, a single evaluation of the forward model (flow simulation) may require several hours of CPU time; therefore, classical Bayesian inversion methods can not be applied to most real cases.

(2) The multi-Gaussian distribution is the only case for which the decomposition (Eq. 2.7) is analytically known (likelihood, posterior and priors are all multi-Gaussian if g is a linear function or can be sufficiently linearized). For this mathematical convenience, the multi-Gaussian is popular in Bayesian inverse work (e.g. Tarantola, 1987; Mosegaard and Tarantola, 1995; Hegstad and Omre, 1996).

a) The Gaussian framework requires a covariance structure for the prior as well as the likelihood. The prior covariance is estimated from any direct measurements of the model parameters **m**; however, no measurements are available for the likelihoods, and their covariance structures often have to be assumed. The PPM can be used in a Gaussian context, but only the prior covariances are required because likelihoods are never used in PPM.

b) Homoscedasticity of the error distribution, namely independence of the error, with regard to the "signal" $g(\mathbf{m})$ in

$$\mathbf{d} = g(\mathbf{m}) + \text{error}$$

is assumed to make inference of the likelihood distribution feasible; however, there is no physical reason why errors need to be independent. The PPM makes no assumptions about the error distribution.

c) The Gaussian framework is too limited to capture the possibly large variety of prior distributions of **m** observed in reality, nor does it properly model the full dependency between **m** and **d** in the likelihoods. Most reservoirs exhibit a strong correlation in the form of non-rectilinear shapes and connectivity (e.g. channels, cross bedding, complex litho-facies distributions) that cannot be adequately described by making multi-Gaussian assumptions. The PPM is not limited by the multi-Gaussian assumptions; therefore, it is able to account for the complex shapes that exist in petroleum reservoirs.

While we have shown that both methods have the same theoretical foundation, the probability perturbation method can be employed in a broad range of applications. In particular, the PPM is a technique where multiple data are integrated seamlessly into self-consistent models. Conditional distributions, P(A|X), can be calculated for any type of data, X. Data that has a linear relationship with the model properties can be directly

included in the model, and data with a non-linear relation can be iteratively included using the PPM. Any redundancy among the various types of data can be accounted for using the tau model (Journel, 2002; Krishnan, 2004). The PPM does not address the issue of tau parameters; developing a method for finding suitable tau values is still an area of considerable research.

Markov chain

Markov chains are used extensively in Bayesian inversion. Markov chains are a sequence of random values whose probabilities at a point only depend upon the value at the previous point and not on any points before the previous one. Bayesian inversion techniques use Markov chains to sample the posterior distributions because it allows the entire space to be sampled, yet the next point is always relatively close to the previous point. Coupling Markov chains with Monte Carlo sampling methods allows the high probability parts of the distribution to be sampled more frequently. The probability perturbation method can be described in terms of Markov chains as well.

One can define a non-stationary Markov chain on the entire set of random variables $I(\mathbf{u})$ $\forall$ $\mathbf{u}$, that starts from an initial model, $i^{(0)}(\mathbf{u})$, and generates iterations, $i^{(l)}(\mathbf{u})$, until convergence. To define such a Markov chain, consider the single random variable $I(\mathbf{u})$ at a specific location $\mathbf{u}$. Since $I(\mathbf{u})$ is binary, four transition probabilities (Eqns. 2.12-2.15) are defined for moving the chain from state (l) to state $(l+1)$ at location $\mathbf{u}$. In the PPM, these transition probabilities are parameterized by a single parameter r_D, where $r_D \in [0,1]$ and depends on the data D.

$$\text{Prob}\{I^{(l+1)}(\mathbf{u}) = 1|\text{D}; i^{(l)}(\mathbf{u}) = 0\} = r_D\text{Prob}\{I(\mathbf{u}) = 1\} \quad (2.12)$$

$$\text{Prob}\{I^{(l+1)}(\mathbf{u}) = 0|\text{D}; i^{(l)}(\mathbf{u}) = 1\} = r_D\text{Prob}\{I(\mathbf{u}) = 0\} \quad (2.13)$$

$$\text{Prob}\{I^{(l+1)}(\mathbf{u}) = 1|\text{D}; i^{(l)}(\mathbf{u}) = 1\} = 1 - r_D\text{Prob}\{I(\mathbf{u}) = 0\} \quad (2.14)$$

$$\text{Prob}\{I^{(l+1)}(\mathbf{u}) = 0|\text{D}; i^{(l)}(\mathbf{u}) = 0\} = 1 - r_D\text{Prob}\{I(\mathbf{u}) = 1\} \quad (2.15)$$

The first two transition probabilities (2.12) and (2.13) are the probabilities of changing states (facies) from step (l) to step (l + 1). The relative probability of the state changing is given by r_D, relative to the prior P(A), where r_D is a taken as a function of the conditioning data D. The degree of freedom from r_D allows the model to more closely match the data, D. At each iteration, a one dimensional optimization is carried out to find the value r_{Dopt} that matches best the data D. For any given $r_D \in [0,1]$ at any given current iteration (l) and for all grid cells **u** the conditional probability P(A|D) is obtained using the above transition matrix as

$$P(A|D) = r_D P(A) \text{ if } i^{(l)}(\mathbf{u}) = 0 \tag{2.16}$$

$$1 - P(A|D) = r_D(1 - P(A)) \text{ if } i^{(l)}(\mathbf{u}) = 1 \tag{2.17}$$

r_D is the same for all grid cells **u**. The resulting probability P(A|D) is then combined with the training image-derived probability P(A|B) using Eq. (2.2). This Markov chain is termed “non-stationary” since r_D changes at each iteration (l).

2.4 Discussion of the PPM

The PPM is able to overcome most of the limitations of the history matching methods discussed in Chapter 1. The main advantage of the PPM is that large scale parameters can be perturbed while the geology of the method is honored. None of the artifacts seen in Figure 1.1 occur with this method. Additionally, the method is robust. Although the previous examples only considered facies models, there is no limitation of the type of geology that can be honored (the only requirement is a conditional distribution). Finally, the method is relatively efficient. While 400 flow simulations is a large amount, it is much less than the thousands that are required by traditional Bayesian inversion techniques.

The last limitation stated in Chapter 1 ("…practical for real field applications …") has not yet been addressed. Real examples are large and geologically complex, and production data response may be quite different for different areas of the reservoir. When history matching these reservoir models, some areas of the model likely will need to be perturbed differently than another areas. However with the current PPM, the change is "in probability" the same everywhere. The PPM parameterizes the entire model with only one parameter, r_D. This reduces the dimensionality of the problem but does not allow for a difference in the magnitude of perturbation over different areas of the reservoir model. Even the relatively simple 2D example required 400 flow simulations to match only two different water cut curves. For real applications, 400 flow simulations will be much too costly.

Often in research, the initial idea alone is not enough, and a second step is needed to take the idea (with all its advantages) and make it work in real practical settings. The next chapter explains how the PPM is extended to be practical for realistic applications.

"Theoretical principles must sometimes give way for the sake of practical advantages."

- William Pitt

Chapter 3

Multi-Parameter PPM

The probability perturbation method is extended to allow for more realistic reservoir models to be matched. Rather than having one perturbation parameter for the entire reservoir, a multi-parameter method is proposed. This allows different amounts of perturbation to be applied to different areas of the reservoir model. A method is presented to optimize multiple parameters without losing the simplicity of the one parameter method. The new multi-parameter PPM is supported with a number of examples. The 2D example from Chapter 2 is revisited and used to study the new method in great detail. The multi-parameter method is compared with the single-parameter method as well as other techniques. Several approaches to optimize the multiple parameters are also evaluated with this example. Additional examples show the robustness of the method; a variety of parameters are perturbed and both pressures and rates are matched. The chapter ends with a three facies 3D example that shows a more practical implementation of the new multi-parameter method.

3.1 Defining Multiple Parameters

Petroleum reservoirs have many locations (wells) where production data is recorded, and the response is uaually quite dissimilar in different areas of the reservoir. Additionally, the numerical models that represent the reservoirs are often large and complex, and the production can be extremely varied throughout the model due to a spatial variation in the geological complexity and the well configuration. The highly variable nature of production response from the reservoir and from the numerical model makes history matching difficult. Typically, production data from one area may be matched while production data from another area is not. In the PPM of Chapter 2, the perturbation parameter, r_D, quantifies the same amount of desired change over the *entire* reservoir, so the whole model is perturbed, at least in probability, by a similar amount. The match of wells in one area might be destroyed as the method tries to match the data from other areas. This will dramatically reduce the speed of convergence for the method, or it may not be able to converge at all. To overcome this problem, the various areas will need to be perturbed by different amounts.

To accomplish this, first, the reference (field) production data is split into K distinct units, $D = \{D_1, D_2 \ldots D_K\}$. A "unit" is a set of production data that can grouped together because all of that data is influenced by properties in the same area. A logical choice would be to have one unit per well; however, nothing mandates that approach. If multiple wells are in a geologically distinct region, production from all of the wells may be lumped into one unit. Also, if there are two separate zones producing from one well and the production split was known, two units could be assigned to one well. However, these situations would be less likely and usually, there would be simply one unit per well. Similarly, the simulated production is divided into corresponding units, $D^S = \{D_1^S, D_2^S \ldots D_K^S\}$. A single perturbation parameter is assigned to each unit, and they are denoted $\{r_{D1}, r_{D2} \ldots r_{DK}\}$ where K is the total number of distinct units to be matched.

The r_{Dk} values are coupled to the reference data and simulated production through the objective function. The objective function previously was defined for $O(r_D)$ as stated by Eq. (2.4). It can be restated as $O(r_{D1}, r_{D2} \ldots r_{DK})$ where it is now a function of multiple parameters, and it can be further decomposed into

$$O(r_{D1}, r_{D2} \ldots r_{DK}) = \sum_{k=1}^{K} O_k(r_{Dk}) \qquad O_k(r_{Dk}) = \| D_k^S - D_k \| \tag{3.1}$$

where each unit has a separate objective, $O_k(r_{Dk})$. The objective is simply some measure of mismatch between the simulation and reference production data, but now it is calculated for each specific unit k. In this work, the mismatch is calculated by the least square error, but any measure can be used. The idea is to perturb the properties that influence production for one unit differently than properties that affect production for another unit by optimizing r_{Dk} for each unit.

Now instead of one parameter to optimize there are multiple parameters, and an important aspect of this work is how to optimize these multiple parameters. The next section provides ample details on how to do this in an efficient and practical manner. However, a discussion of how models are built with multiple perturbation parameters, $\{r_{D1}, r_{D2} \ldots r_{DK}\}$, needs to be completed first.

As previously stated, there are K separate r_{Dk} values, and they are attached to the "units" of production data. In the single parameter method, the same r_D value was attached to all gridblocks and their corresponding property values. A method to assign the multiple r_{Dk} values to all gridblocks is presented. Each gridblock can be assigned only one r_{Dk} value, and the following options are proposed.

(1) r_{Dk} values are constant for all gridblocks within a given region. This requires defining the region geometries.

(2) r_{Dk} values are interpolated from the well locations (where the production data, D_k, is measured) to all gridblocks.

When using regions, one must specify the geometry of the regions. The regions could be defined from geologic information, from streamlines, from sensitivity coefficients, from Euclidian distances (Voronoi polygons), or simply from the engineer's judgment. If the regions overlap, a technique to average the parameters in the overlapping areas will have to be determined. There is a significant amount of work available on defining regions for history matching (Milliken *et al.*, 2001; Bissell, 1994), and a number of different techniques are investigated and discussed throughout the dissertation.

The r_{Dk} values are calculated based on production data, and the production is measured at wells, so the r_{Dk} values are assigned to the gridblocks where the production well is located. When this is done, a method must be specified to interpolate the values. Splines, inverse distance weighting, kriging, or other methods can be used. For a couple of examples in this work, kriging is used. For the different types of interpolations, various parameters must be defined; for example, in kriging, the variogram and search radius need to be specified.

Recall from Chapter 2 that a soft probability, P(A|D), is utilized to integrate production data information into the model. The multi-parameter perturbation technique relies on the same idea but differs in how P(A|D) is calculated. A perturbation parameter is assigned to every grid location, $\mathbf{u}_n$ where n = 1, 2 … N, and N is the total number of gridblocks in the reservoir model.

When *regions* are used, the following P(A|D) is proposed.

$$\forall\, \mathbf{u}_n \in R\,,\ P(A|D) = (1-r_{Dk})\, i^{(0)}(\mathbf{u}_n) + r_{Dk}P(A), \text{ with } k : \mathbf{u}_n \in R_k \qquad (3.2)$$

where A is for example, "channel occurs," or in geostatistical terminology: $\{i^{(0)}(\mathbf{u}) = 1\}$. P(A|D) is defined for the entire reservoir, R, but its local value depends on the region definition: when $\mathbf{u}_n$ is located in region R_k, the perturbations parameter takes on a value

of r_{Dk}. Therefore P(A|D) will have different values for different regions of the reservoir model. Similar to the single parameter method, P(A|D) is combined with P(A|B) using Journel's method (Eq. 2.2) to obtain a new probability model, P(A|B,D). This probability model is then used to create a new single realization, $i^{(r_{Dk})}(\mathbf{u})$.

When *interpolation* is used, some function to interpolate the r_{Dk} values must be specified. As discussed previously, this function can be of many different types, but it will be represented as $f(r_{D1} \ldots r_{DK})$. There for the definition of P(A|D) becomes

$$\forall\ \mathbf{u}_n \in R\ ,\ P(A|D) = (1 - f(r_{D1} \ldots r_{DK}))\ i^{(0)}(\mathbf{u}_n) + f(r_{D1} \ldots r_{DK})P(A) \qquad (3.3)$$

The value of $f(r_{D1} \ldots r_{DK})$ may be different at each gridblock. P(A|D) will not be constant over a region, but will smoothly vary throughout the reservoir. Once P(A|D) is calculated, the subsequent aspects of the method remain the same.

A visual illustration of the multi-parameter method is provided in the following sub-section. The illustration compares the region based method with the interpolation based approach. The reason that artifacts and discontiuities are not created when regions are used is also explained. Finally, the similarities and differences between the single parameter method and the multiple parameter method are described.

Illustration of the MP-PPM

A conceptual example is provided to explain the details of the MP-PPM. Some of the information used in the example is provided in Figure 3.1. The model includes elliptical bodies that extend from top-right to bottom-left (45º) as displayed in the training image. There is one injection well and two production wells, and the production data is divided into two units. The "true" production data is displayed as dots for the two production wells in the plots labeled Well 1 and Well 2 (Figure 3.1). An initial realization has been created, and production from the initial realization is also included in the plots (as lines). The right-most well (Well 2) matches exactly, but the well on the left (Well 1) is not matching. Thus, the properties around Well 1 require more change (and a higher r_D) than

the ones around Well 2; consequently, the value for r_{D1} equals 0.80 while r_{D2} is 0.05. These r_D values are simply assumed for demonstration purposes (i.e. they were not calculated) since this is only a conceptual example.

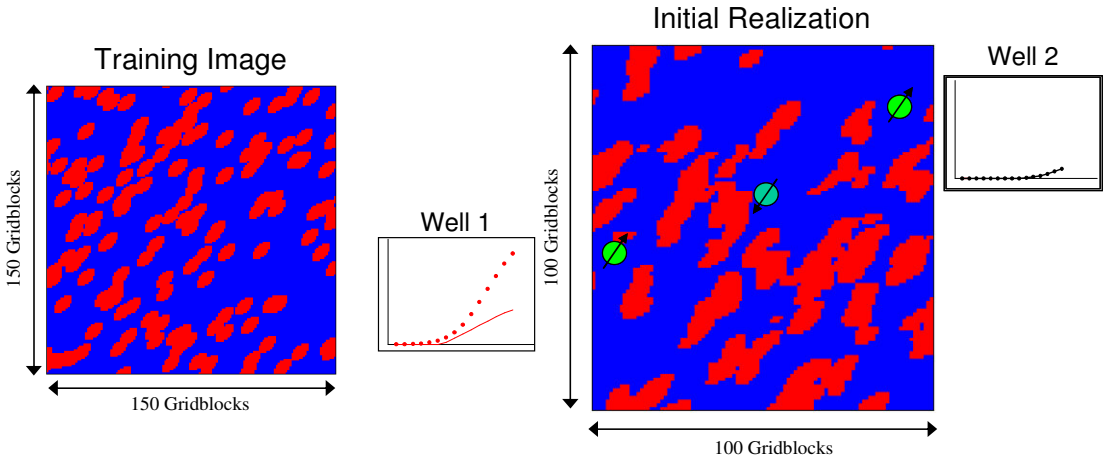

Figure 3.1: Training image, initial realization and well data for illustrative example.

These two r_{Dk} values need to be assigned somehow to all the gridblocks. With the regional method, all gridblocks in a region receive the same r_{Dk} value, and only the geometry of the regions needs to be defined. Three different ways to define regions are presented in Figure 3.2. If Voronoi polygons (closest gridblocks) are used to define regions, Figure 3.2A will be the result. However, the principal direction of the elliptical bodies is going in the other direction, and the regions may need to be aligned with the bodies (as defined geologically). In this case, Figure 3.2B would be the region geometry. Another method would be to simply split the model down the middle as in Figure 3.2C. This appears as a compromise between the two previous region definitions, and the regions in Figure 3.2C will be used through the rest of the example.

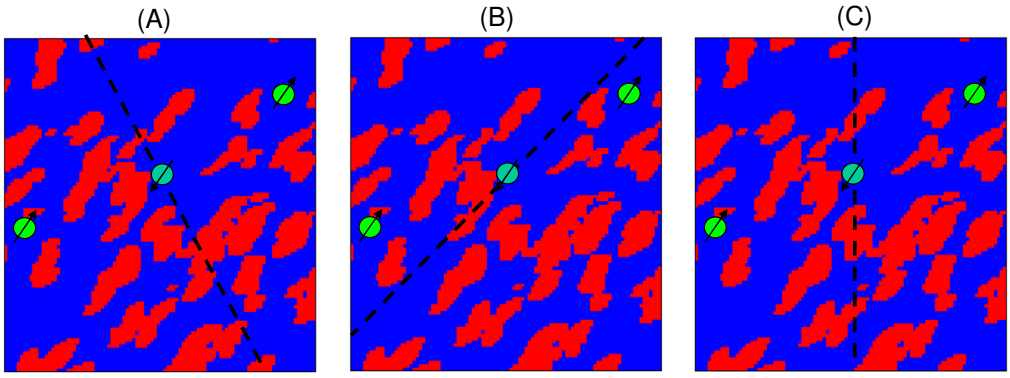

Figure 3.2: Three different methods to define regions.

A region does not have to perfectly reflect the drainage zone of a well included in that region. As the model is perturbed, the flow paths change and the gridblocks that affect flow likely change. Production is primarily influenced by the set of gridblocks that are close to the well. As long as these principal gridblocks are included in the region, a history matched model can be achieved. For example, all three region definitions in Figure 3.2 would likely lead to a history match.

In fact, a single gridblock may influence production at more than one well, and this idea is recognized when using interpolation to create regions. Although other interpolators can be used, kriging is used in this work and will be discussed. Kriging requires a variogram. The variogram defines how the data are correlated to each other; the variogram range gives the extent of the correlation, and if anisotropy exists, the direction of preferential correlation can also be provided. For interpolation of r_D values, any arbitrary variogram could be used as long as it is smooth (e.g. Gaussian), it fits the main direction of flow and it has a small nugget effect. The variogram ranges are taken to be approximately equal to the production well spacing. The anisotropy ratio and angle are taken to follow the values used for the petrophysical properties in the reservoir model.

For the previous r_D values ($r_{D1} = 0.80$ and $r_{D2} = 0.05$), Figure 3.3 shows r_D maps for three different variograms ranges. The first one (A) has a relatively short range (20

gridblocks). The r_D value for most gridblocks equals the average of the two r_D values (0.425). The second image (B) has a larger range (100 gridblocks), and most blocks have a value that is different from the mean. The third one (C) has the same range as (B), but it has increased continuity in the 30º direction.

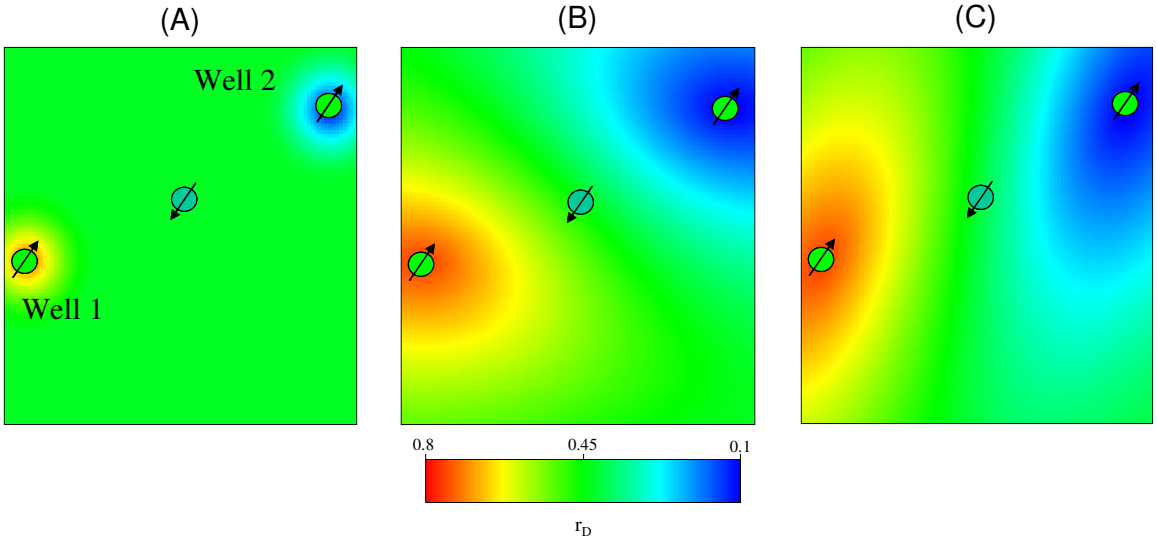

Figure 3.3: Three different variograms to interpolate r_D values.

The r_D values from the regions in Figure 3.2C and from kriging in Figure 3.3C are used in Eq. (3.) to generate soft probability maps P(A|D). Then a new realization is generated for each method (Figure 3.4).

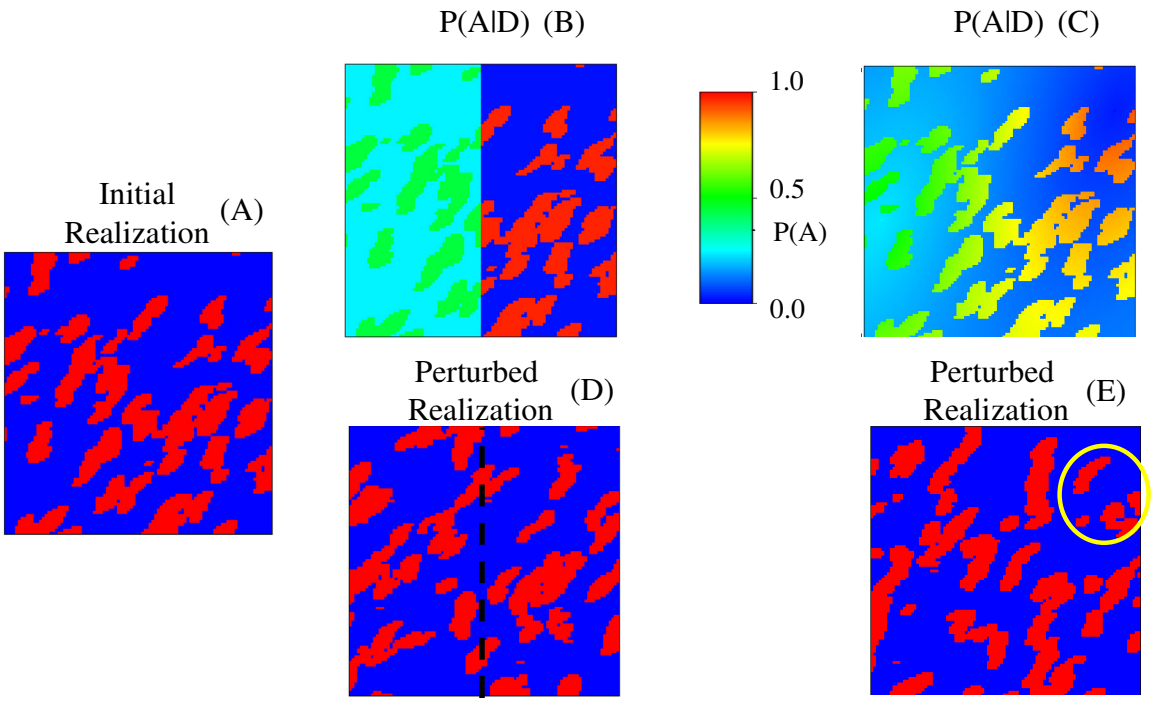

Figure 3.4: (A) Initial realization (B) P(A|D) map for regions, (C) P(A|D) map for kriging, (D) new realization for regions, and (E) new realization for kriging.

There are areas where high confidence exists in the current model (e.g. areas where production data matches closely), and there are areas where there is less confidence in the current model (e.g. poor match of production data). Where there is little confidence (r_{Dk} is high), P(A|D) is close to the marginal (0.30) for both facies, and significant change is expected. While for the high confidence region, the facies model is very close to the initial facies model, and only minor changes occur. For kriging this only happens very close to the well locations. There is very little change close to Well 2. In the circled region in Figure 3.4E, the bodies are in the same place and only the shapes have somewhat changed. The amount of change for the regional method is more dramatic; the right has little change, but the left is changed significantly. When r_{Dk} is close to zero (well matched region), the model does not change very much, such as is Region 2; conversely, when r_{Dk} is close to one (poorly matched region), the model changes

significantly as shown in Region 1. Although the perturbations in the two regions are quite dissimilar, discontinuities along the border of the regions are not observed.

The reason for not creating artifacts is explained by the nature of the sequential simulation algorithm and by the perturbation method applied. The bodies always maintain their geological continuity, and the hard data is always honored because the P(A|B) is the same for all regions. Discontinuities are not created across region boundaries due to the nature of sequential simulation. In sequential simulation, each grid block is simulated based on any reservoir data and on any previously simulated grid block properties. The method searches for any such previously simulated grid location in an elliptical search neighborhood. This search neighborhood may (and should) cross the region boundaries. When simulating a gridblock in one region, the gridblock properties in any other regions are used to determine P(A|B,D); hence, continuity is created across the boundaries. It should be emphasized that at no instance does the method construct realizations per regions or run the flow simulator per region.

For further emphasis, the E-type estimate (average) of 100 realizations of Figure 3.4C is calculated and displayed in Figure 3.5. Since there is high confidence in Region 2 (r_{D2}=0.05), the facies bodies in this region change only slightly from their original positions causing the facies body positions in the initial realization to be distinguishable. Conversely, in Region 1 the large perturbation of r_{D1}=0.8 is reflected by a large spatial variability; the placement of facies bodies of the initial realization is hardly discernable.

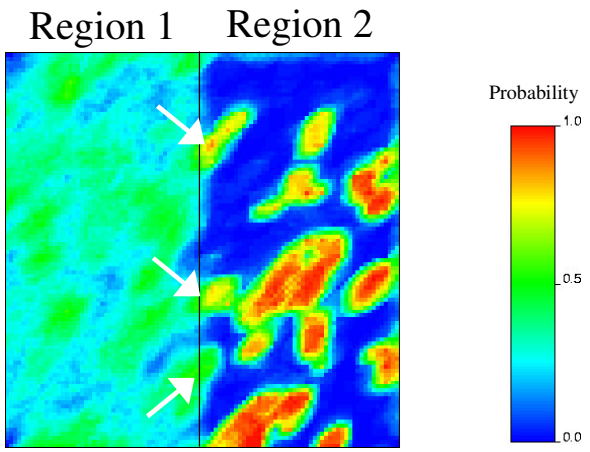

Figure 3.5: E-type estimate (average) of 100 realizations of Figure 3.4C. Arrows point to three locations where bodies cross the border in the initial realization.

There are three places in the initial realization (Figure 3.2C) where the elliptical bodies cross the region boundary, as indicated by the arrows in Figure 3.5. The probability that these bodies will occur in any perturbed realization is less than the other bodies in Region 2. This is especially true for the lowest body that crosses the border (lowest arrow). If the body is present in the perturbed realization in Region 1, then it will extend across into Region 2. If it is not present in Region 1 then it will not be in Region 2, thus enforcing the geologic continuity of the realization.

Figure 3.6 illustrates the similarities and differences between the single parameter perturbation method and the multiple parameter perturbation method. Regions are used for the multiple parameter method. Both methods use the same initial guess as shown in Figure 3.6A and Figure 3.6B. The next images (Figure 3.6C and Figure 3.6D) are the P(A|D) maps for the two methods. Figure 3.6C is created using Eq. (2.3) with r_D=0.50 while Figure 3.6D is created using Eq. (3.) with different r_{Dk} values (r_{D1}=0.80 and r_{D2}=0.05). Then the last step is the same for both methods; P(A|D) is combined with the P(A|B) and used to create one single realization (Figure 3.6E and Figure 3.6F).

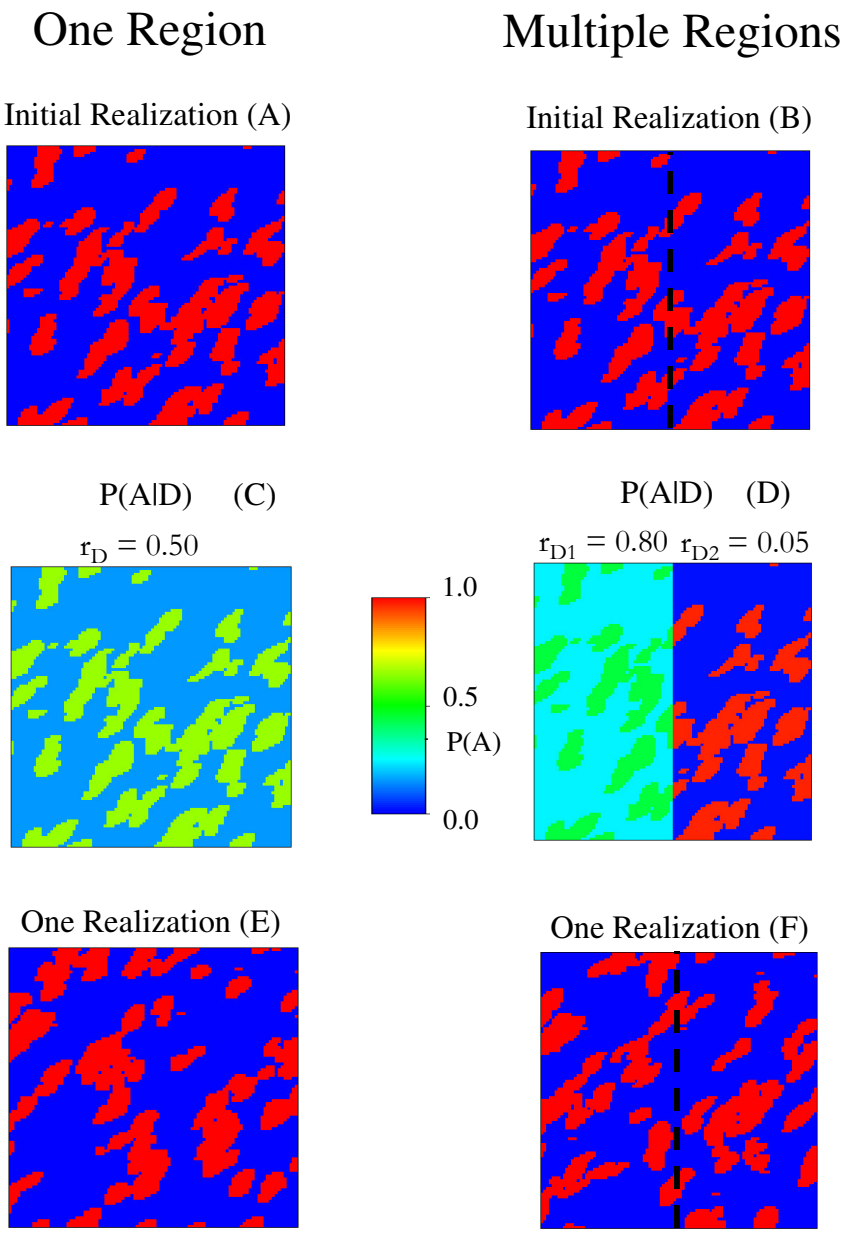

Figure 3.6: Comparison of the global method and the regional method.

The realization in Figure 3.6E shows a change that is similar everywhere in the reservoir, and as the bodies on the right are perturbed, the match in Well 2 is lost. However Figure 3.6F clearly shows more change in Region 1 than in Region 2. In Region 1 the bodies are

in completely different locations, but in Region 2, the bodies are all essentially in the same place other than the shapes have changed a little. Therefore, the match in Well 2 is maintained while the match in Well 1 is improved.

3.2 Optimizing Multiple Parameters

Now that the method to create a perturbed realization has been covered, the procedure to find the optimum values for r_{Dk} is discussed. Finding optimal perturbation parameters for all regions creates a multi-dimensional optimization problem that is more difficult to solve than the method with one parameter. Both methods have the same initial steps; make an initial guess and run the simulator for the entire model, but the objective function is calculated somewhat different for the two methods. For the single parameter method, the objective is directly calculated from all the data (Eq. 2.4), but for the multi-parameter method, the objective is calculated first for each unit separately. This allows multiple r_{Dk} values to be modified. The stopping criteria or convergence is based on the overall mismatch calculated by summing the multiple individual mismatches (Eq. 3.1).

The simulated production at each well may be influenced by gridblocks anywhere in the reservoir model; therefore, the values of $O_k(r_{Dk})$, in general, depend on all $\{r_{D1} \ldots r_{DK}\}$. However, the production from a well is principally dependent on the petrophysical properties that are close to that well; thus, we will assume, at least implicitly, that $O_k(r_{Dk})$ only depends on r_{Dk}. Based on this assumption, *all* K r_{Dk} values are updated based on the production data from *one* single flow simulation. The r_{Dk} for each unit is determined by performing one step of a one-dimensional optimization routine, but since there are multiple parameters, this "one step" is done K times (once for every unit). After all r_{Dk} values are updated, a new reservoir model is generated and flow simulation is performed.

In order to jointly find all r_{Dk} values, the Brent algorithm is modified slightly from its original form. For the multi-parameter (MP) method, K objective function values (mismatches) are needed. We start with K initial r_{Dk} values (one for each unit) and use regions or interpolation to determine the r_{Dk} values for all gridblocks in the model. A new realization is generated from the r_{Dk} values, and flow simulation is run on that new realization. The mismatches, $O_k(r_{Dk})$, are then calculated for each distinct region. This is where Brent is modified. In the new routine (MP-Brent), the part that updates the

function points (r_{Dk} values) based on mismatch is repeated K times. Each time a different mismatch, $O_k(r_{Dk})$, is used to determine a different r_{Dk} value. Recall that one does not have to run a flow simulation for each region because the mismatches are calculated from the full flow simulation. The updated set of r_{Dk} values is calculated using bisection (Figure 3.7). With the new r_{Dks}, a new realization is generated; flow simulation is performed, and a new set of objective function values are calculated. Now the two previous $O_k(r_{Dk})$ are used to determine the next r_{Dk} values for each unit. Once there are three sets of objectives and three sets of r_{Dk} values, either parabolic interpolation or another bisection may be used to find any of the next r_{Dk} values (see Section 2.2).

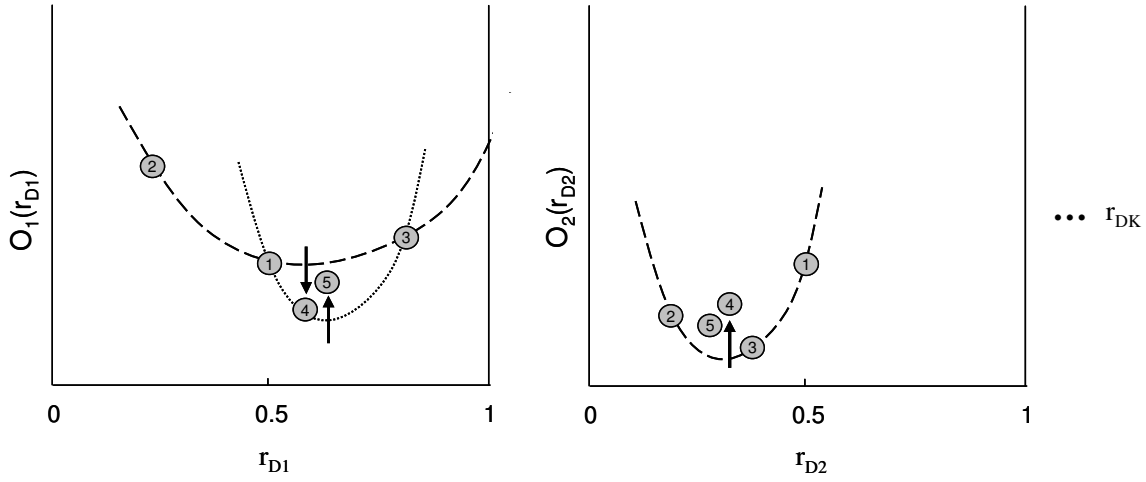

Figure 3.7: Perturbations parameters and objectives for two of possibly K different production units. The circled numbers indicate the step number of MP-Brent. The dashed lines are the fitted parabolas and the arrows point to the objective function value found at the minimum of the parabolas.

Figure 3.7 shows how the optimization might occur for two of possibly K different units of production. For r_{D1}, the objective value for point 2 is higher than point 1, so the MP-Brent algorithm will search the other half of the space and will place the third point there. A parabola is fit through the three points, and at the minimum another objective (4) is calculated. A new parabola is fit through points 1, 4 and 3 to get the next point (5). For r_{D2}, the objective values of the second point is lower than the first, so the bisection is

completed between points 1 and 2, and a parabola can be constructed through points 1, 2 and 3 to get the fourth point. Since the objective values for point 4 is greater than 2 or 3, the parabola through these three lowest points would have a maximum, not a minimum. Therefore, bisection between the two lowest values is used to get the next point (5).

Although multiple parameters are being updated, the example makes it clear that this is not a traditional multi-dimensional optimization method. The r_{Dk} values are updated without any explicit reference to each other; indeed, some may be updated by bisection while others are updated with parabolic interpolation. However, dependence is implicitly maintained by creating the models and running flow simulation over the entire model.

Since convergence is based on minimizing the overall objective, the proposed multi-parameter optimization will determine the set of parameter values, $\{r_{D1} \ldots r_{DK}\}$, that results in the best production match for the total reservoir. Therefore, each individual r_{Dk} may not be at its own minimum. For example, if there were only the two objectives shown in Figure 3.7, then the best match would be the fourth point, even though the third point is the lowest for r_{D2}.

Similar to the single parameter method, the multiple parameter technique consists of an inner loop and an outer loop. The inner loop finds the optimum realization, $i^{(r_{Dkopt})}(\mathbf{u})$, between the initial realization, $i^{(0)}(\mathbf{u})$, and an equiprobable realization, $i^{(1)}(\mathbf{u})$. The outer loop consists of replacing the initial realization with the previous optimum realization and changing the random seed. Once the overall match is below some tolerance, the method is stopped and a history matched realization is obtained. Figure 3.8 displays a flow chart of the multi-parameter method to incorporate production data into a geologic model in a consistent manner. The Fortran computer code is provided in Appendix A.

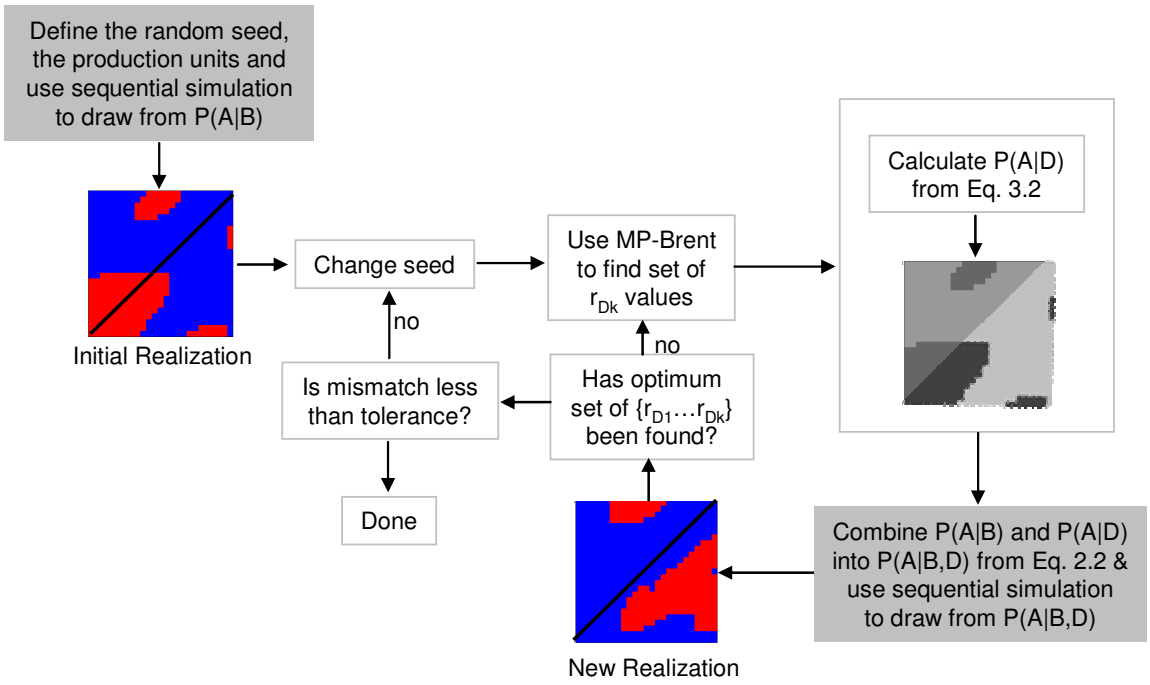

Figure 3.8: Multi-parameter probability perturbation flowchart.

It is important to note that the model is never broken into K independent problems where the flow simulation or model generation would be completed on each one separately. Rather, the realizations are always generated over the entire reservoir, and the flow simulation is always completed over the entire reservoir. The objective function calculations (mismatches) and perturbation parameter (r_{Dk}) updating are done separately. The most time consuming (CPU consuming) aspect of the method is the flow simulation. For the multi-parameter method, this is completed only once per iteration (the same as the single parameter method). Therefore, much faster convergence is possible because the MP-PPM allows different areas to be perturbed by different amounts (e.g. matching areas can be perturbed by a small amount while non-matching regions are perturbed by a large amount) for the same number of flow simulations.

Algorithm

A summary of the multi-parameter probability perturbation method is provided:

- Define a training image depicting the desired geological continuity
- Assign observed production data to respective units.
- Generate an initial realization, $i^{(0)}(\mathbf{u})$ for $\mathbf{u} \in R$
- Change random seed
- Outer loop: until history match, $O(r_{D1}, r_{D2} \ldots r_{DK}) < \varepsilon$:
  - Inner loop: perform optimization to get the r_{Dks} that result in the best match to the production data, D_k, for k = 1, 2 … K. The objective function is evaluated as follows:
    - Use regions or interpolation to go from r_{Dk} at wells to r_{Dn} at all gridblocks.
    - Calculate $P(A|D) = (1-r_{Dn})i^{(0)}(\mathbf{u}_n) - r_{Dn}P(A)$ and combine it with P(A|B) to get the probability model P(A|B,D).
    - Generate new model, $i^{(r_{Dk})}(\mathbf{u})$ and run flow simulation to get production data, D_k^S, for k = 1, …, K.
    - Calculate regional objectives, $O_k(r_{Dk}) = \| D_k^S - D_k \|$, for k = 1, …, K and overall objective, $O(r_{D1}, r_{D2} \ldots r_{DK}) = \sum_{k=1}^{K} O_k(r_{Dk})$.
  - If $O(r_{D1}, r_{D2} \ldots r_{DK}) < \varepsilon$–end; else–change random seed and set $i^{(0)}(\mathbf{u}) = i^{(r_{Dkopt})}(\mathbf{u})$

3.3 Simple 2D Example

The 2D flow model from Chapter 2 is used to better understand the multiple parameter probability perturbation method. Due to the simplicity of the model, multiple parameters would not be required to history match this reservoir model; however, that simplicity allows numerous aspects of the model to be studied. For example, the multi-parameter (MP) method is compared in terms of efficiency with the single parameter (SP) method and also with random search. The MP-Brent optimization is compared with a more traditional multi-dimensional optimization, the polytope (simplex) method. Other aspects of the MP-PPM are explored such as the initial r_{Dk} values and the convergence tolerance. The section concludes by looking at predictions from the history matched models.

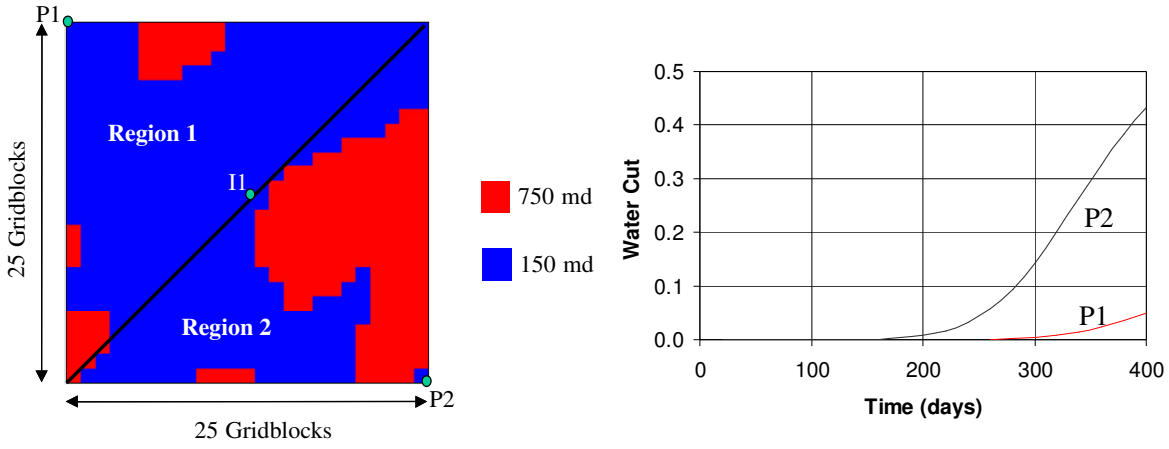

Figure 3.9: The reference model is shown divided into two regions and the reference production data is also displayed.

For the multi-parameter method, there are two units (2 r_{Dks}), one for Well P1 and one for Well P2. The grid will be filled using regions that are defined by Voronoi polygons (Figure 3.9). The top-left will be Region 1 and the bottom right will be Region 2. The reference model and production data are the same as used for the single parameter method (Figure 3.9); the training image is also the same (Figure 2.2).

MP-PPM and SP-PPM

At this point the multi-parameter method is used to history match the example. The same initial realization that was used for the single parameter method is used to start the process for this case. The multi-parameter Brent method that was discussed in the previous section is used to optimize r_{D1} and r_{D2}. Once again, the standardized least square error between the reference water cut and the simulated water cut is calculated. Once the overall error is below 0.001, the matching procedure is stopped. The required number of flow simulations per outer loop (~7 flow simulations) is the same as for the SP method. Figure 3.10 shows the matched realization and the matched water cut. The history match required 40 outer iterations and around 280 flow simulations to complete. This is considerably less than the SP-PPM (57 outer iterations and 400 flow simulations). Although 280 flow simulations are too many for real field applications, this example is simply used to compare the methods. It is a fast running case, so 280 simulations is not unreasonable. Further discussion on how to improve the efficiency of the method so that it can be extended to real cases is provided later in this section.

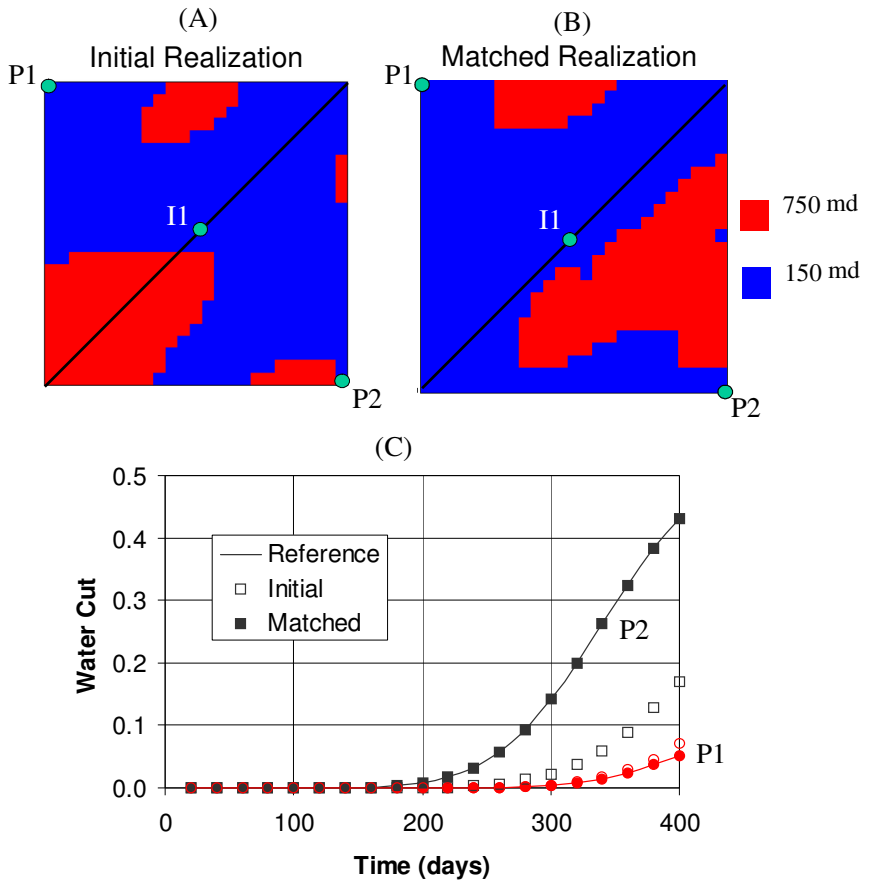

Figure 3.10: An initial realization and matched realization using MP-PPM and water cut data for the reference, and the initial and matched realizations.

Figure 3.11 gives the objective function for the outer iterations for both the MP-PPM and the SP-PPM and the r_{D1} and r_{D2} values for one inner iteration. Although the overall convergence is more rapid for MP-PPM than for SP-PPM, the shape of the objective function decrease is similar for both methods. Much of the convergence is obtained early in the process, and a large portion of the CPU time is spent to reduce the objective by the last small amount.

The inner iterations for r_{D1} and r_{D2} have the same high variability as the inner iteration for the SP-PPM method given in Figure 2.10. As noted in the previous chapter, the

objectives are not decreasing monotonically due to the non-continuous nature of the bodies, and the optimum realization from the inner loop is not always the last one. Also the best realization retained is based on the overall objective, so the lowest objective value for any single unit may not be obtained. For this example, the fourth iteration has the lowest overall objective and is saved, even though the lowest objective for r_{D2} is the seventh (Figure 3.11). Further, the behavior of the r_{Dk} values from this example provides interesting insight. Production from P1 is improving its match as r_{D1} is getting closer to zero (small changes), while P2 production improves as r_{D2} gets larger (large changes). This is because production from Well P1 is initially matching satisfactorily while Well P2 is not.

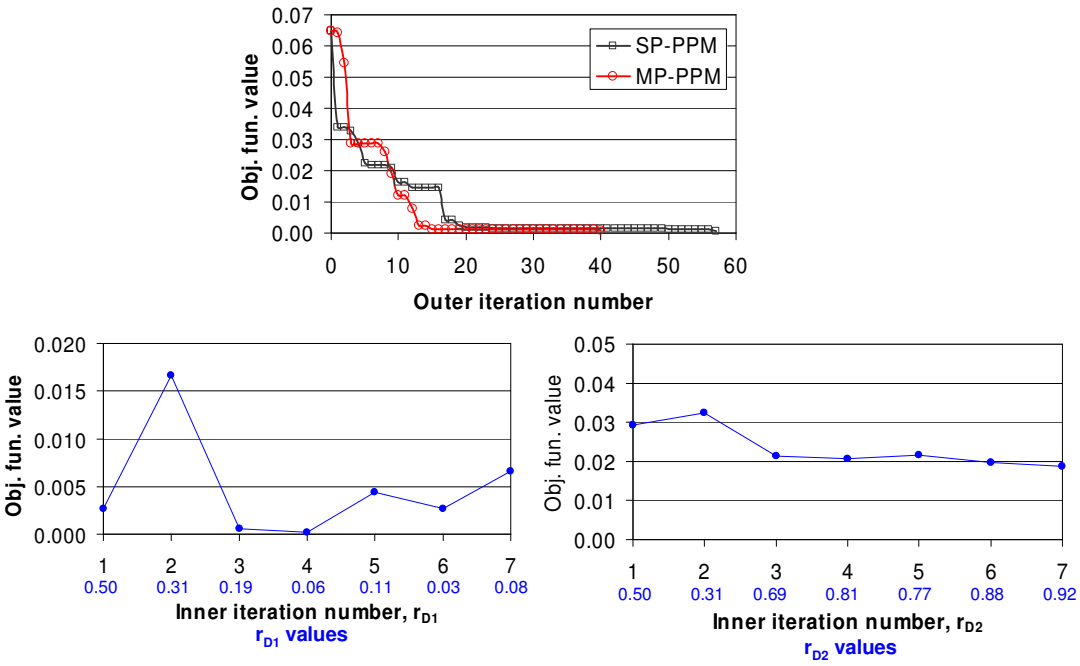

Figure 3.11: Outer iteration for both methods, and the inner iteration of two r_{Dk} values for one outer loop of MP-PPM.

To limit the influence of ergodic fluctuations in stochastic realizations, ten separate history matched realizations were generated by matching to the same reference data. The results are averaged to allow a more robust comparison of the two methods. The ten

procedures start with ten different initial realizations and use ten different sets of random numbers. Both MP and SP methods use the same initial realization and random seeds for each run to ensure the methods are compared fairly.

Table 3.1 shows the number of outer iterations for both methods for each realization. The MP-PPM is much faster for all realizations except two realizations (#3 and #8). For two other realizations, (#4 and #10) the SP-PPM is not able to converge in 100 outer iterations, the maximum allowed for this example.

Table 3.1: Comparing the number of outer iterations for SP-PPM and MP-PPM to converge for 10 realizations.

	Number of outer iterations to converge to a tolerance of 0.001										Average	Average
Realization	1	2	3	4	5	6	7	8	9	10	Iteration	Flow Sim.
MP-PPM	40	25	42	21	7	21	47	58	20	31	31.2	216.1
SP-PPM	57	44	34	>100	90	31	81	40	90	>100	66.7	456.9

On average the MP-PPM converges in about 31 iterations, more than twice as fast of the SP-PPM, which takes about 67 iterations. When the SP method was not able to converge in 100 iterations, a value of 101 was used for averaging. There are around seven flow simulations in each outer iteration, and the average number of flow simulations for the two methods is also listed in the table. For this example, the MP-PPM takes on average just over 200 flow simulations to match the production while SP-PPM takes about 450.

Figure 3.12 compares how the outer iteration objective function decreases for MP-PPM and SP-PPM for all ten realizations. Notice the long tail for the majority of the realizations, especially for the SP-PPM. For both methods, the majority of the reduction in objective function occurs in the first few iterations, but it takes many more iterations to converge below the final tolerance.

In the reference, there is significantly more high permeability facies between the injector and the P2 well than the P1 well, and this causes the water cuts for the two wells to be quite different. This is why the two regions need to be perturbed differently. If the production for one well is matching while the other is not, the matching regions can be

maintained as the other region is improved, and this is the reason the MP-PPM is faster than the SP-PPM for this example.

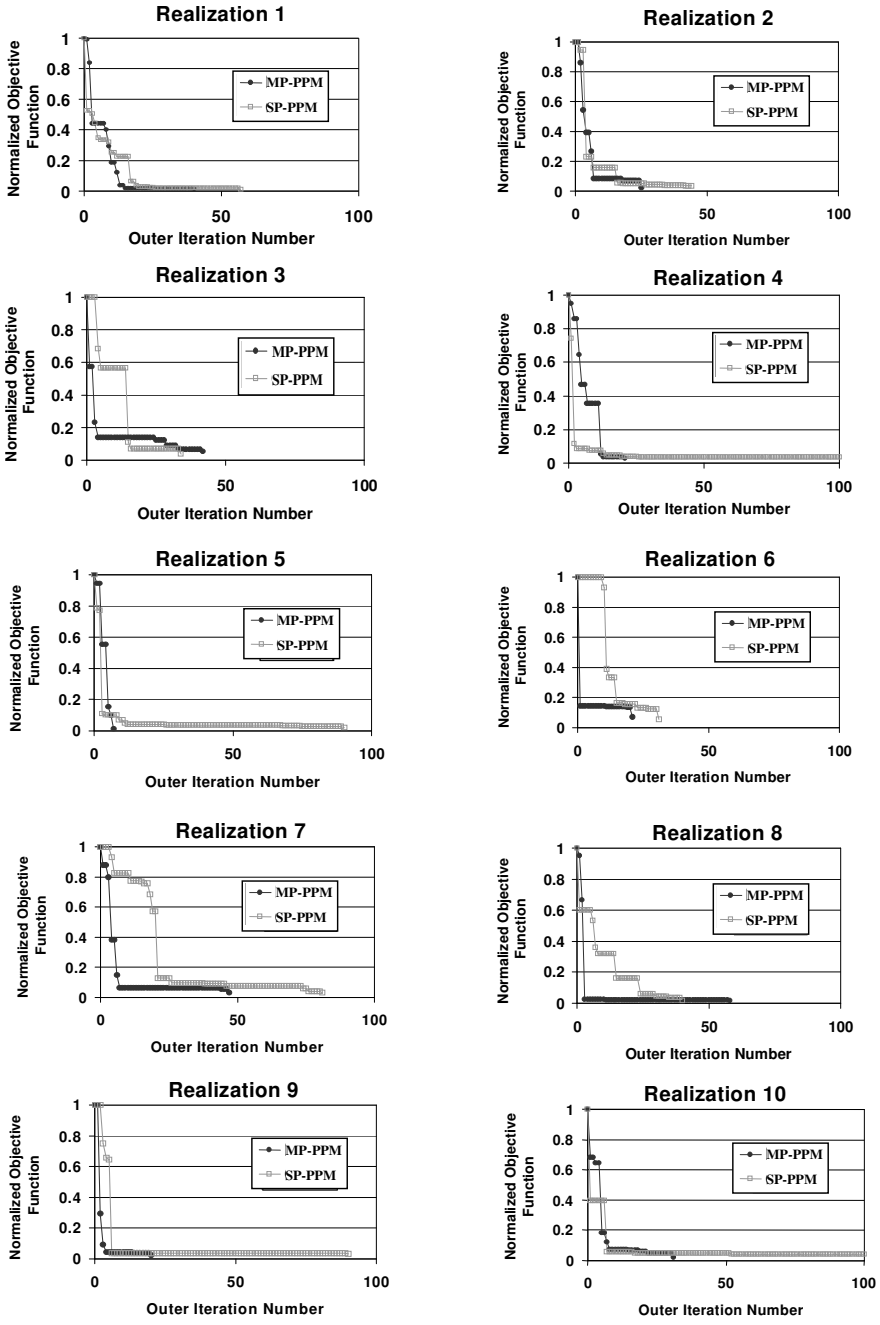

Figure 3.12: Objective function decrease for outer iteration for ten realizations.

Random Search

The probability perturbation method is a stochastic technique; thus, it is evaluated against a standard stochastic method, random search. In random search, a series of realizations are generated; flow simulation is completed, and the mismatch is calculated for each realization. The realizations are only conditioned to the static data, and the only difference among the realizations is the random seed. Therefore, each realization is similar to the initial realization of the PPM. As soon as the mismatch of one realization is less than the convergence tolerance, the process is stopped and that realization is the "history matched" one. In this method one is essentially searching a large number of realizations to find the ones that match production data. This is done for ten realizations with the 2D example to compare random search with the PPM. The averaged results are displayed in Table 3.2.

Table 3.2: Comparing random search and the probability perturbation methods on the average number of flow simulations for ten realizations.

Method	Average Flow Sim.
MP-PPM	216.1
SP-PPM	456.9
Random S.	3022.2

Random search takes an order of magnitude longer to converge than the probability perturbation methods with an average of over 3000 flow simulations required. Random search is one of the least efficient algorithms, so it is expected that the PPM would do better; however the fact that the PPM is so much faster is encouraging. By perturbing probabilities (intelligently modifying the realizations), we are ensuring that we are searching in an advantageous direction.

Polytope (Simplex) Optimization Method

Since multiple parameters, r_{Dk}, are being optimized, this problem could also be solved with a full multi-dimensional optimization routine. Therefore the MP perturbation

method is compared with one such routine, the polytope (simplex) method (Press *et al.,* 1989), for both efficiency and robustness. Polytope is chosen for comparison because it is one of the few optimization techniques that does not require derivative information, only function evaluations are needed. In the polytope method, one starts from K+1 sets of r_{Dk} values, $\{\{r_{D1} \ldots r_{DK}\}^1 \ldots \{r_{D1} \ldots r_{DK}\}^{K+1}\}$, where K is equal to the number of units. K+1 initial realizations are generated and flow simulation is completed on each one. This constitutes the initial guess for the polytope method (initial members of the polytope). The realization with the highest error (worst match) is perturbed into a new set of r_{Dk} values. A new realization is generated; flow simulation is performed, and a new error is calculated. The previous K unchanged sets of r_{Dk} and this new set become the new members of the polytope. This is repeated until the method has converged to a minimum. Since K+1 initialization runs are needed, this method becomes very time consuming as K increases.

Both the MP-PPM and the SP-PPM take approximately 7 function evaluations per outer iteration while the polytope method uses around 9 or 10 per iteration. The increase is due to the number of initialization runs. Since there are two regions, three initialization runs are needed for the polytope method, while the other methods need only one initialization run. Table 3.3 displays the results for the three different cases.

Table 3.3: Comparing the number of outer iterations for polytope.

	Number of outer iterations to converge to a tolerance of 0.001										Average	Average
Realization	1	2	3	4	5	6	7	8	9	10	Iteration	Flow Sim.
MP-PPM	40	25	42	21	7	21	47	58	20	31	31.2	216.1
SP-PPM	57	44	34	>100	90	31	81	40	90	>100	66.7	456.9
Polytope	60	30	2	26	>100	64	50	>100	33	7	47.4	453.2

The polytope method takes less outer iteration than SP-PPM, but since there are about 2 - 3 more flow simulations per iteration for the polytope method, the total number of simulations are approximately the same as SP-PPM and many more than MP-PPM. In some realizations the polytope method converged extremely fast (#3 and #10), but in

others it could not converge within the maximum number of iterations (100 iterations) that was set for this problem (i.e. it is not robust for this type of problem).

Error Tolerance with Prediction

All of the previous runs were matched to an overall tolerance of 0.001 (standardized water cut mismatch). This is a very restrictive tolerance; in essence, all points of production data have to be matched exactly for the method to be converged. In most real petroleum situations, the reservoir models can not be matched this closely, so the current model is also matched to a less restrictive tolerance, 0.01. When the models are matched to that less restrictive tolerance, there are some differences in how efficiently the models converge. Table 3.4 shows the results for the new criteria.

Table 3.4: Number of outer iterations for a higher tolerance, 0.01

	Number of outer iterations to converge to a tolerance of 0.01										Average	Average
Realization	1	2	3	4	5	6	7	8	9	10	Iteration	Flow Sim.
MP-PPM	12	6	3	12	5	1	4	3	2	5	5.3	36.0
SP-PPM	17	4	15	2	4	11	21	15	6	7	10.2	73.4

The MP-PPM still converges about twice as fast as the SP-PPM, but all models converge about 6 times faster as their lower tolerance counterparts. For the MP-PPM, the number of flow simulations is reduced significantly from about 216 per realization to about 36 per realization. While over 200 flow simulations is too many for most real applications, thirty-six is reasonable.

Figure 3.13 displays the difference for both MP-PPM and SP-PPM for two realizations. The solid horizontal lines depict the 0.01 tolerance. Notice how most of the objective function reduction is completed before the higher tolerance, but numerous additional runs are required to match to the lower tolerance.

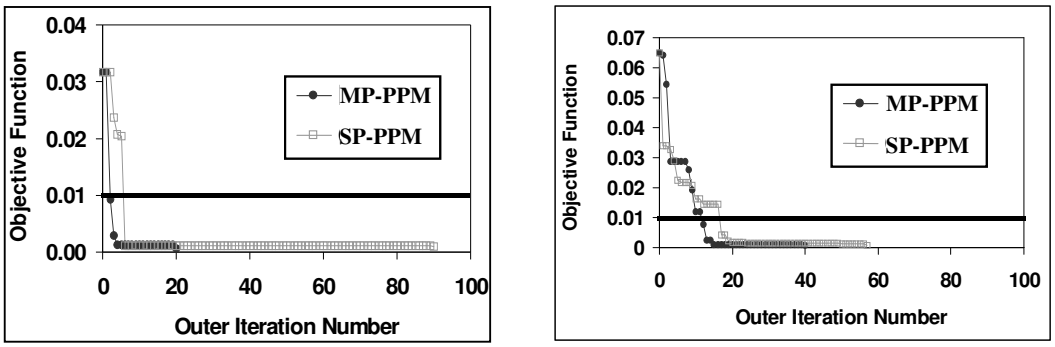

Figure 3.13: Objective function (non-standardized) decrease for two different convergence tolerances for realizations 9 and 1.

With a less restrictive tolerance, the match will of course be worse, and Figure 3.14 shows how the water-cut profiles differ for one realization. For the 0.001 tolerance, the production match is almost exact for both wells, but for the 0.01 tolerance, Well P2 does not match as effectively. Recall that the convergence is based on the overall objective, so one well may be matched better than another. For this realization almost all of the error results from P2, but for different realizations converging to 0.01, an equivalent amount of error is allocated more evenly between the two wells. In many practical reservoir applications, the match displayed in Figure 3.14 would be acceptable.

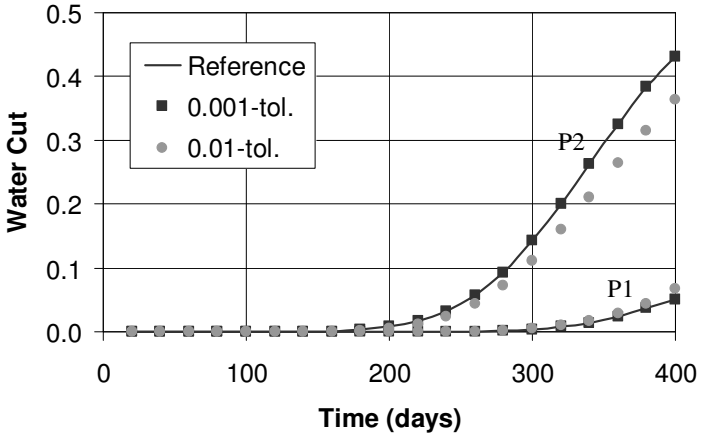

Figure 3.14: Water cut match for two tolerances.

In fact, the quality of the match is not what is desired when constructing and history matching reservoir models. The goal is to have reliable predictions from the models, so the matched models are run for an additional 400 days. Predictions from the 10 realizations that are matched to a tolerance of 0.001 will be compared with predictions from the realizations that are matched to 0.01. When only the original two wells are used for prediction, the 0.001 tolerance models continue to outperform the 0.01 tolerance models. The average mismatch for the 0.001 tolerance is 0.015 while it is 0.027 for the 0.01 tolerance models. This is expected when the history match and the prediction is completed on the same wells. The next data points are highly correlated with the previous ones, so if the previous ones are already matching better, the next ones should also be matched better.

However often models are built to predict production from new wells drilled in the reservoir. Two new wells are added to verify the prediction power. The new wells are included at the end of the history matching period (400 days), and they are produced for another 400 days. One example is examined where the new wells are placed in the corners of the model, and another is studied where the new wells are closer to the original wells. The locations of the wells are displayed in Figure 3.15, and the reference water cuts for the new wells are also presented in that figure.

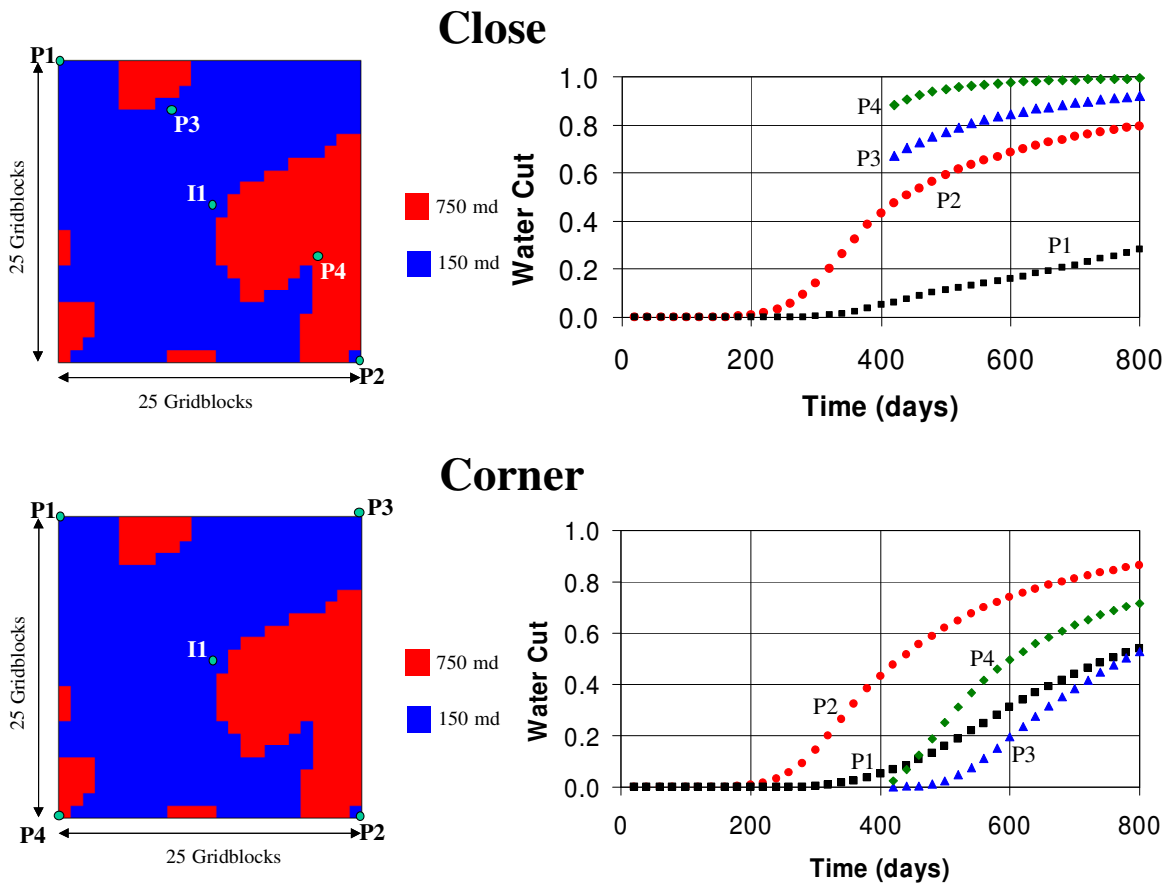

Figure 3.15: Reference model with new well locations (corner and close), and reference water cut for all wells.

Again, the two different error tolerances are compared for the 10 realizations. The average mismatch (least-square error) for the two new wells and the two original wells is presented in Table 3.5. They are displayed for both history matched tolerances and the initial realization (before any history matching).

Table 3.5: The least square error for predicting from two new wells and two original wells. The values are averaged over 10 history matched realizations (MP-PPM). It is presented for three case and three different matching tolerances.

Case	Two New Wells			Two Original Wells		
	0.01	0.001	Initial	0.01	0.001	Initial
No new wells	--	--	--	0.027	0.015	0.090
New wells - Corner	0.043	0.047	0.067	0.031	0.020	0.091
New wells - Close	0.012	0.011	0.031	0.034	0.022	0.111

The two original wells always are better predicted for the closer tolerance match due to the reason previously discussed. However, the two new wells have about the same amount of error for both tolerances. If the model is constructed to get accurate values for future production from the current wells, then converging to a tight tolerance can improve the prediction significantly; however, if what we desire from the model is prediction from newly drilled wells, a less restrictive matching tolerance will be as good with much less time spent history matching.

When the new wells are “close” to the original wells, the average error is much lower than when the new wells are in the “corners” simply because the area close to the original wells is better constrained by the history match. Nonetheless, all cases are improved over the initial error, so even for the case where the new wells are in the corners, history matching this model enhances the prediction. Figure 3.16 has water cut plots for wells P2 and P3 for the initial case and the 0.01 and 0.001 tolerance cases. The symbols represent the reference data and the lines are for the ten history matched realizations.

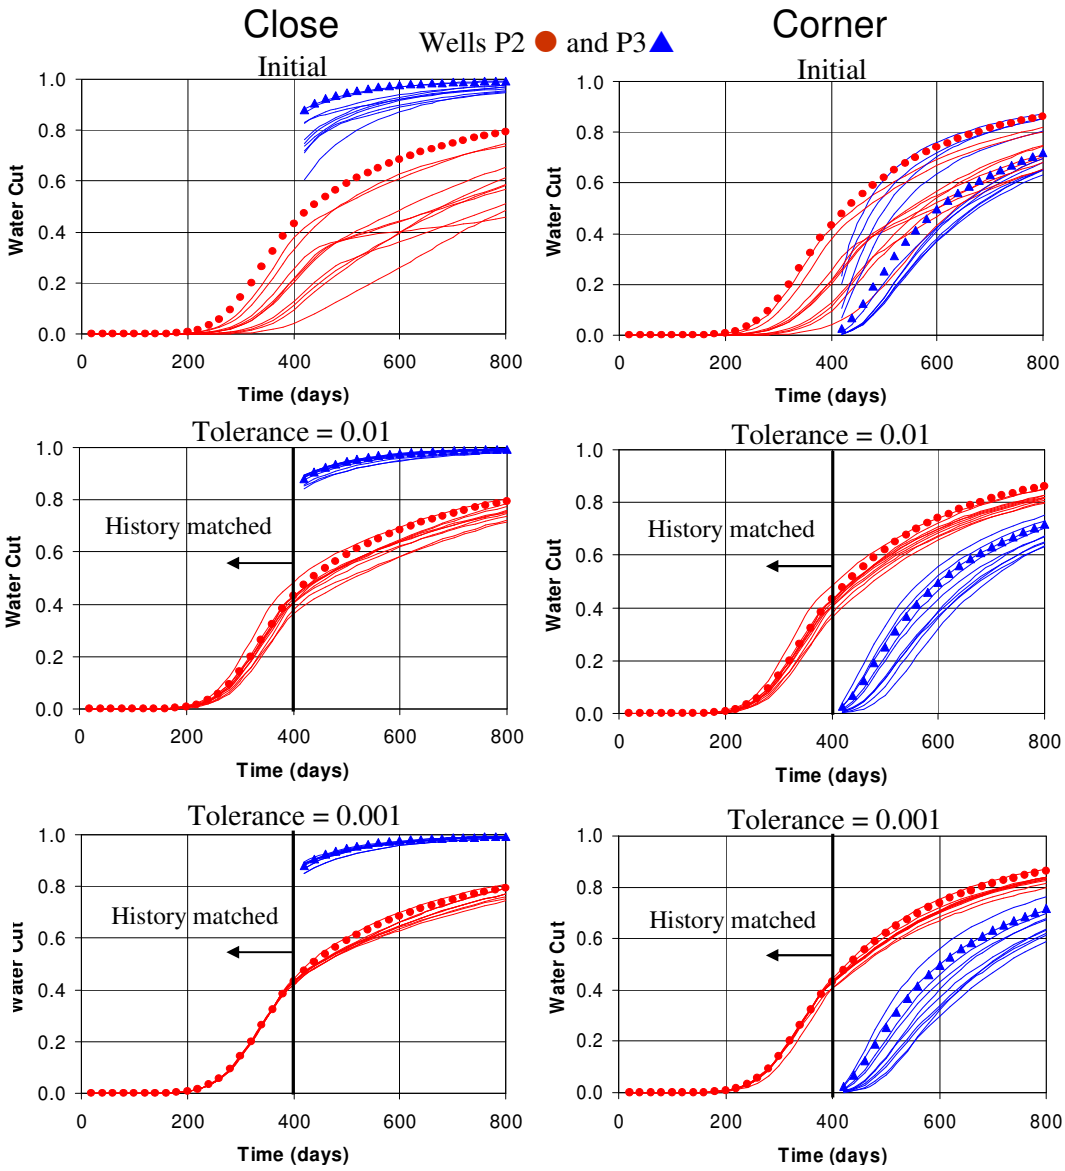

Figure 3.16: Match and prediction for an old well (P2) and a new well (P3). The reference (symbols) and ten realizations (lines) are shown for two cases. One case where the new well is "close" to the old wells, and one where the new well is in the "corner" of the model.

Initial Perturbation (r_D) Values

Whenever a new inner loop is started, a set of initial r_{Dk} values needs to be specified. For the previous methods, the initial parameter value always equals 0.50 simply because it is

the middle of the range [0, 1] for r_{Dk}. The initial values can have an impact on the convergence rate, so the sensitivity of the initial r_{Dk} values is further studied. Initial values of 0.75 and 0.25 are used as alternatives. Table 3.6 presents the sensitivity of the MP-PPM for different initial perturbation parameters.

Table 3.6: Comparing initial r_{Dk} values for efficiency.

	Number of outer iterations to converge to a tolerance of 0.001										Average	Average
Realization	1	2	3	4	5	6	7	8	9	10	Iteration	Flow Sim.
MP-PPM	40	25	42	21	7	21	47	58	20	31	31.2	216.1
Init = 0.75	11	30	22	33	37	25	38	30	98	20	34.4	241.1
Init = 0.25	12	31	12	13	39	6	14	22	62	48	25.9	179.6

When the initial parameter is set to 0.75 there is a slight decrease in the efficiency of the method, and when the initial parameter is 0.25 there is a slight increase in the efficiency. For realizations that are converged to very low tolerances, a smaller value for the initial parameter is preferred. In these types of models, a large portion of the model converges in the first couple of steps (Figure 3.12), but most of the time is spent to reduce the objective the last small amount. During this time, only small changes in the model are needed to reach the desired tolerance. Therefore, a smaller initial value performs better because it causes a smaller initial change in the model.

Table 3.7: Comparing initial r_{Dk} values for efficiency when converging to 0.01.

	Number of outer iterations to converge to a tolerance of 0.01										Average	Average
Realization	1	2	3	4	5	6	7	8	9	10	Iteration	Flow Sim.
MP-PPM	12	6	3	12	5	1	4	3	2	5	5.3	36.0
Init = 0.75	11	4	6	4	5	1	2	8	2	2	4.5	30.9
Init = 0.25	7	10	3	4	5	1	5	5	3	14	5.7	39.3

However, when converging to a less restrictive tolerance (0.01), the results are quite different. The higher initial value (0.75) causes faster convergence (Table 3.7) for the 0.01 tolerance case. Early in the optimization when production match is very poor, large perturbation may be desirable to get the simulated production close to the reference production. Once the match is close, however, only small changes are needed to get the simulated data to match exactly the observed data. Moreover, the objective function

decrease also demonstrates this behavior (Figure 3.13). It decreases very fast early (significant model changes), and at the end it reduces very little (small changes in the model). The choice of the initial r_{Dk} values is able to reduce the number of flow simulations by about 20%, and while 20% is not a large amount, it can be important in real reservoir models that take hours of CPU time to complete.

Improving Efficiency

The MP method reduces the number of flow simulations by half compared to the SP method; however, when converging to a tight tolerance (0.001), 180 flow simulations are required even for the best case (initial r_{Dk} values equals 0.25). This is a large number of flow simulations, and for most real field applications, this is too many to be practicable. Therefore, another approach to reduce the number of flow simulations was investigated.

Figure 3.12 showed that to converge to a tight tolerance there is usually a long tail where the objective function does not improve many iterations. However, during this period each inner loop continues to require around seven flow simulations. When examining the inner iterations in detail, we found that if one of the first three flow simulations does not improve at least some over the previous best, then most of the time (85%), the next four simulations will not improve over the previous best either, and this iterations will not decrease the overall objective.

This new insight caused the optimization routine to be modified slightly. After three flow simulations, the best overall objective from the current inner loop is compared with the previous overall best. If one of the three current matches is better than the previous best, then the method is continued as normal; however, if none of the three are better, the current inner loop is ended, the random seed is changed and a new inner loop is started.

For the ten realizations, the average history match is achieved in almost half as many flow simulations with this modified routine. For the case where the initial r_{Dk} values are 0.50, the average number of flow simulations is 129.7 and for the case where the initial

r_{Dk} is 0.25, the average equals 105.1 flow simulations. Some of the subsequent examples and case studies use the "new" optimization method, while others use the method with a full seven flow simulations in every inner loop.

Response Surface

By examining the inner loop optimizations in Figure 2.3 and Figure 3.11, it is apparent that the objective function in the inner loop is not a smooth function. As discussed previously, this is caused by the discrete nature of perturbing facies bodies. As bodies disappear and reappear, the production response does not change smoothly. Since only two parameters (r_{D1} and r_{D2}) are being optimized for this case, the response surface (objective function value surface) can be displayed, and the fluctuations are even easier to observe when the surface is examined.

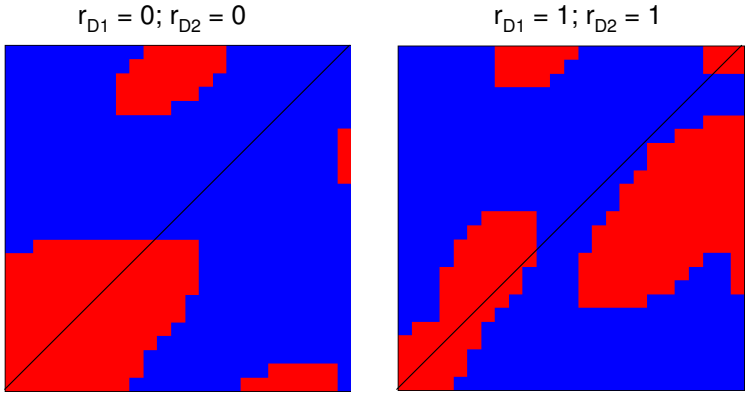

Figure 3.17: The two realizations that are combined to generate the response surface.

The surface is constructed by considering the two realizations displayed in Figure 3.17 as end points of the perturbations. Each r_D value is discretized into segments of length 0.02, so there are 51 values in the range of r_{Dk} [0, 1]. Combining the two r_{Dks} requires 2601 models to be generated, flow simulations to be performed, and objectives to be calculated. Since the example is so simple (each process only takes about 15 seconds), the entire surface can be generated in less than a day. The response surface is displayed

in Figure 3.18; r_{D1} and r_{D2} are on the x and y axes and the objective function values are on the z axis.

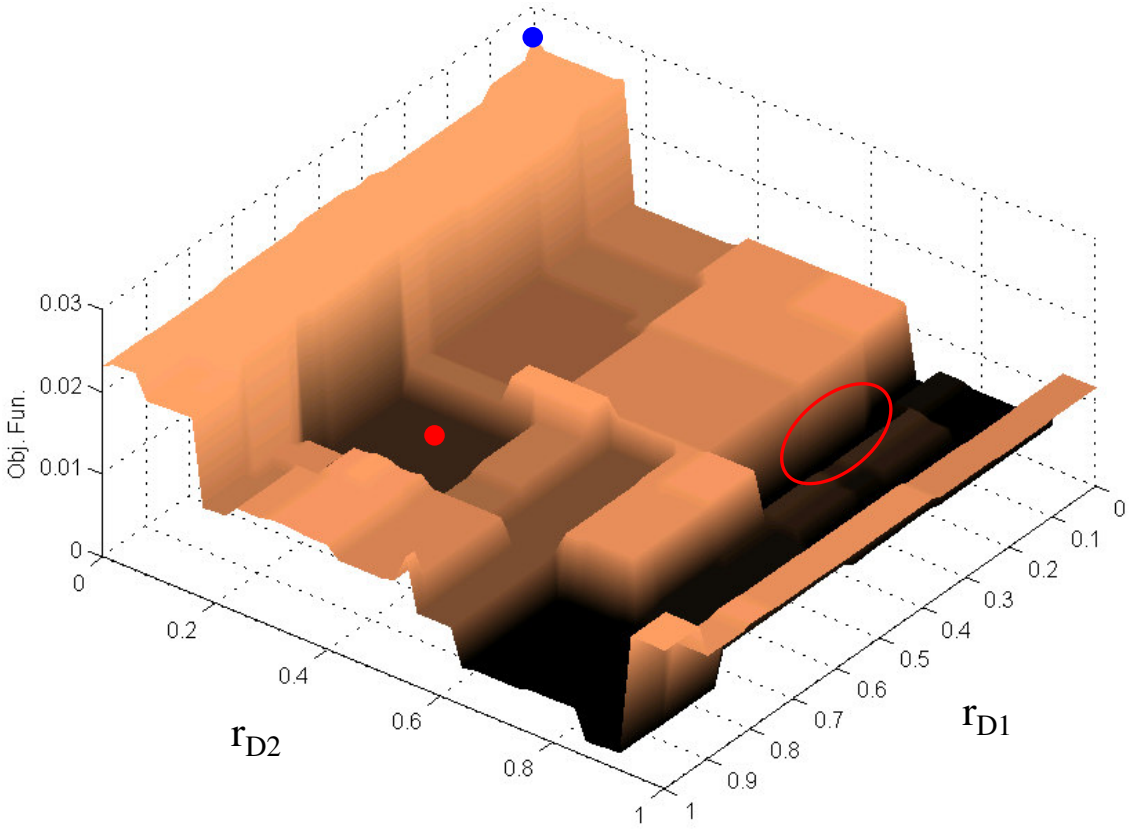

Figure 3.18: Response surface for combining two realizations at different r_{Dk} values.

The surface appears like stair steps; there are many flat parts of the surface where the objective value is the same. The large flat regions exist because the model is relatively small, and there are only so many different configurations that are possible. For a larger model, the surface likely would be more peaked. Due to the discrete nature of the bodies, there are sharp changes between the flat areas. As the r_{Dk} values are changed, bodies may move around, appear or disappear, and these changes cause large jumps in the objective function.

The location on the surface of the minimum is circled in Figure 3.18. This occurs where r_{D1} is between 0.14-0.28 and where r_{D2} is 0.68-0.70. The MP-PPM does not find that minimum, but converges to a local minimum at $r_{D1} = 0.64$ and $r_{D2} = 0.29$ (highlighted by a point in Figure 3.18). While the MP-PPM did not find the global minimum, it did considerably improve the match over the initial point ($r_{D1} = 0.0$ and $r_{D2} = 0.0$) that is marked with a blue dot in Figure 3.18. The MP-PPM need not necessarily find the optimum realization in every inner loop; it is simply too costly to find it. Furthermore, the optimum realization is not required because the current best realization becomes the next initial realization and the inner loop is rerun with a new random seed. In fact, changing the seed makes the method a stochastic (global) search.

Reproducing "true" r_D values (and prediction)

The 2D example is used for one final numerical experiment aimed at answering the following question. If the reference solution is in the search space between the two equiprobable realizations, will the probability perturbation method be able to find that reference?

To build the "true" reference, we select two realizations and choose randomly an r_D value (or set of r_D values) between zero and one. One of the realizations will be dubbed the initial realization so the P(A|D) can be calculated for the r_D values using Eq. (3.2). Thus, a new realization is generated that has aspects of the two previous realizations. Flow simulation is completed on this model, and the corresponding production data is used as the reference data in the objective function.

Then the Brent algorithm is started using the same two realizations and an initial r_D value of 0.50. Essentially one inner iteration of the probability perturbation method is completed to see if the "true" solution can be recovered. The process is repeated starting with two new realizations and a new random r_D value. After 1000 trials, the following results were ascertained.

When one r_D value is used (SP method), the PPM is able to find the true solution 37% of the time, and when the MP method with two r_{Dk} values is used, the true solution is only found 10% of the time. The definition for finding the true solution is: the objective equals zero. From Figure 3.18, the surface is seen as a series of flat steps, so the r_D does not have to be exactly the same as the one specified. It only has to be close enough for its response to be on the same flat area. In practice, this means that the r_D values should be within approximately ± 0.03 of the solution.

Most of the time, the PPM is trapped in a local minimum. The response surface is very discrete (Figure 3.18), so it is easy to see why this happens, especially when there are multiple r_{Dk} values. The goal of history matching is not to find the "true" solution because the true solution is never known. The goal is to create a model that improves the prediction of future response from the model. Even if the true solution is not found, a model that matches the production data can still be found by running the PPM. The question remains, however, if the matched model is any good (has any prediction power).

If the true solution is not reproduced for the two parameter case, the MP-PPM is continued until it converges to a tolerance of 0.001. Then two new wells are added to the model, and the production is compared to the previously generated reference. The two new wells are added at the same location as the "close" case from the Error Tolerance example (Figure 3.15).

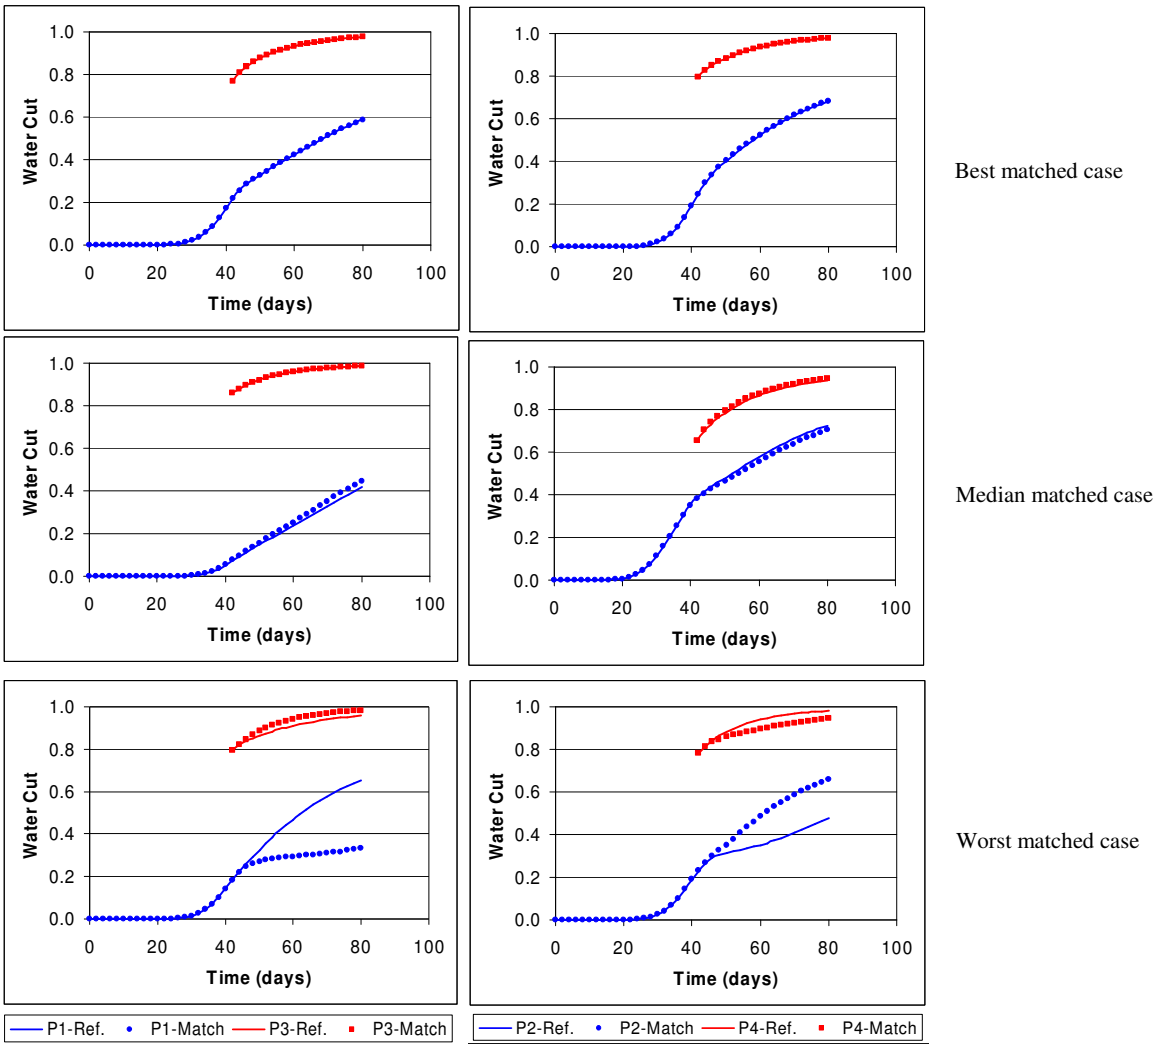

Figure 3.19: Best median and worst matched prediction realizations out of ten realizations.

Ten runs where the true solution was not recovered are history matched, and eight of the cases match very well, and only one matches poorly; the other match is adequate. Figure 3.19 shows the results for the best, median, and worst matched cases. The best and median matched realizations are extremely good. The worst matched case is sufficient for the two new wells (P3 and P4), and only for the two original wells is the match poor.

This shows that even if the true solution is not found, the PPM is able to find a model that has good prediction.

Summary

The 2D example in this section showed that the MP-PPM is more efficient than SP-PPM. For each perturbation, the efficient MP-Brent optimization requires the same number of flow simulations as the SP method, but the new method allows different areas to be perturbed by different amounts greatly improving its ability to match. Both probability perturbation methods are an order of magnitude faster than random search. Although improvement is expected, the magnitude of improvement is impressive. Also a number of sensitivities were examined to see the affects on efficiency of the MP-PPM. This includes sensitivity to initial r_D values and to convergence tolerances. For tight tolerances, a lower initial r_D value was faster, and for less restrictive tolerances, a higher initial r_D was more efficient. Further improvements to efficiency are achieved by checking for improvement early in the inner iteration. Because the example is simple with only two wells and two facies, we were able to explore the response surface and search for a true solution. Although the method is not always able to find the true solutions or best location, the PPM is still able to find a solution and the predictions for the resulting models are quite good.

3.4 Additional Examples

Many different aspects of the MP-PPM were studied with the previous 2D example; however, the flexibility and applicability of the method were not fully showcased since there was only one model. This section explores four diverse examples. The first synthetic example is similar to the previous example, but a slightly different geologic setting is used. The second example is a fractured reservoir where pressure is being matched. The third one demonstrates how continuous parameters such as gridblock permeability can be perturbed, and the last example illustrates how the method works for multiple facies and in 3D.

Similar Geology Everywhere

Similar to the example from the previous section, this is a two facies model and water cut from two wells is being matched, but there are some important differences between this example and the previous one. In this example, the facies distribution is fairly similar everywhere in the model, whereas in the previous example (Figure 3.9), the distribution was quite different in the two regions. Figure 3.20 shows the reference model with the three well locations (one injector and two producers) and the water cut for the two producers. The water cut curves for both wells have almost the exact same profile; this is further evidence that the amount of connectivity between the injector and each well must be similar.

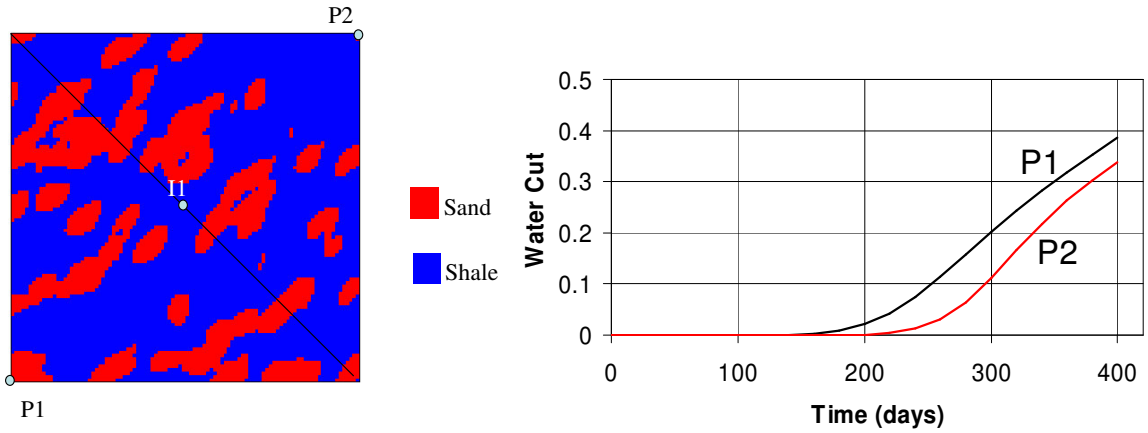

Figure 3.20: Reference model and production from two wells.

This example uses a 100 by 100 block model. The red facies is a sand facies with a permeability of 750 md while the blue facies is shale with a permeability of 150 md. The porosity is constant for both facies and equals 20%. The top depth of the model is at 1000 feet and the initial pressure is 655 psi. Also, no movable water is initially present in the model although the initial water saturation equals 15%. The production wells are producing at a fixed liquid rate of 100 barrels per day, and the injection well is injecting 200 barrels of water per day. The fluid properties and relative permeability curves (Figure 2.6) are the same as the previous example.

For this example the two types of the probability perturbation method (MP and SP) are compared. For the MP examples, regions of constant r_{Dk} values are used, and Figure 3.20 displays the two regions (again Voronoi polygons are used). In this case, the sand bodies from the reference cross the regional dividing line. Table 3.8 shows the results for ten different realizations.

Table 3.8: Number of outer iterations for 100 by 100 model to converge.

	Number of outer iterations to converge										Average	Average
Realization	1	2	3	4	5	6	7	8	9	10	Iteration	Flow Sim.
MP-PPM	4	13	3	46	6	15	2	9	1	12	11.1	77.3
SP-PPM	15	3	3	11	24	30	4	13	2	53	15.8	107.2

Similar to the previous example, the MP-PPM does better than the SP-PPM, but for this case there is less incremental improvement. For the previous example, the MP-PPM converged in about half as much time as SP-PPM, but for this example MP-PPM took about two-thirds as much time. Since the model is similar in both regions and the water-cut profiles are similar for both wells, the amount of perturbations required for each region will often be the same. Therefore, the SP method does well for this example, but the MP method is still slightly better.

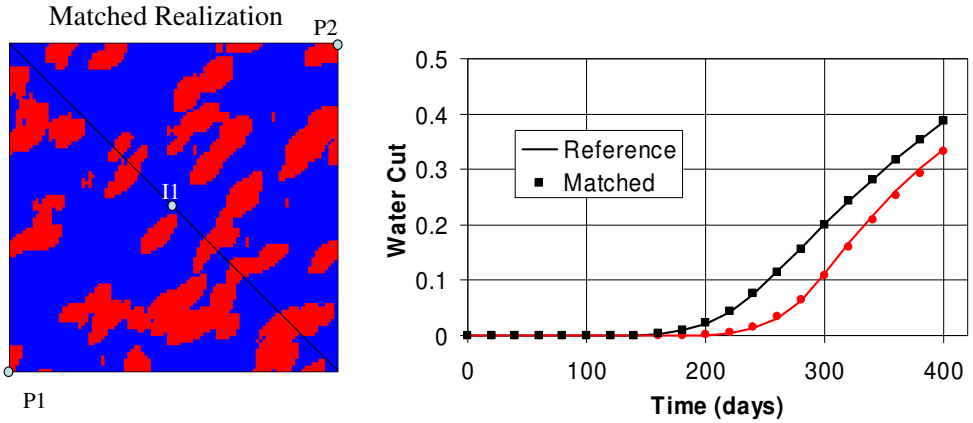

Figure 3.21: Matched realization and water cut profiles.

Even though the two regions are being perturbed separately, there are no geologic discontinuities formed along the boundary. Figure 3.21 shows a matched model that has continuous sand facies across the dividing line. No artifacts are created because the search neighborhood crosses the region boundary and the prior geologic information (P(A|B)) is always used. Further discussion is provided in Section 3.1 about reproducing body shapes across regions boundaries.

Pressure Match

This example displays some of the robustness of the proposed method. The model for this example is geologically more complex than the previous cases. It is a fractured reservoir that has three separate fault blocks. The fractures represent 15% of the model.

The center fault block has been fractured in the orthogonal direction to the two outside fault blocks. Furthermore, each fault block has different physical properties. The left block has 100 md matrix and 2 md fracture permeability, the middle block has 50 md matrix and 25 md fracture permeability, and the right block has 150 md matrix and 3 md fracture permeability. The two outside blocks have essentially sealing fractures while the center block fractures allows some flow through the fractures.

The initial pressure is 1000 psi and the reservoir is 2000 feet deep. The grid model is 100 by 100, and there are 9 production wells and no injectors (Figure 3.22). Each well is producing at a fixed oil rate of 250 barrels/day. The proposed method can match different types of production data, and for this example the bottom hole well pressure is matched for the nine wells. Each fault block has 3 wells and is treated as a separate region (geologically defined regions) for the MP-PPM. In this case, three r_{Dk} values must be optimized with the MP-Brent routine.

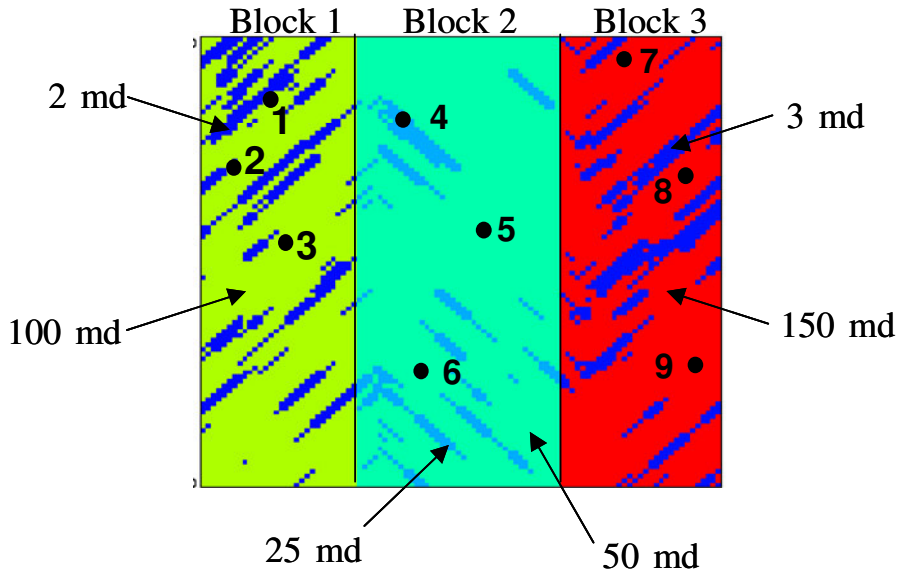

Figure 3.22: Fractured model reference within three fault blocks. Nine well locations and facies permeability displayed.

Starting with the same initial model, both the SP-PPM and the MP-PPM attempt to match the reference pressure data. Realizations are generated using the program snesim (Strebelle, 2002). The training image is displayed in Figure 3.23, and to account for the different directions of fractures, the angle function is used from the snesim program. This allows the training image to be rotated for different parts of the reservoir model.

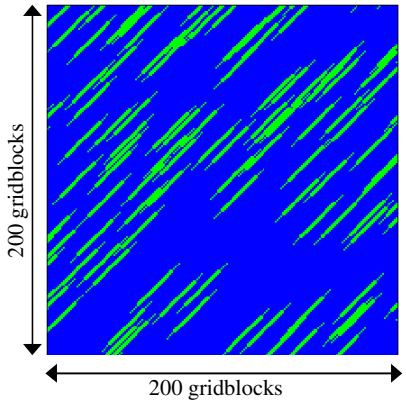

Figure 3.23: Training image for simple fracture example.

The realizations are conditioned to the hard facies indicator data at the nine well locations. Realizations are displayed in Figure 3.24 for the MP and SP methods. While they visually appear similar, the pressure response from the two models is quite different.

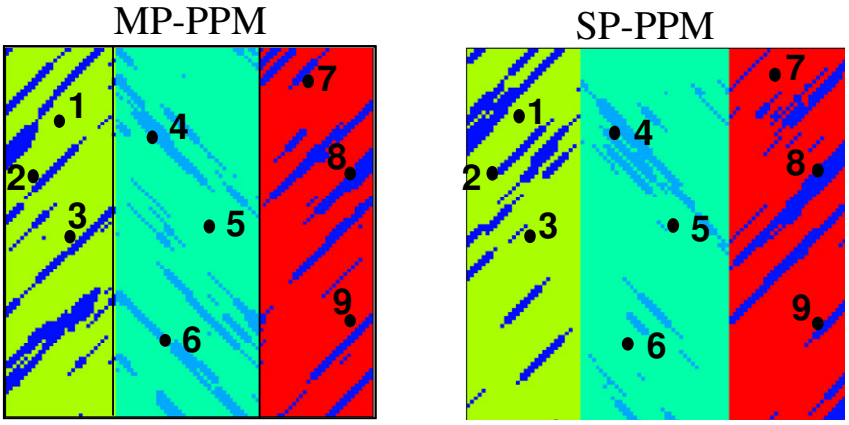

Figure 3.24: Two realizations for fractured example; one from each method.

After only 9 outer-iterations, the regional method converges to a solution; however, after 100 iterations, the SP-PPM method has not converged. The match for three wells is shown in Figure 3.25 for both the MP-PPM and the SP-PPM methods. The other six wells have a similar response. All wells are matched for the MP-PPM method, and while some wells are matched for the SP-PPM method (e.g. Well 5), others are not (e.g. Well 2 and Well 8).

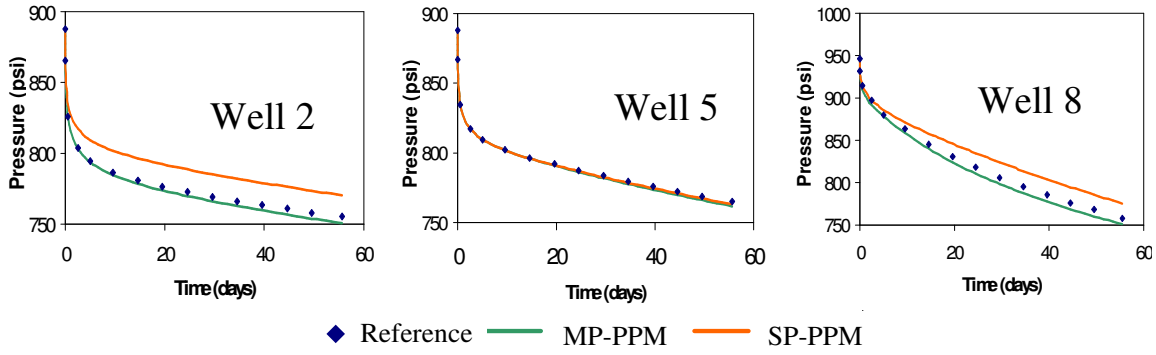

Figure 3.25: Pressure match comparison between MP-PPM and SP-PPM for one realization.

For ten realizations history matching is completed to a low tolerance. The SP method and the polytope method are compared to the MP method. Figure 3.26 shows the outer objective function for all three methods for the ten realizations. Polytope uses about three more flow simulations per outer iteration than the other two methods, but still converges slower than MP-PPM for most of the realizations. For all realizations, SP-PPM converges much slower than MP-PPM.

Table 3.9: Number of outer iterations for fractured model to converge.

	Number of outer iterations to converge to a tolerance of 0.001										Average	Average
Realization	1	2	3	4	5	6	7	8	9	10	Iteration	Flow Sim.
MP-PPM	9	33	14	14	6	14	15	5	16	37	16.3	113.3
SP-PPM	>100	92	62	47	76	33	58	14	20	19	52.1	366.9
Polytope	38	18	36	42	21	45	16	2	32	7	25.7	263.9

From Table 3.9, one can observe that MP-PPM converges about three times faster than the SP method. The polytope method converges slightly faster than the SP-PPM, but it converges much slower than MP-PPM, especially when the average number of simulations is considered. Due to its poor performance, this is the last example that the polytope method will be used for comparison.

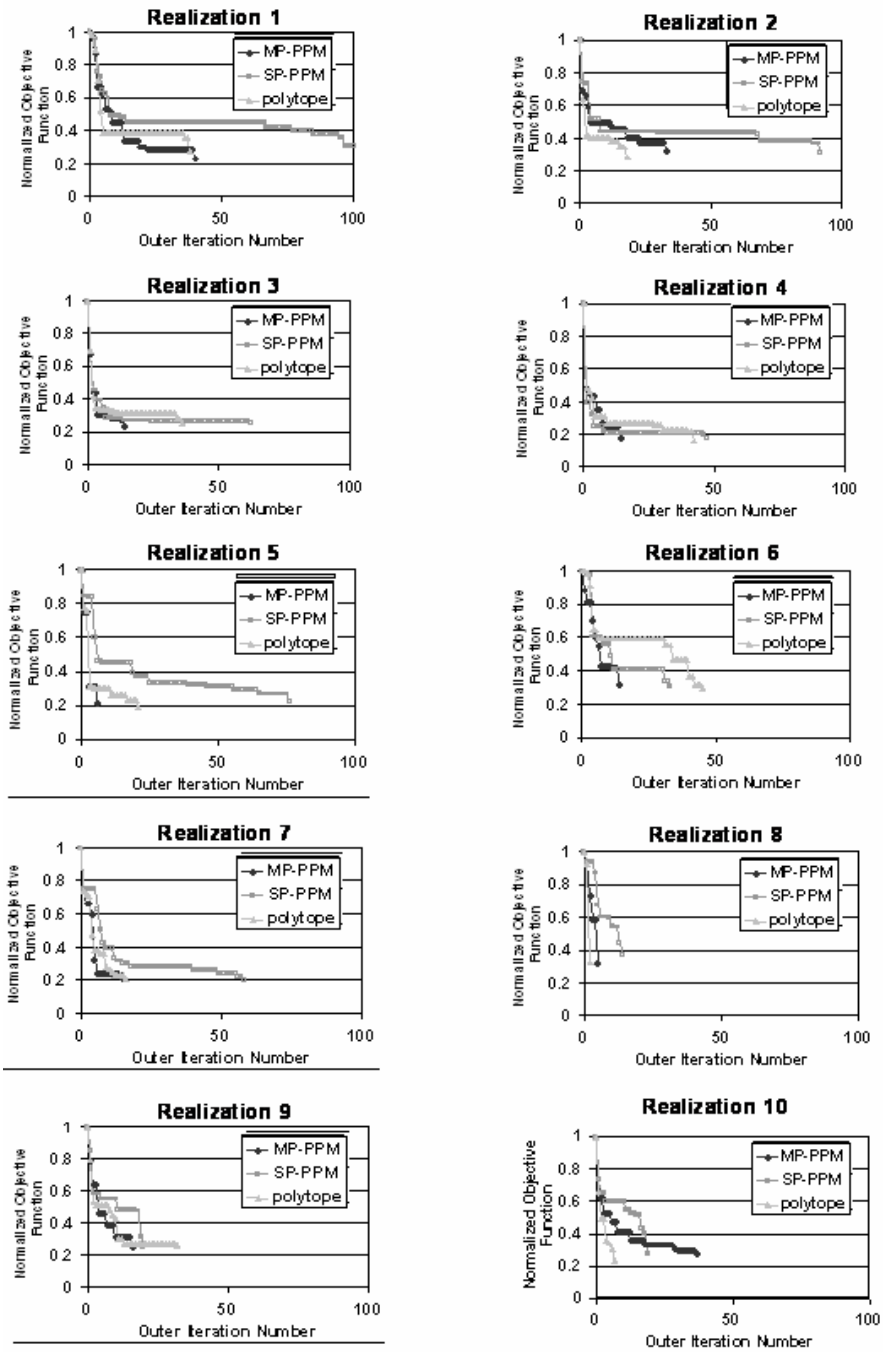

Figure 3.26: Outer iteration convergence for ten realizations for MP, SP and polytope methods.

Although this example is still geologically very simple compared to most field examples, using the MP-PPM is clearly superior to matching the entire model using a SP perturbation or by using a full multi-dimensional optimization routine. As the models become geologically more complex, the speed of convergence for the MP-PPM likely will increase over the SP-PPM.

Continuous Properties

To further demonstrate the diverse applicability of the probability perturbations method, an example is matched by perturbing a continuous property, gridblock permeability. Also this is the first example where kriging is used to interpolate the r_{Dk} values, and it is compared with the region based approach.

The reference for this example is a 2D model that is 100 by 100 gridblocks. The permeability is constructed from sequential Gaussian simulation (Deutsch & Journel, 1998). The correlation range is 40 in the x direction and 8 in the y direction. The data is transformed from the standard normal distribution using the following equation:

$$\text{Perm} = 10*3^{st} \tag{3.4}$$

where st stands for the standard normal Gaussian values. The porosity is constant everywhere in the reference at 0.25. There are five wells that have been producing on primary production for ten years. The oil rates vary for the five wells and are given in Table 3.10.

Table 3.10: Fixed production rate for five wells in continuous parameter example.

Well Name	Rate (bbls/day)
1	65
2	60
3	90
4	50
5	40

There is no water or gas initially present in the reservoir nor is any injected, so there are only oil rates and they are fixed in the history matched models. Therefore, bottom hole well pressures at the five wells are matched. The initial pressure in the reservoir is 1000 psi, and once a well pressure drops to 40 psi, that well is shut in. The well locations, permeability from the reference model and pressures from the five wells are displayed in Figure 3.27.

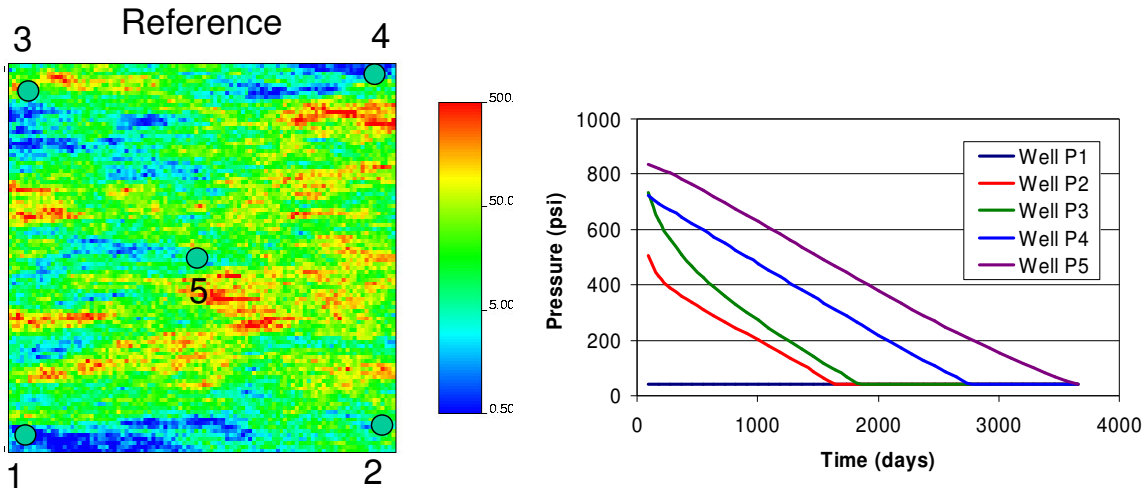

Figure 3.27: Reference model and pressure data from continuous example.

The PPM is implemented slightly different for continuous parameters than indicators. In the continuous case PPM, the "next" realization continues to be conditioned to soft data that comes from production data; however, for the continuous case, P(A|D) is not calculated. Instead, collocated co-simulation (e.g. sgsim) is completed where the soft data is the previous "best" realization and the correlation coefficient that describes the correlation between the primary and secondary data is equal to $(1-r_D)$. The correlation coefficient relates the amount of weight to give to the previous best realization and the amount to give to the new realization (Section 2.1). One correlation coefficient is tied to each production data unit, and MP-Brent is used to find the set of correlation coefficients and consequentially the realizations that best match the production data.

Results are compared for the single parameter (SP) method and the multiple parameter (MP) method. When the MP method is used, two different ways to fill the grid with r_{Dk} values will be presented. One method is similar to previous examples where Voronoi polygons are used to define regions and the r_{Dk} values are constant within the region. The other way is to use kriging to interpolate the values between the wells. The horizontal to vertical range ratio is still 5:1, but overall the ranges need to be five times larger (about three times the interwell distance) so that a smooth transition occurs between the wells. Figure 3.28 portrays the region definitions for the Voronoi polygons and one possible example of how the kriging defined values would appear.

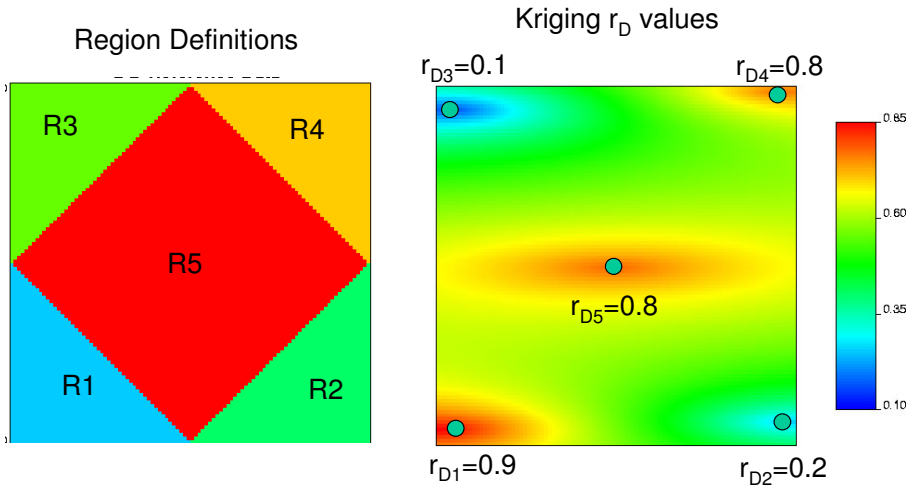

Figure 3.28: Two methods to interpolate r_{Dk} values: regions of constant value and kriging.

The permeability is perturbed until the pressure match is satisfactory. All three methods are matched to the same tolerance, and they are all completed for ten realizations. The average number of outer iterations and flow simulations is presented in Table 3.11. The MP methods are about 3 times faster than the SP method, but both MP methods require about the same amount of time to converge. Similar results are produced when using regions or kriging interpolation because both methods fill the model essentially the same.

The corners are filled with r_{D1} - r_{D4}, and the center is filled with r_{D5} (Figure 3.28). A more complex example may produce greater differences between the two techniques.

Table 3.11: Convergence comparison for three methods.

Method Type	Average Iterations	Average Simulation
SP-PPM	52.8	179.2
Kriging	16.9	62.9
Voronoi	16.1	60.2

The average number of flow simulations per iteration is reduced compared to the previous examples. Instead of running around seven for each inner iteration, now after three flow simulations we check to see if any of the three are better than the previous best. If at least one is better, then the method continues as normal, but if it is not, then the inner loop is stopped, and a new random seed is selected and the method starts a new inner loop. This allows a high quality pressure match to be achieved with a reasonable number of flow simulations.

Figure 3.29 presents an initial and matched model for one history matching realization along with the reference model. This model used the voronoi polygons for r_{Dk}, and once again, no artifacts at the region boundaries are noticeable in the matched realization.

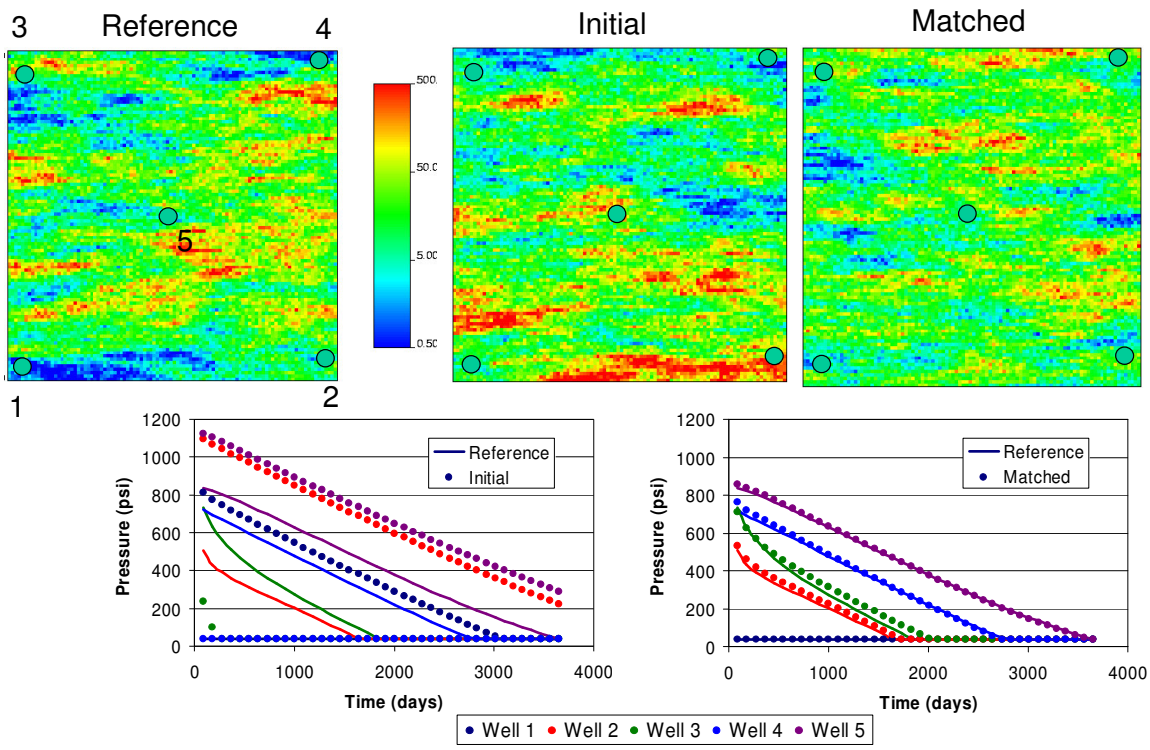

Figure 3.29:Reference, initial and matched models and pressure for continuous case.

The reference, initial and matched pressure data are also presented in Figure 3.29. The initial pressures for all five wells are far from matching, and by examining the initial model, the reason is easily understood. Wells 1 and 2 are in much higher permeability features in the initial model than in the reference. This allows the same amount of rate to be produced from a smaller drawdown, thus, the pressures for these two wells are much higher for the initial than the reference. Conversely, wells 3 and 4 are surrounded by much lower permeability in the initial than the reference; therefore, they need a larger drawdown to meet their target rates. The matched model better represents the permeability distribution in the reference and provides an extremely good match.

Multi-Facies 3D Example

So far only 2D examples have been used to explain the method, which is sufficient to showcase the details of the technique. However, most reservoir models in the petroleum

industry today are fully 3D in complexity; to be practical, the method must also work well in 3D. This example and all of the remaining ones are in 3D. Three facies types are present, each facies has the same constant and known porosity of 25%, but each facies has a different, constant but known permeability. A background facies has a permeability of 10 md, a NE elliptical shaped high permeability facies of 1000 md is eroded by NS elongated low permeability barriers of 0.1 md. The proportion of each facies is assumed to be known and is 45%, 35%, and 20% respectively for the three facies. A Cartesian grid of 100 x 100 x 10 cells is used.

The model is a two phase oil-water system. The viscosity of the oil is 1.0 centipoise (cp), and the water is 0.30 cp. The density of the oil is 54.20 lb/ft^3 and the water density equals 62.97 lb/ft^3. The oil-water relative permeability curves are provided in Figure 3.30.

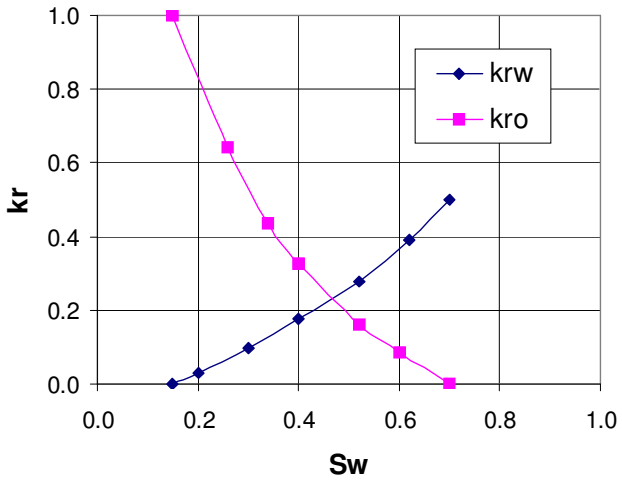

Figure 3.30: Relative permeability for 3D example.

As prior information we consider that the style of facies architecture is known, including the proportion of each facies. A multiple-point statistics algorithm is used to simulate the facies model. This algorithm requires a 3D training image depicting the style of heterogeneity present; Figure 3.31A shows the training image used to generate the realizations.

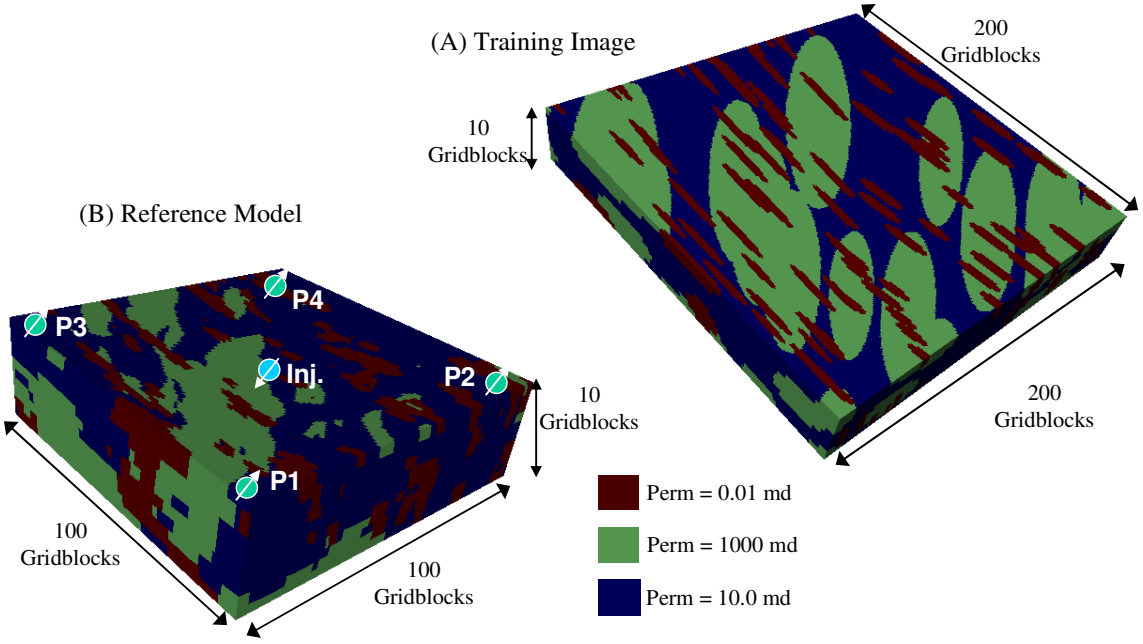

Figure 3.31: Training image, reference model with well locations, and permeability for each faices.

The reservoir is initially saturated with 85% oil and 15% water, no gas is present. The water table is below the reservoir and remains so during production. The reservoir contains four production wells in the corners of the reservoir and one injection well in the center (Figure 3.31B). The injector pumps water in the reservoir at constant rate of 400 barrels per day. The data that is matched consists of 10 years of water cut at the four producing wells. Figure 3.32 displays the reference water cuts for the four wells. One single flow simulation takes about 90 minutes.

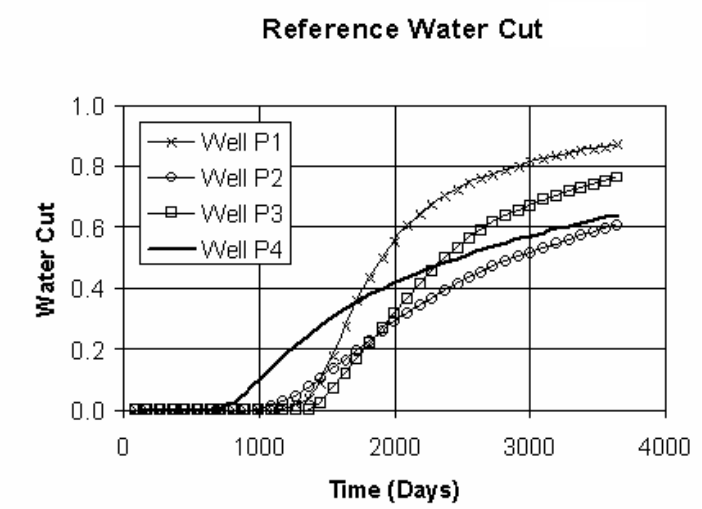

Figure 3.32: Reference production data for four wells.

The locations of the three facies are perturbed to match the water cut. Figure 3.33A shows an unconditional initial realization of the prior model; in this example, no hard conditioning facies data is used. The initial realization is perturbed iteratively with the single parameter method until a reasonable match is achieved. The realization obtained after 15 outer iterations is shown in Figure 3.33B, or a total of 75 flow simulations.

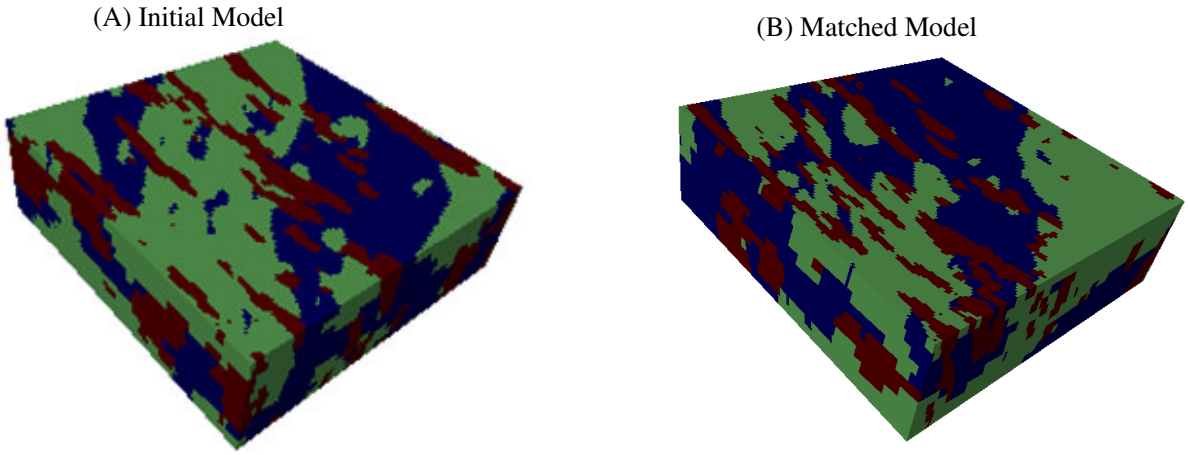

Figure 3.33: Initial and matched realization for single parameter method.

A reasonable match to the water-cut in the four producing wells is obtained, and the match is much improved over the match of the initial model (Figure 3.34).

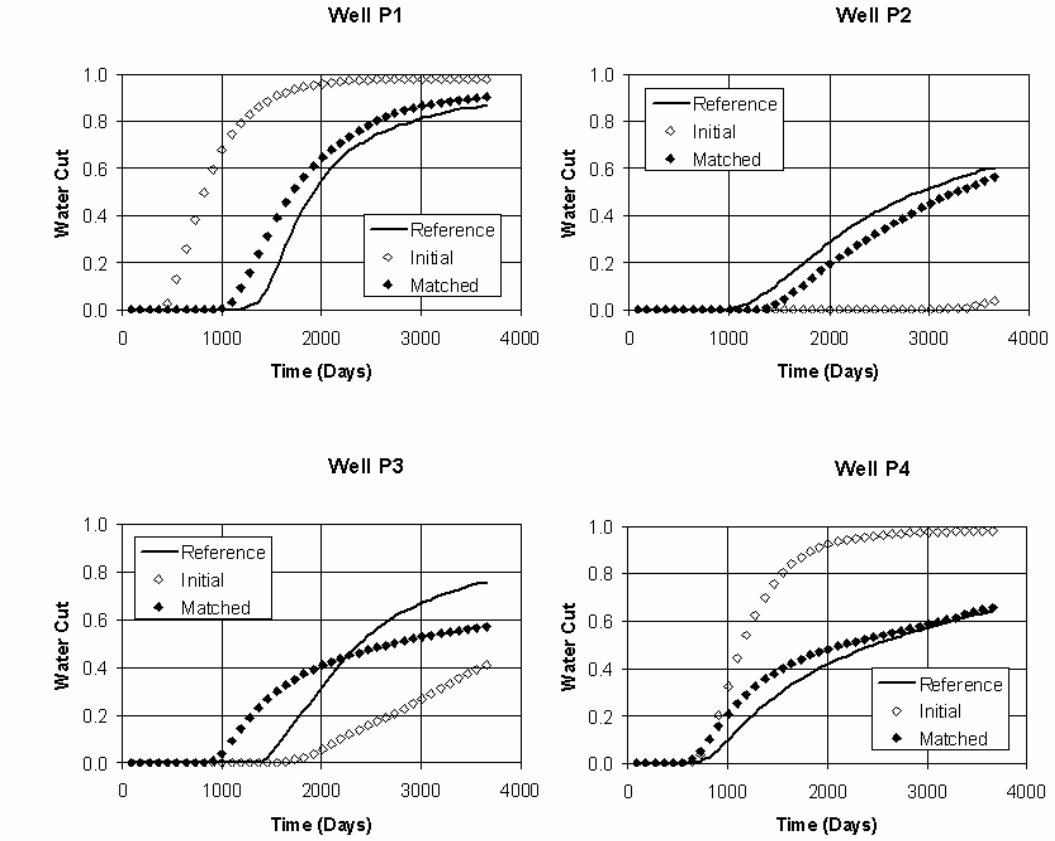

Figure 3.34: Production data from the reference, initial and matched models.

Four rectangular regions are created with the injector at the center (Figure 3.35A). Since flow between each producer and injector takes place mostly in the quadrant occupied by the producer, the facies distribution in each quadrant is mostly consequential to the flow data in that quadrant; hence, efficiency gain is expected over the single parameter method. Indeed, the number of iteration required to reach the same accuracy of the match as in the one-parameter case equals three (21 flow simulations). A single matched model is shown in Figure 3.35C, note that despite the existence of regions, the final model displays the same style of continuity as the reference.

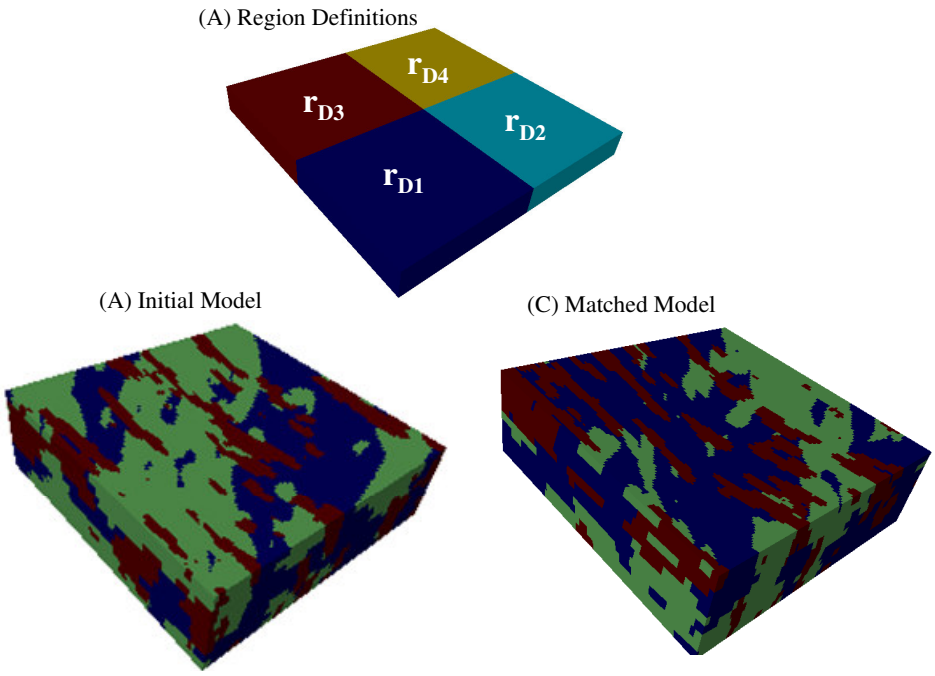

Figure 3.35: Regions, initial and matched models for the multi-parameter method.

Summary

These four examples demonstrate that the PPM is applicable to many different situations. Continuous parameters such as permeability as well as discrete properties such as facies or fractures are perturbed. The method is not limited to two facies given that multiple facies also are modified easily. The PPM is able to match well pressure or rate information. For the multiple parameter case, one well can be assigned to a production unit, or multiple wells can be combined in the same objective. While 3D examples require more time to run the simulator, the probability perturbation method is applied to 3D cases without any difficulty. As the geologic complexity increases, the improvement in efficiency of the MP-PPM increases over the SP-PPM. The improvement ranged from 1.5 times faster for the Similar Facies Distribution Everywhere case to 3.5 times faster for the Multi-Facies 3D example. In the next two chapters, the MP-PM is applied to two different real case studies.

"I am a true believer in luck, I believe the harder you work the more luck you have."

- Thomas Jefferson

Chapter 4

Channel Reservoirs

Static reservoir models are built with a single pass (non-iterative) and honor hard point-location data, soft seismic data and available geologic information. By perturbing properties in the static model, a new reservoir model is created that matches the historical production data. One of the primary advantages of the probability perturbation method over other history matching methods is that any property can be modified to effectively match a reservoir model. This includes large scale properties such as facies that often have a large impact on flow in a reservoir. At the same time, the static data including the geologic information is always maintained. This chapter demonstrates the practicality of the method for one particular type of facies model that is present in many clastic reservoirs: channel systems. Two different reservoirs are examined, a 3D synthetic reservoir and a real offshore African reservoir.

4.1 Stanford V Reservoir

In the examples from the previous chapters, the reference and the history matched realizations are all drawn from the same geological description (e.g. same variogram or training image). This was done for convenience (it is easier to build and history match), and for the previous cases where we are merely demonstrating the method, it is acceptable to have the same geological model. However, it is not realistic in practice to assume that we know perfectly the geological heterogeneity. The Stanford V examples attempts to replicate at least some of aspects of this reality (at a minimum acknowledge it exists). The Stanford V reference has a finer grid and more complex property modeling than the reservoir model we history match and use for prediction.

The example uses a synthetic 3D fluvial channel reservoir model as the reference. The reference model was created by Mao (1999) and is based on a North Sea reservoir. It is made up of three distinct facies: channel sand, crevasse splays, and mudstone. The locations of the facies were determined by using an object-based simulation technique. The petrophysical properties were generated independently for each facies using Sequential Gaussian Simulation. The model has aerial dimensions of 10,000 ft by 13,000 ft and is 450 ft thick, and it is divided into 100 gridblocks in the x-direction, 130 in the y-direction, and 30 in the z-direction. The z-direction has three major horizons each with 10 layers. The three horizons have channels of different size, different directions, and different sinuosity (Figure 4.1). The top horizon has thick channels extending east to west; the second horizon has thinner channels that stretch north-east to south-west, and the third horizon has very thin, highly sinuous channels that run north to south.

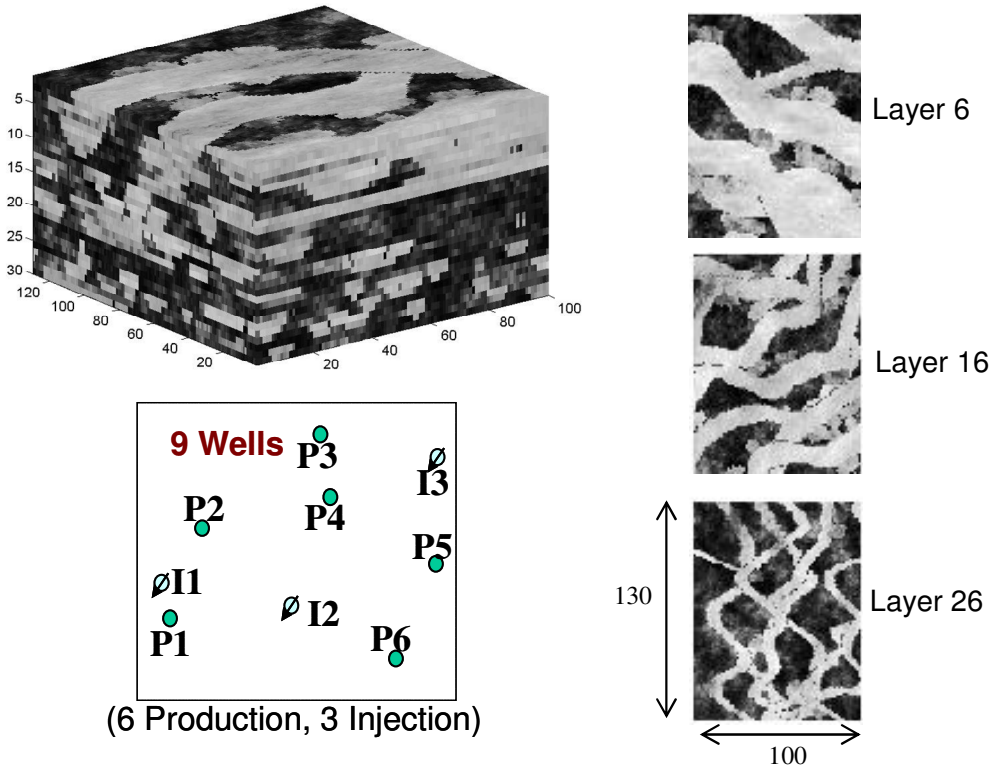

Figure 4.1: Reference model for 3D example. Three layers are shown where channel differences can be distinguished.

In this example nine vertical wells (six production and three injection) were added to the model for flow simulation. The areal locations are provided in Figure 4.1. The completion intervals are not the same for all wells. Production wells 1 and 4 and injection well 3 are completed in the top 20 layers; production wells 3 and 5 and injection well 1 are completed in layers 6 through 25; and production wells 2 and 6 and injection well 2 are completed in layers 21 through 30. The injectors are each injecting 12000 barrels of water per day, and wells 1, 2, 3 and 5 are producing 6400 barrels of liquid per day, and wells 4 and 6 are producing 6200 barrels of liquid per day.

A real-life situation of incomplete information is established in order to generate a reservoir simulation model. A multiple point geostatistical algorithm (snesim) is used to

create the realizations for the model. The geological parameters required for the realizations differ from the reference truth as follows:

- The facies proportions are calculated from the well-logs directly. The proportion of channel sand encountered by the wells is 27% and the proportion of crevasse encountered is 6%. Note that this is different than the reference, which has channel and crevasse proportions equal to 41% and 5%, respectively. The difference in the amount of facies also affects the pore volume. While the reference pore volume is 1.76×10^9 barrels, the generated models are from 1.32×10^9 to 1.44×10^9 barrels.
- Although the reference has different channel shape parameters for different layers, the channel shape parameters for the geostatistical realizations are the same for all layers. The training image used for all of the layers is shown in Figure 4.2. It was generated with an object-based technique and is five layers thick, and each layer is 200 by 200 gridblocks. The proportions for each facies in the training image come directly from the well-log information discussed in the previous bullet.

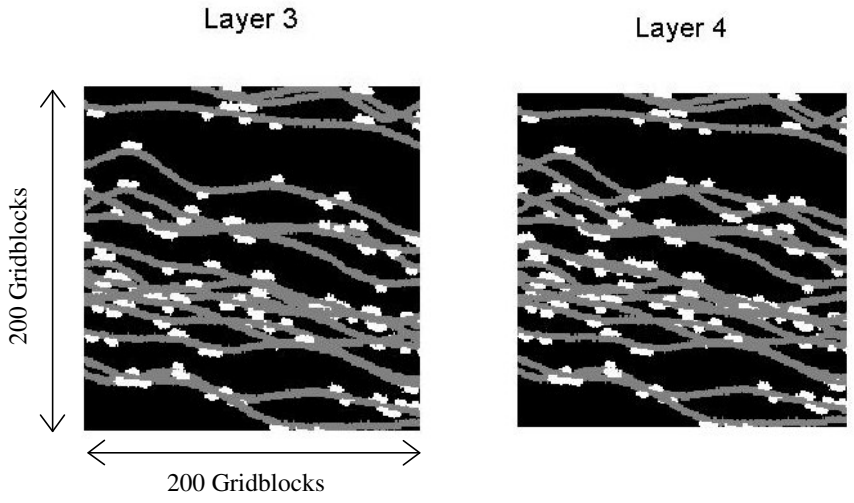

Figure 4.2: Two layers from a five-layer training image

- The porosity and permeability are assumed constant for each facies and obtained by averaging porosity and permeability per facies from wells. Table 4.1 presents the porosity and permeability values for the three facies.

Table 4.1: Porosity and permeability for each facies

	Porosity	Permeability
Channel	29 %	456 md
Crevasse	23 %	253 md
Mudstone	7 %	0.03 md

- The gridding of the model is also different than the reference. There are 50 gridblocks in the x-direction, 65 in the y-direction, and 15 in the z-direction for a total of 48,750 cells. In each direction there are half as many gridblocks as the reference model, but the overall dimensions of the model remain the same. The dimensions of the cells are twice as long in each direction; therefore, the cells are 200 ft by 200 ft and 30 ft thick. The wells will be completed for the same depth interval as the reference model.

A single static realization that is conditioned to the well data is created using these model parameters and is displayed in Figure 4.3.

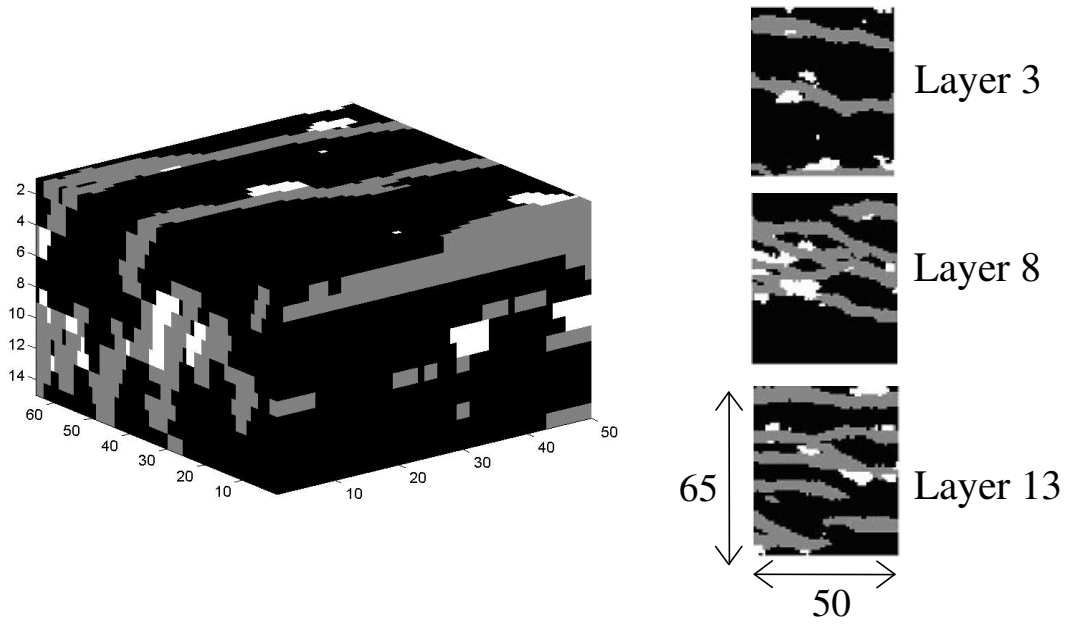

Figure 4.3: Initial realization for 3D example. Three layers show that there are no channel property differences among layers.

The model is a two phase oil-water system. The viscosity of the oil is 1.0 centipoise (cp), and the water is 0.325 cp. The density of the oil is 49.94 lb/ft^3 and the water density equals 62.48 lb/ft^3. The oil-water relative permeability curves are provided in Figure 4.4.

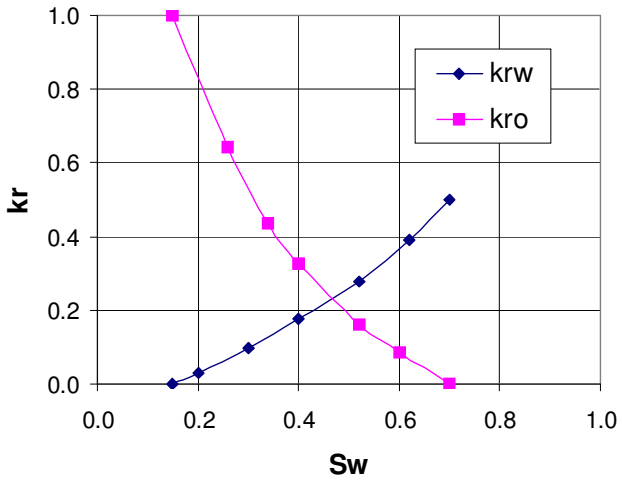

Figure 4.4: Relative permeability for Stanford V example.

In this example the SP probability perturbation method is compared with the MP probability perturbation method. Hence, the matter of how to fill the reservoir grid with r_{Dk} values needs to be addressed. Three different methods are compared for this example. The first method uses streamlined defined regions. This technique to define region geometries has not yet been used in the dissertation, so a short description is provided on why the regions are utilized and how the regions are determined.

As history matching takes place, it might be desirable to change or update the region geometry definition based on the modifications made in the reservoir model. Such dynamic region definition may be achieved using streamlines. Streamlines are well suited for the job because they show the direct paths by which fluid will enter a production well. These paths identify the gridblocks that, if changed, will have a direct affect on production from that well. Furthermore, streamlines have been used successfully to define regions in traditional history matching schemes (Milliken *et al.,* 2001).

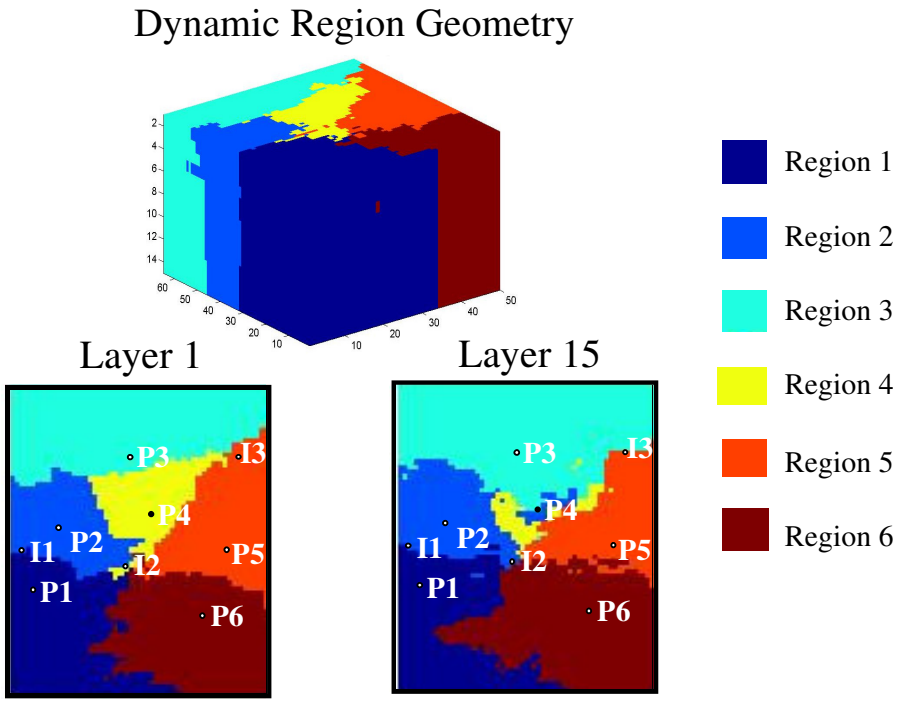

Figure 4.5: Dynamic region geometry for one iteration of one realization defined from streamlines.

Every production well has a set of streamlines entering that well. All blocks hit by this set of streamlines define the "drainage zone" for that well. The various drainage zones define the geometry of the regions used for history matching. As the model is perturbed, the facies geometry changes; consequently, the drainage area of the production wells also changes. The region definitions for one iteration of one realization are displayed in Figure 4.5.

Notice that region 4 is hardly represented in the lower layers and recall that Well 4 was only completed in the top 10 layers. Few streamlines that ended at Well 4 traveled in layers lower than it was completed in. Well 1 was also completed in layers 1-10, but its region does include area in the lower layers. This shows that the regions can be defined differently for different areas; thus, the necessity for the streamline based approach is confirmed.

To investigate the importance of using streamlines to define regions, an alternative "fixed" region definition is used. The fixed regions are defined by the set of gridblocks that are closest to a production well (Voronoi polygons). Figure 4.6 displays the shape of these regions, and while the regions are in the same general locations as the streamline defined regions in Figure 4.5, there are some differences. The key point is that these regions are fixed for whatever facies geometry is present in the reservoir.

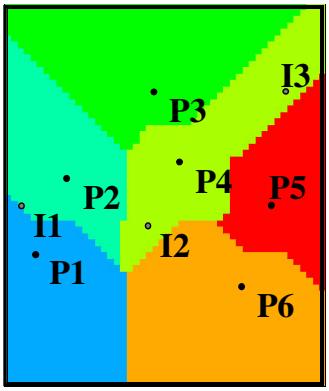

Figure 4.6: Static region geometry defined by Voronoi polygons.

Another method to establish the r_{Dk} values at all gridblocks uses kriging. The r_{Dk} values are interpolated from the six values calculated at the production well locations, and they are displayed for two layers of the model in Figure 4.7. The variogram used to perform kriging is also in Figure 4.7. The isotropic range has a very low nugget and it is roughly equal to the well spacing. This causes a smooth transition of r_{Dk} values throughout the reservoir.

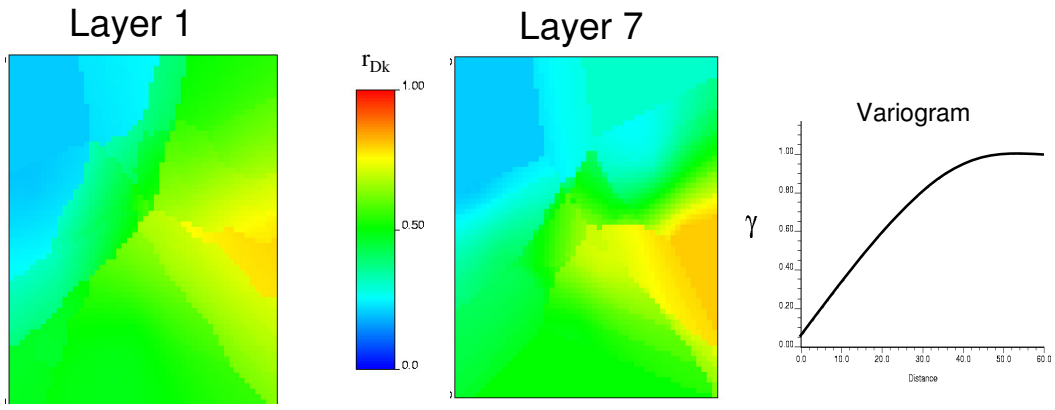

Figure 4.7: Kriging defined r_{Dk} values for two layers, and the variogram used to generate the values.

The reservoir simulation is performed using eclipse (Schlumberger, 2003), and the streamline simulator used is 3DSL, (StreamSim, 2002). Water cut information at the six production wells is the data used for history matching. Ten realizations are completed for all cases, and all realizations for the three cases are converged to the same tolerance. A history-matched model is displayed in Figure 4.8.

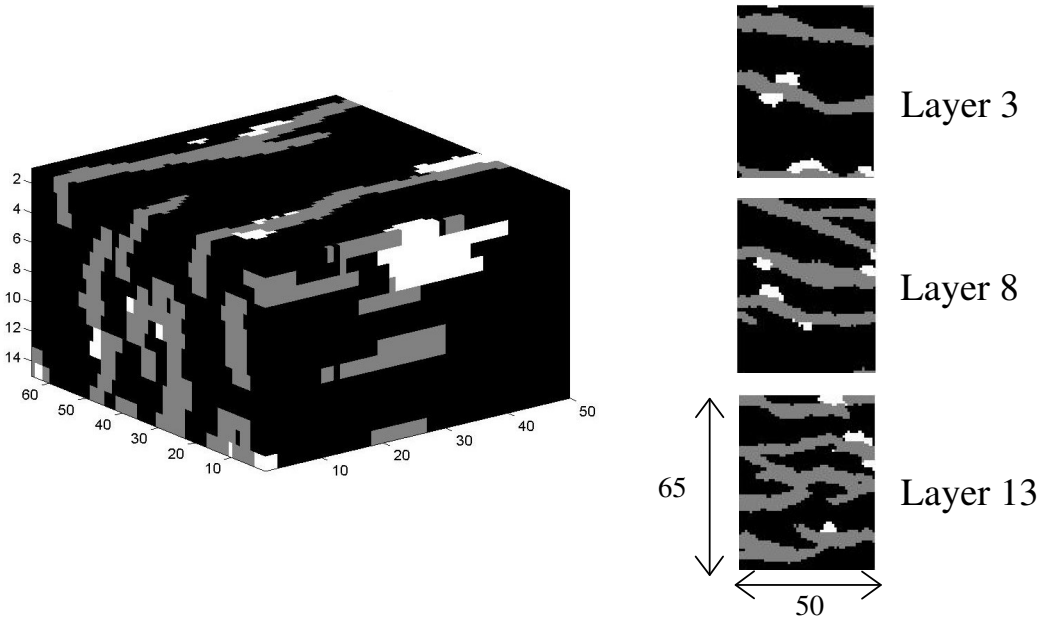

Figure 4.8: History matched realization for 3D example.

Although the generated models have different channel shapes than the reference, history matching is consistently achieved for all cases. Figure 4.9 shows the improvement of the history match for one realization.

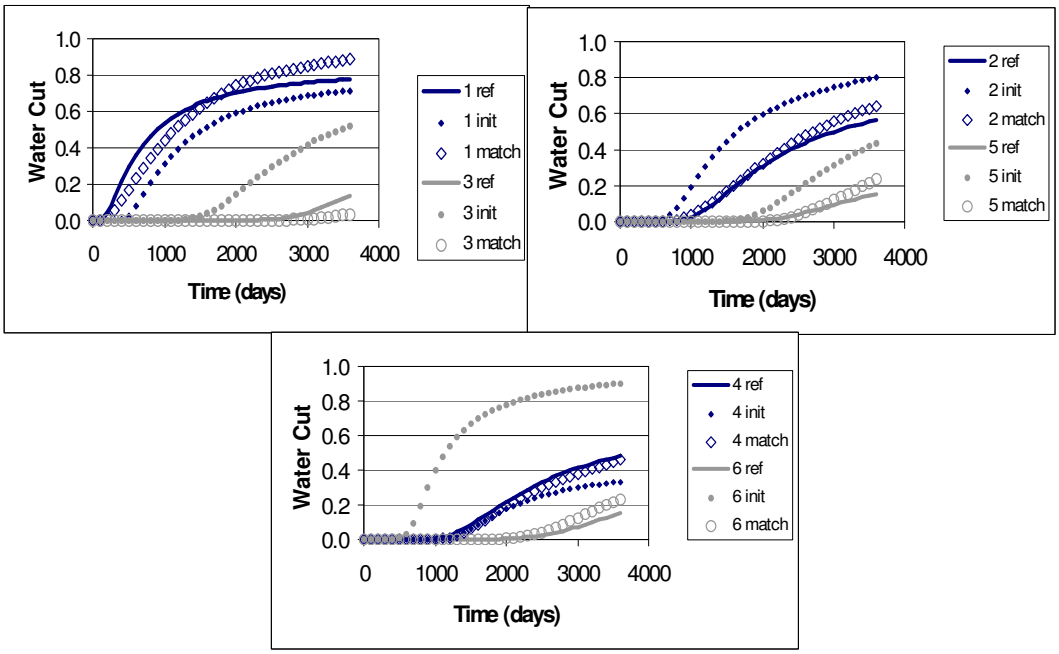

Figure 4.9: Water cut reference, initial match and final match of six production wells for one realization.

The reference data (solid line) along with the initial (filled symbols) and matched data (open symbols) are displayed for the six wells (Figure 4.9). Exact matches are not achieved, but significant improvement is observed in all wells, some with almost perfect match. The matches obtained would be adequate for most true field applications.

Table 4.2: Number of outer iterations for 3D model. (S) stands for streamline defined regions, and (V) stands for voronoi polygon defined regions and (K) stands for kriging.

	Number of outer iterations to converge										Average	Average
Realization	1	2	3	4	5	6	7	8	9	10	Iteration	Flow Sim.
MP (S)	5	3	1	1	4	1	1	4	4	3	2.7	16.1
MP (K)	8	2	4	1	4	8	1	5	3	8	4.4	26.3
MP (V)	11	4	5	13	2	9	4	4	3	2	5.7	34.0
SP	8	4	4	10	4	8	7	1	9	6	6.1	36.5

Table 4.2 compares the efficiency of the four different cases where the cases differ in how the r_{Dk} values are defined. MP (S) has streamline based regions that change (once per outer iteration) as the history matching process continues. It is the most efficient converging in about half as many outer iterations as the other methods. Although the method that uses Voronoi polygons to define regions, MP (V), perturbs multiple regions separately, its convergence efficiency is about the same as the SP method. This is because the region definitions for MP (V) are not optimal for this problem; however, even for this simple region geometry, the multiple region method performs no worse than the one region method. Using kriging to define the r_{Dk} values improves the efficiency over the Voronoi region method, but it is not as good as the streamline defined regions. For types of models where water is injected and water cut is to be matched, streamline defined regions seem to be the best regions.

Traditional history matching method comparison and prediction

The Stanford V example is also history matched with a more traditional history matching method. The method is from a commercial history matching software package and is termed SimOpt (Schlumberger, 2003). SimOpt used sensitivity coefficients to determine the region definitions and the amount of change required in the model.

The regions for the modification were determined using an internal program in SimOpt called Gradzone. An initial run is completed to calculate sensitivity coefficients for every gridblock and the blocks with similar properties are grouped together into regions. For this example, four regions are selected to be perturbed in the history matching process. The regions selected for Layer 10 are depicted in Figure 4.10. The white gridblocks are where no perturbation is done.

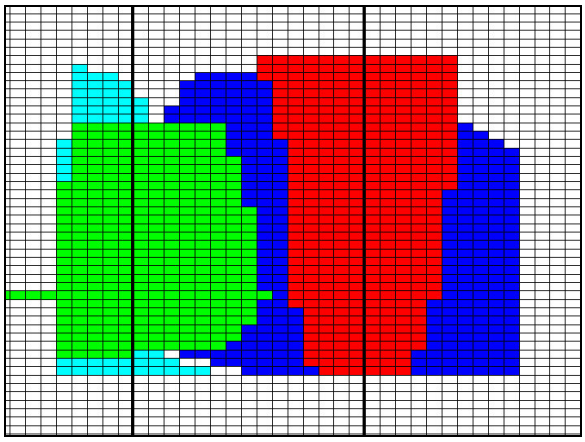

Figure 4.10: Four regions selected to be perturbed for history matching for Layer 10.

In the SimOpt method, the horizontal permeability is multiplied by a different constant in each region until a history match is achieved. The method was able to achieve a high quality history match for this field with only eight flow simulations, so it is a very fast method (compared to 16 flow simulations for the fastest PPM); however, the geological continuity of the SimOpt model is not maintained (Figure 4.12). Figure 4.11 displays the production profile for true data, the history matched model and the initial model for the six production wells, and the following paragraph discusses why geology is not maintained.

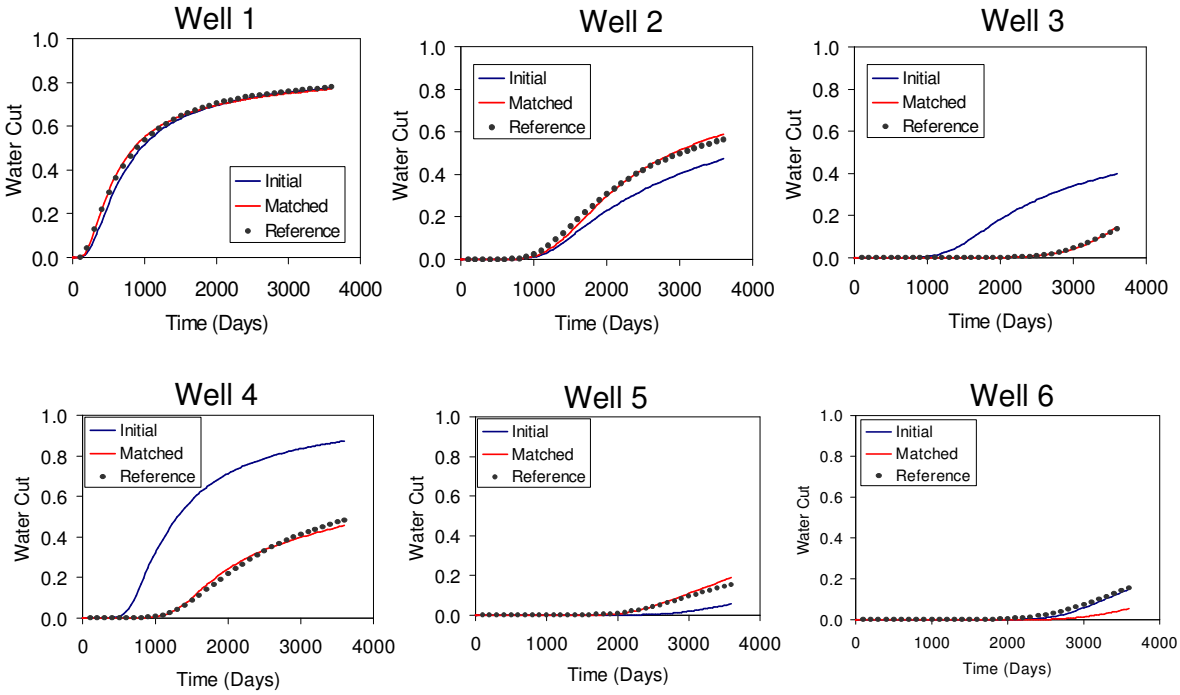

Figure 4.11: Initial, matched and reference water cut data for six production wells when SimOpt is used for history matching.

Wells P3 and P4 are significantly over-predicting water cut in the initial model. To remedy the over-prediction, a very low number, 0.05, multiplies permeability in the central region (red region of Figure 4.10) of the model. In the right map of Figure 4.12, this can be visualized where the channels have been virtually eliminated and appear blue. This reduces the connection between the injector I2 and the producers P3 and P4. The black region in Figure 4.12 corresponds to values off the scale and is approximately 3500 md. The increase adds connection between I2 and P2 where the initial model was under-predicting the observed data.

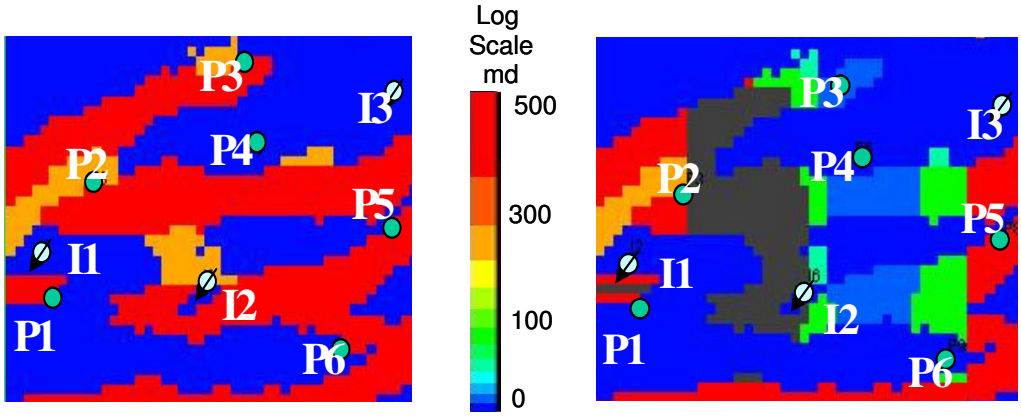

Figure 4.12: Initial (left) and matched (right) permeability profile for Layer 10.

The matched model of Figure 4.12 clearly demonstrates that the channel nature of the reservoir is no longer honored in the middle of the model. While the history match is very good, the matched model may not be representative of the true reservoir. Consequently, the prediction from the model may be less accurate than a matched model that honors both the production data and the channel nature of the reservoir.

To compare prediction from the SimOpt method with the PPM method, the reference and history matched models are produced for an additional ten years. The PPM model used is the first history matched realization obtained from the streamline defined regions. In the prediction runs we also added three new wells, one injector and two producers. The prediction response for three of the original wells and the two new production wells is provided in Figure 4.12Figure 4.13.

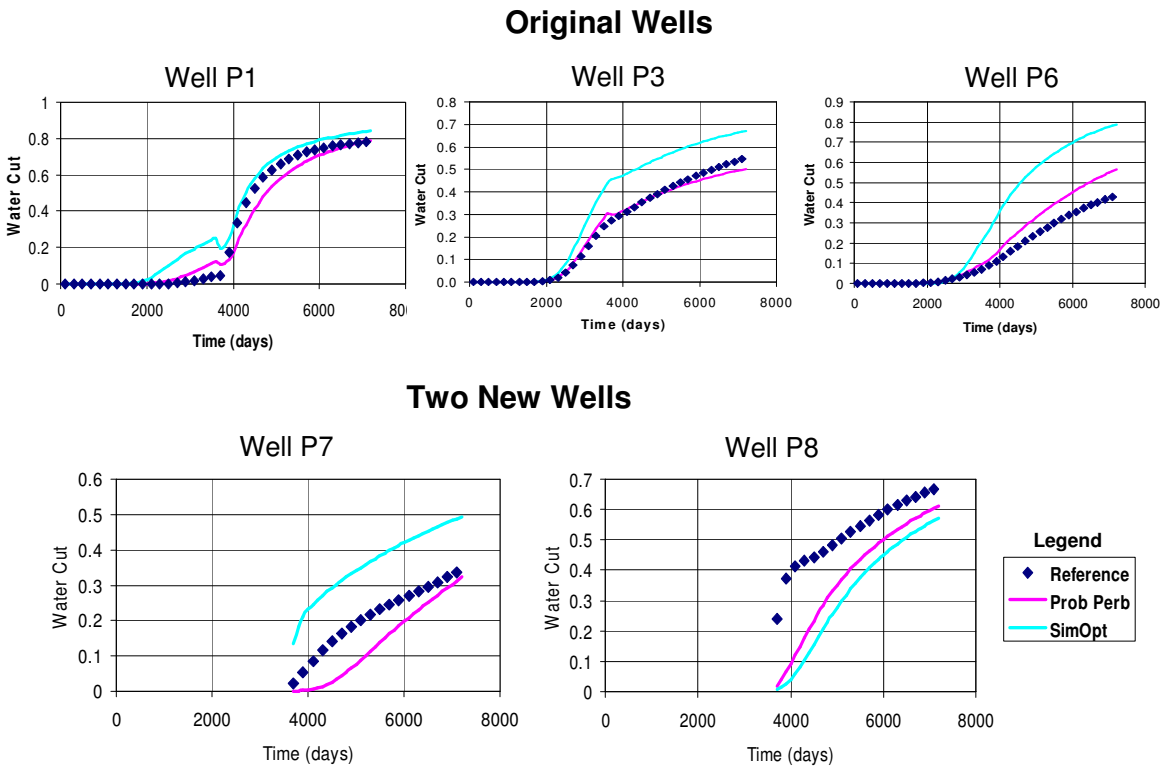

Figure 4.13: Prediction for three original wells and two new wells.

Some of the original wells predict quite well for both the SimOpt model and the PPM model (Well P1), but the PPM model performs much better for other wells (Wells P3 and P6). Both new wells have a better match for the PPM model, although for Well P8, neither prediction is very good. The prediction error (average least square error for the prediction part of all eight wells) equals 0.21 for the SimOpt model and 0.10 for the PPM model. The PPM model predicts two times better than the SimOpt model because one piece of information, geologic continuity, is represented in the PPM model but not in the SimOpt model.

4.2 West Coast African Reservoir

The West Coast African (WCA) reservoir is the first real reservoir to be history matched in this dissertation. History matching actual reservoirs is much more complicated than synthetic cases. Although not directly related to the probability perturbation method, many aspects of history matching real reservoirs need to be understood, and the first part of this section describes the additional work that is required for real reservoirs. Next, the geological description of the WCA reservoir is provided, and the procedure to build realizations that are conditioned to the static data is explained. Similar to the Stanford V reservoir, the WCA reservoir will be history matched by perturbing channel locations. Two different geologic scenarios are used, and multiple history matched models are generated for each scenario. The reservoir has been producing for almost five years, but the models are matched only to the first three years of production. Afterwards, the matched models are run for the last one and a half years of data, and the results compare favorably with the field data.

History Matching Real Reservoirs

In real reservoirs, the production data does not have the smooth curves that are typically found in synthetic data. Real data usually has a significant amount of noise; there may be short or long shut-in periods, and there may be errors in the data especially for the non-sales fluids. Hopefully, these errors are detected and corrected when the data is quality checked prior to the simulation study, and we assume that any remaining errors in the data are insignificant.

To handle the noise in the data, one must decide how often to change the well constraints in the simulation model. Should all the peaks and valleys be matched, or should the data be averaged over some time interval? Most current commercial simulators accept as frequent changes as desired, so all the variation could be captured. However, every change reduces the efficiency of the simulation, and may not be relevant for modeling the reservoir properties. Data is usually averaged over monthly or quarterly intervals, and

the simulation model is updated that frequently. This is in contrast to synthetic data where constraints are usually constant for the entire study.

Additionally, one must also decide how the wells in the simulation model are constrained to the production data. Each well is usually constrained to a fluid rate, and this rate will be honored exactly in the model output. The wells can be constrained to oil rate, total liquid (oil and water) rate, or total fluid (oil, water, and gas) rate. Oil rate is almost always the most accurately recorded data, and one may want to exactly honor this data. However, honoring the total liquid or total fluid rate also has advantages. If the model is producing too much water, then the model will be producing too little oil (or gas); thus, the correct volume of fluid is always extracted from the reservoir model. This also helps the pressure match in the reservoir model to be maintained.

Besides fluid rates, a minimum bottom hole (BHP) pressure constraint must be specified. Once the well pressure equals the minimum BHP, the well no longer will meet its target rate. The minimum BHP can usually be obtained from field information as the minimum pressure that the wells will produce. A primary BHP constraint could also be used, but this is less common in real field applications.

Once we decide how to constrain the model, we need to determine what data or combination of data to match. Should water rates or water cuts be matched? Is produced gas included? Do we match pressure information? The answers depends on what the history matched model is used for and what type of data is available. Typically, we would want to match both pressure data and some rate information (e.g water cuts and GOR or water and gas rates). Matching the total field rates is important, but if the model is going to be used for any local predictions, the main focus of the history matching endeavor should be to match the individual well data.

For synthetic cases, the matching tolerance is clearly defined, and once the mismatch is below the tolerance, the model is considered “matched”. For real cases, the mismatch tolerance or stopping criteria cannot be defined that simple. Real data is difficult to

match and there is usually a deadline for completion, so the matching criterion often becomes "as good as one can get it in the time that one has." Although a mismatch may be calculated, it is usually not appropriate to specify some tolerance where the reservoir is "matched". Visual comparisons often provide the most appropriate stopping criteria.

Choosing the constraint type or the matching data generally does not determine *if* the model can be history matched. However, choosing which properties to modify will have a major impact on the success of the history matching endeavor. The selection procedure is not trivial since a large number of properties may affect production data and since all of the properties have some amount of uncertainty. To achieve a match, the properties with the most significant impact on the production (within their range of uncertainty) need to be perturbed. Every reservoir is different and these parameters usually are not known when the history matching process is started. Therefore, a sensitivity study must be performed prior to the history match to determine which parameters are to be perturbed.

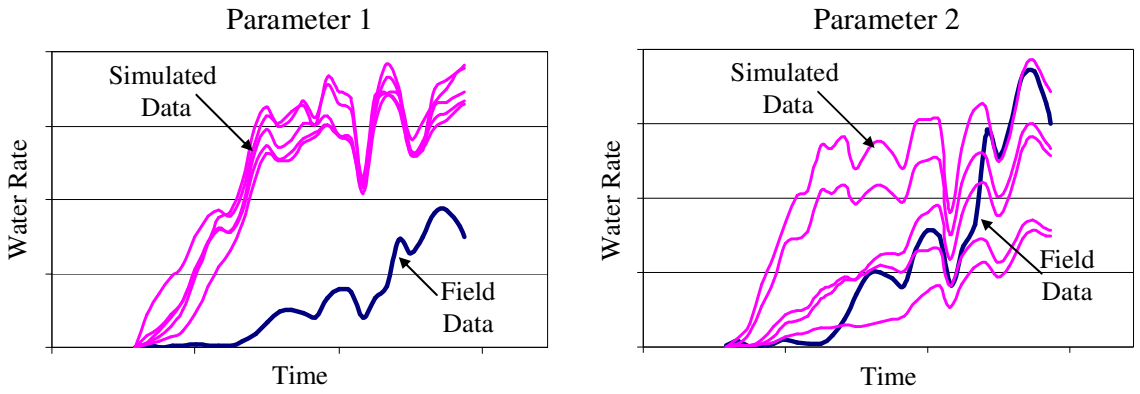

Figure 4.14: Sensitivity of two parameters and their affect on production data

For example, Figure 4.14 shows the influence of two different parameters on water production. This is simply a descriptive example, and these could be any two parameters from a reservoir model. For the parameter on the left, there is a relatively small impact on production for its range of uncertainty; conversely, the parameter on the right has a

significant impact when varied. By perturbing the parameter on the right, the water rate for this well can be matched; however, perturbing only the parameter on the left will never result in a history match regardless of the number of iterations. There are tools available to help select the parameters with the most impact. One method used in industry consists of setting up a list of potential parameters, assessing the range of uncertainty about each parameter, then using experimental design (Damsleth *et al.*, 1992) to flow simulate a few models and identify the most sensitive parameters. Another common method is to use sensitivity coefficients (Schlumberger, 2003) to determine the parameters that are most important; however, neither of these tools were used in the current work. We simply completed a number of flow simulations for each parameter, and then manually compared the various responses.

During this "pre-history matching" process, flow simulations are completed to test the sensitivity of various properties; however, in many instances, the current parameterization is ineffective. Regardless of what parameters are perturbed, a history match can not be achieved. This "ineffective parameterization" could be caused by an actual aspect of the model such as an existing fault not being modeled or an aquifer not being included or included incorrectly in the flow simulation model. Additionally, it could be something more intangible such as the range of uncertainty for a property is larger that initially thought. To achieve a match in these cases, the engineer must discover what part of the model is not parameterized correctly. This is one of the most challenging features of history matching, and it is not restricted to any one history matching technique or algorithm. To find the appropriate element for the model, the engineer must revisit the various data and communicate with other engineers and geoscientists working on the field.

The information from this section (e.g. searching for missing elements and defining what properties to perturb and what data to match) can be combined with other aspects of history matching that were discussed in previous sections with regard to the probability

perturbation method (e.g. defining the static model and generating an initial realization, calculating an objective function, and perturbing properties in the model) to be able to history match real reservoirs with the probability perturbation method.

Description of WCA Reservoir

The WCA reservoir is a deep-water depositional system that occurs on the steepest part of the slope (Figure 4.15). This is a high energy location where steep canyons cuts are formed. Petroleum reservoirs are found here less often than further down the slope where it is less inclined, and when reservoir are found in these areas, they tend to have a more heterogeneous nature.

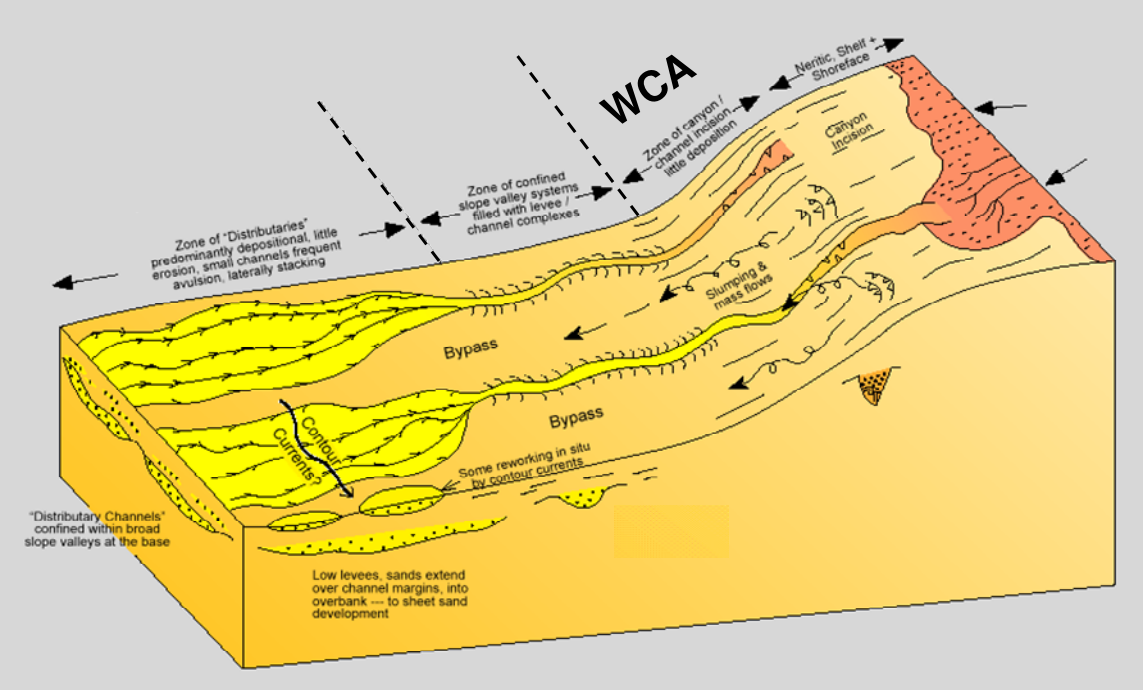

Figure 4.15: Geologic schematic describing zone of canyon where WCA reservoir is.

The WCA reservoir is an amalgamated and somewhat aggradational channel complex that fills the canyon cut. The uppermost channels are more aggradational and have definite recognizable meander patterns from seismic data (Figure 4.16).

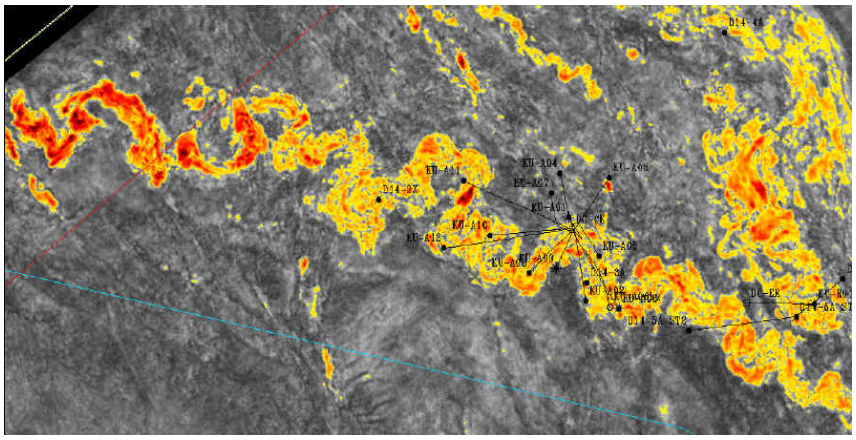

Figure 4.16: Interpreted seismic map of upper layer in WCA reservoir. Meandering channels are clearly visible.

Vertically there are three massive bodies where many individual channels meander within each body (Figure 4.17A). Although a shale barrier divides the A and B geobodies, tortuous channels have been mapped between them, and pressure data indicates some communication between the two bodies. Therefore, the A and B bodies are modeled as a single unit (Figure 4.17B). Geobody CD is not in communication with AB and is modeled as a separate unit.

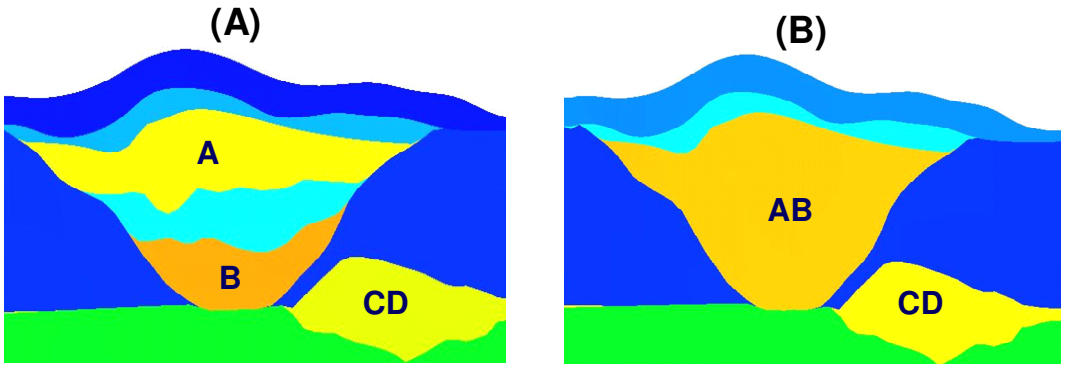

Figure 4.17: A depiction of the geobody definition in a cross-section conceptual model of the WCA reservoir.

The reservoir is located offshore in 1600 feet of water and is 4600 feet below sea level. The reservoir is dipping slightly from east to west. A number of faults compartmentalize the reservoir into four areal segments: CW, CC, CD, and CE (Figure 4.18). The three segments CW, CC, and CE make up the massive AB geobody. Although the faults reduce transmissibility, the three segments appear from production data to be in communication with one another.

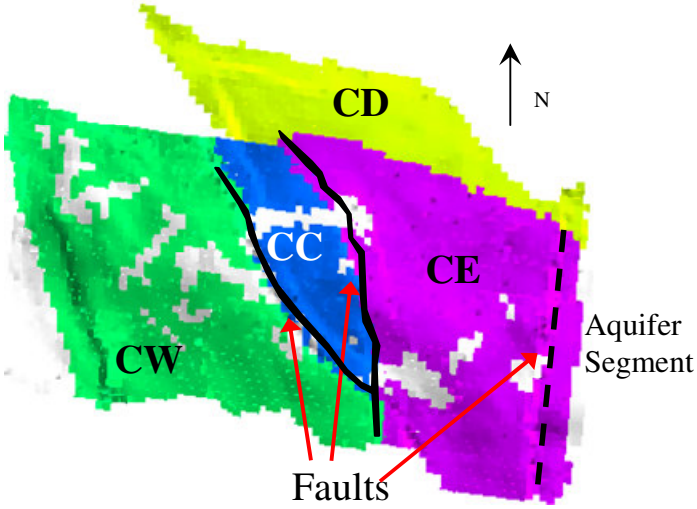

Figure 4.18: WCA reservoir model depicting four segments and three faults.

A fifth "aquifer" segment exists east of the reservoir, and a fault separates it from the rest of the reservoir. The fault is believed to be sealing at initial conditions, but during reservoir production, the fault appears to allow water to cross into the reservoir.

Water support also comes from below the reservoir. The water-oil contact varies for the different segments, but it is located around 5440 feet, so the maximum gross thickness of the WCA reservoir is close to 800 feet.

There are three sand facies of varying quality in the reservoir and a fourth facies that represents the shale or non-pay part of the reservoir. The amount of sand is about 55% of the gross volume. The three sand facies are classified by their petrophysical properties

measured at the wells. The highest quality facies is the channel facies with an average permeability around 2 Darcys. The description of the other two sand facies is less clear. They have been interpreted as either debris flow lobes, overbank levees or poorer quality channels. These two facies each account for 10% to 15% of the total reservoir. The high quality channel represents about 30% of the reservoir.

In some of the upper layers, seismic data is able to position accurately the channels (Figure 4.16), but in most layers the locations of the channel sands are not well known. The placement of the channels in the reservoir model has a substantial impact on production. The primary purpose of this study is to determine the location of the channels that provide the best match to production data.

The reservoir is intercepted by 28 wells. There are 20 production wells and eight water injection wells. The locations where the wells intersect the top of the reservoir are displayed in Figure 4.19. The production wells have ‘R’ and ‘S’ prefixes, and the injection wells begin with ‘W’. The ‘S’ wells began producing approximately 1 ½ years after most of the ‘R’ wells came on line.

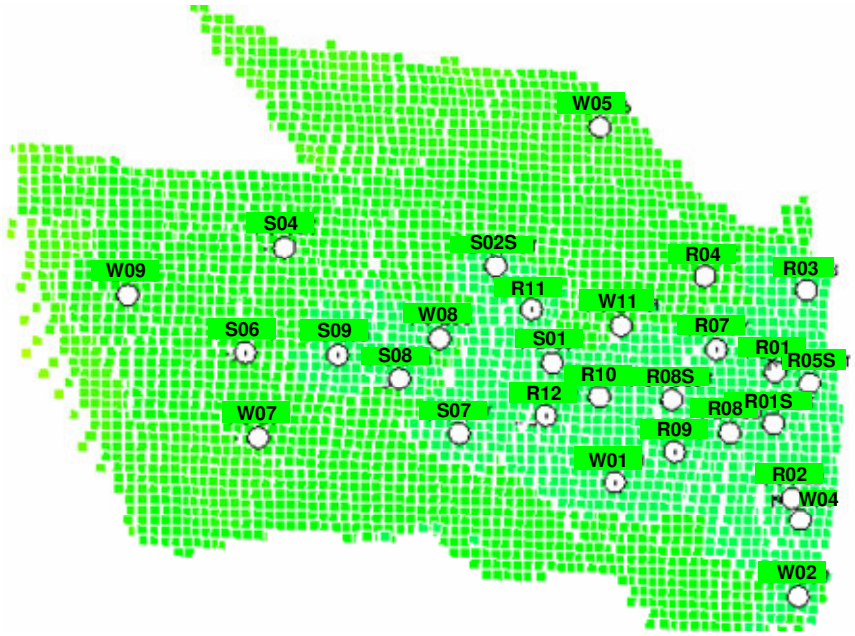

Figure 4.19: Locations of 28 wells. 'S--' and 'R--' are production wells and 'W--' are water injection wells.

A number of seismic attributes are utilized to estimate a value for the shale volume fraction (v_{shale}) throughout the model area. In general this has done an excellent job of capturing sand and shale distribution. Well core and log measurements are used to determine the four facies indicators along the well bore. Additionally, permeability and porosity data from core and logs is available for the four facies (Figure 4.20). For all four facies, the majority of the porosity values are between 0.25 and 0.30. However, the permeability varies considerably for the four facies

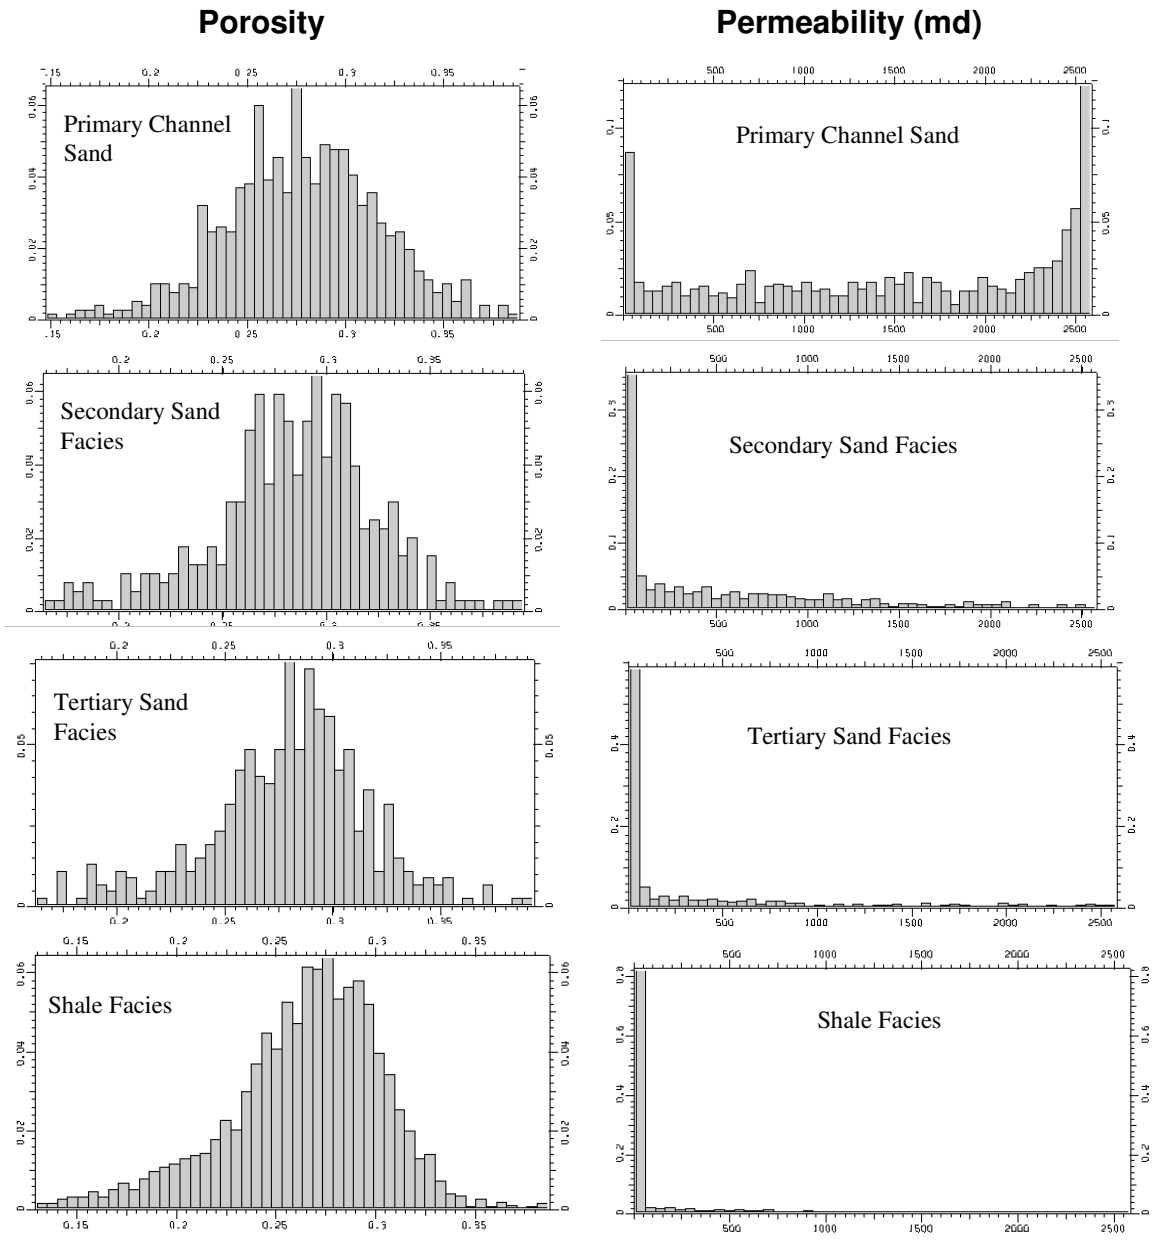

Figure 4.20: Histograms of well porosity and permeability for the four facies measured at well locations.

Since the porosity distribution for all facies is about the same, the porosity does not differentiate between high and low permeability values throughout its range. However,

the relationship between v_{shale} and core permeability is somewhat better in that most of the high permeabilities fall in the low v_{shale} range and most of the low permeabilities fall in the high v_{shale} range.

The Reservoir Model

A reservoir simulation model had been generated previous to the current study; however, this particular model had not been history matched at the well level. The model is termed the "preliminary" model, and it provides the foundation for the models used in this study. The structure of the model is provided by the geological description provided in the previous section: 2 geobodies, 3 faults, and 4 segments. The model is approximately one mile long, one-half of a mile wide and 800 feet thick. The reservoir model has 78 by 59 by 116 gridblocks, but depending on the realization, there are only around 100,000 active gridblocks. The active gridblocks for one realization are displayed in Figure 4.21.

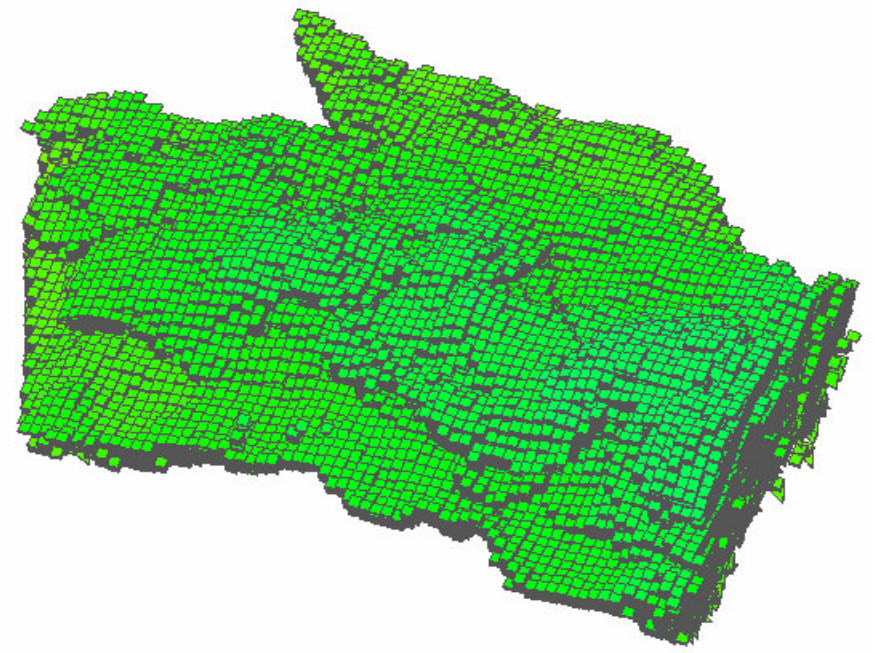

Figure 4.21: Active gridblocks in one of the WCA reservoir models.

The gridblock porosity and permeability is calculated in a three step process. First the facies geometry is simulated using a multiple point geostatistics technique. The facies are conditioned to the available static data (training image, well indicator "hard" data, and for some models seismic data), and during the MP-PPM, the facies are also conditioned to the production derived soft data, P(A|D). The two different geobodies (AB and CD) are simulated independently. Then, four separate v_{shale} simulations are completed (one for each of the four facies) using sequential Gaussian simulation. The values are drawn from the respective v_{shale} histograms for each facies and conditioned to well log data. Finally, the porosity and permeability are determined using co-simulation for each facies. The collocated soft data is the v_{shale} data, and the results are conditioned to the hard well data and the histograms. The facies locations as well as the porosity and permeability for two layers are presented in Figure 4.22.

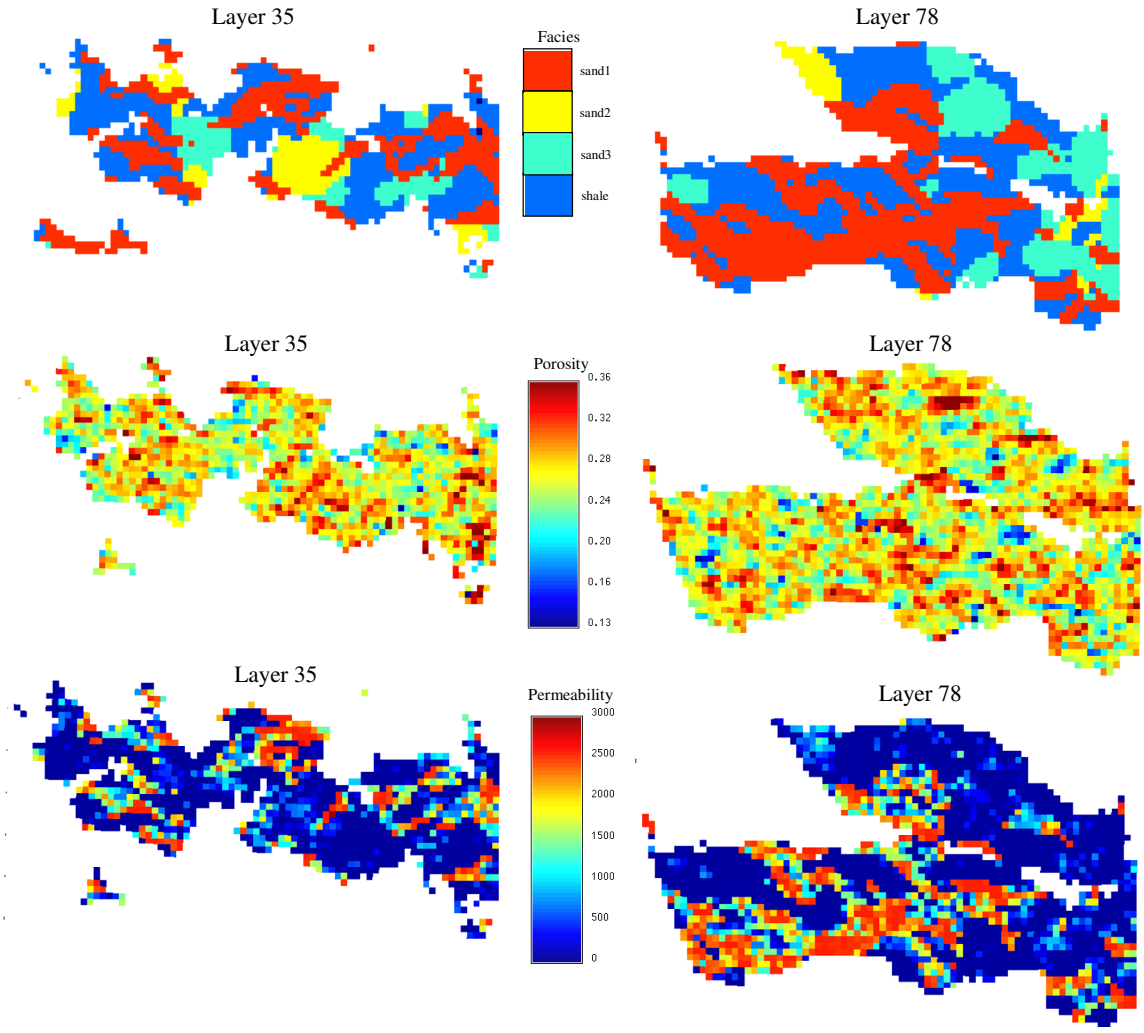

Figure 4.22: Porosity and permeability for two layers of the reservoir model.

The sand and shale facies can be distinguished in the permeability maps, but the porosity distributions do not differentiate sand from shale. These petrophysical properties are calculated using the geostatistical algorithm sequential Gaussian simulation from the reservoir modeling software Gocad (Earth Decision Sciences, 2003); the multiple point geostatistical simulation of facies uses an additional plug-in to Gocad that has been developed by the ChevronTexaco Technology Company.

In the preliminary model, the gridblocks with shale facies were set to inactive status. This is reasonable because the shale contributes very little to flow in the reservoir; however, in some models this caused isolated gridblocks and numerical convergence problems. We reinstated the shale gridblocks to active status but assigned a low permeability value (10 md), so they would not contribute to flow. The shale porosity was taken from the simulated shale porosity.

The porosity and permeability data are exported to the flow simulator (CHEARS). Additionally, other parameters such as fluid properties, relative permeability and initialization are defined in the reservoir simulation model and remain fixed. The initial fluid contacts, pressures and oil properties vary for the five different segments, but the water-oil contact is about 5440 feet, the initial pressure is around 2300 psi, and the oil is around 24 °API. The relative permeability for the reservoir is provided in Figure 4.23.

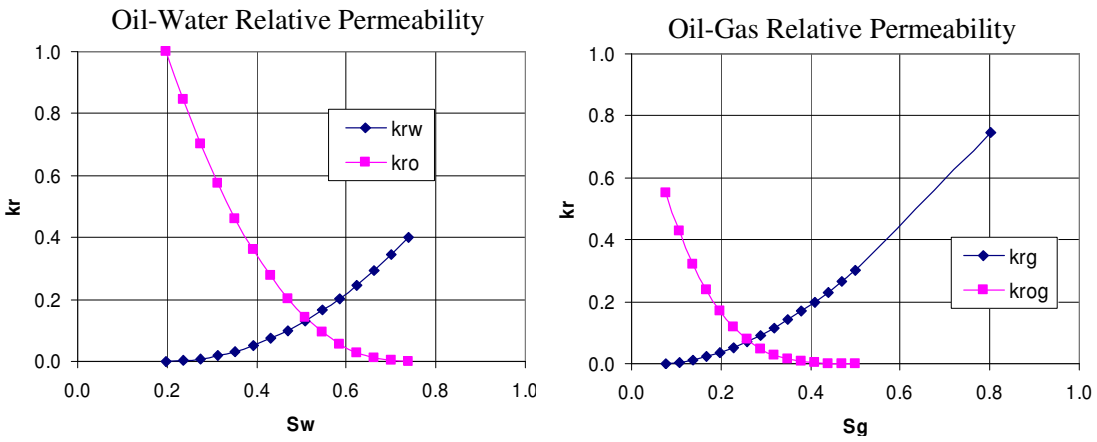

Figure 4.23: Relative permeability for oil-water and oil-gas behavior.

Three major faults are present in the reservoir model. During the pre-history matching (sensitivity analysis) process, various fault multipliers were evaluated, but nothing consistently and significantly affected the match. Thus, during the actual history matching phase, no fault multipliers were used.

The preliminary model does not include an aquifer on the east end of the reservoir. Initially the east end fault was considered completing sealing, but during the pre-history matching process, it became apparent that the fault must be allowing water to flow across it. Figure 4.24B shows the water rate data for the field (blue) and five different realizations (pink) that do not have the east end aquifer included. The well in Figure 4.24B is representative for all the wells circled in Figure 4.24A. None of these wells produce as much water in the simulated models as the field data indicates. Therefore the aquifer was added to the reservoir model. Figure 4.24C shows the water rates for the same well when the aquifer element is included in the models. For this particular reservoir model, the east end aquifer was the one “ineffective parameterization” of the preliminary reservoir model, and it had to be added so the model accurately represented the true reservoir response.

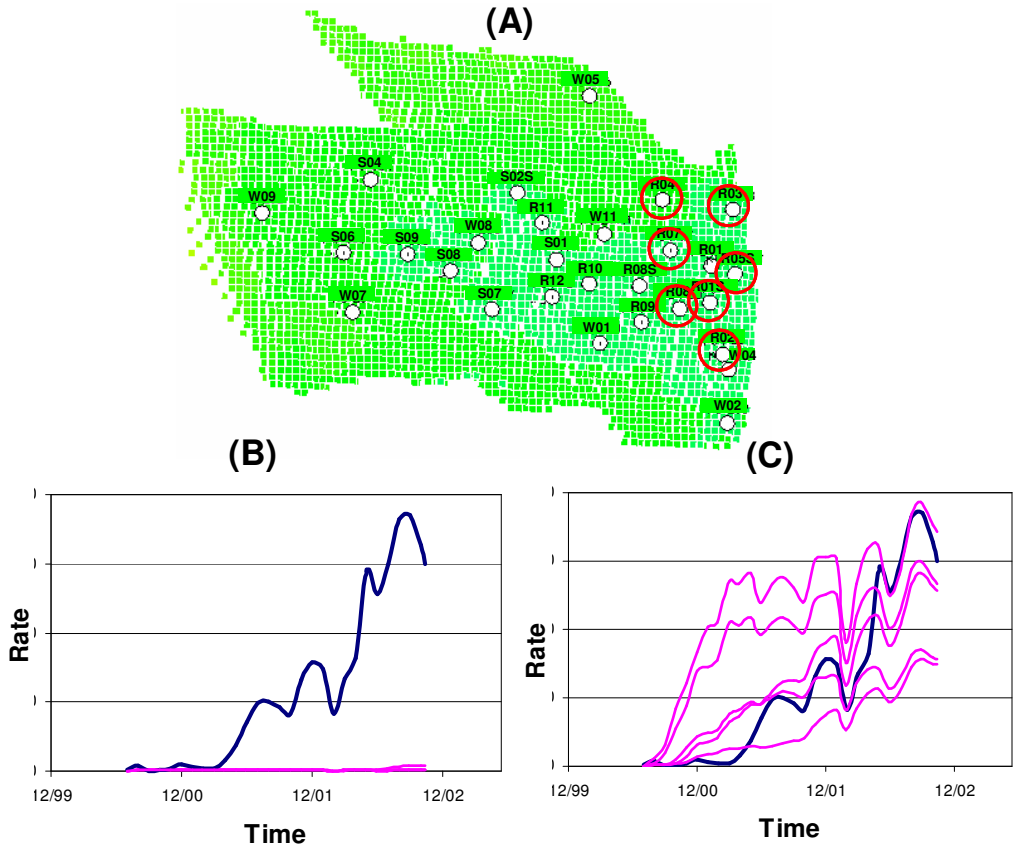

Figure 4.24: (A) Wells circled in red do not produce enough water in the simulation models, (B) water rates from field data and five simulated models that do not have the east end aquifer, and (C) water rates from field data and five simulated models that do have the east end aquifer.

History Matching with the MP-PPM

The MP-PPM has been well developed with the previous examples, and by this point, the most advantageous aspects of the MP-PPM have been established. Therefore, no comparisons of the MP-PPM parameters are made in this example; simply, the aspects that were found to be the best in the previous examples are used.

In this example, there are 28 wells; and 20 of them are producers. Only water rate data is included in the objective function. Since the well production from the model is

constrained to total fluid produced, including other rates information brings little additional information. Pressure data from the reservoir were not readily available, so they were not included in the objective function; however, the matched models are manually examined to ensure the reservoir pressures are realistic.

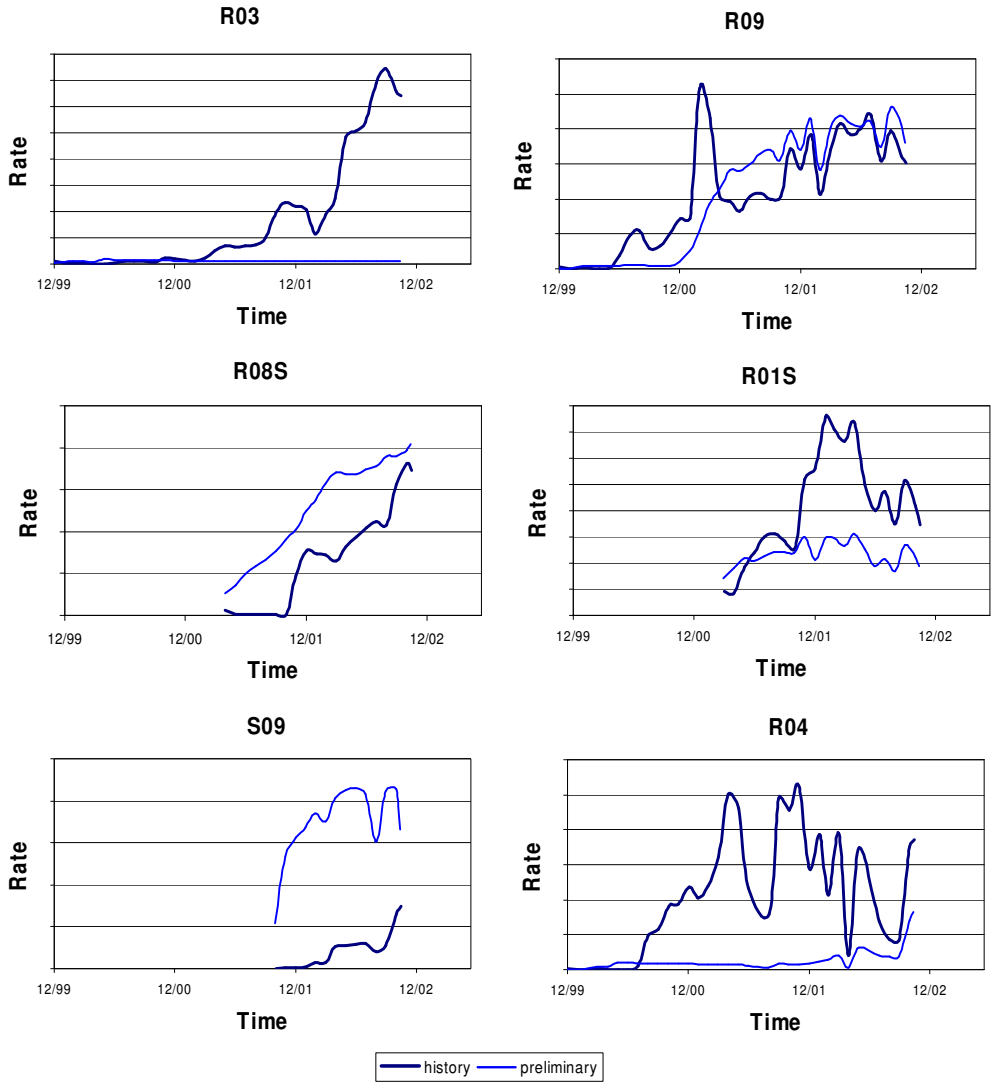

Figure 4.25: Production data from the field and from the preliminary model for six wells in the reservoir.

The water rates from the six production wells in the reservoir are displayed in Figure 4.25 as the solid navy-blue lines. The water rates from the preliminary model are also in the figure as lighter blue lines. Of the 20 production wells, only six are matching satisfactorily in the preliminary model, and one of them (R09) is displayed in Figure 4.25. Some wells in the preliminary model are breaking through too early, but most are not producing enough water and breaking through too late.

A number of parameters were examined in the pre-history matching stage. Some of the parameters considered were fault transmissibility, k_v/k_h ratio, relative permeability, porosity and permeability, and facies distributions. Presently, the comparison of production response is discussed only for two parameters: porosity/permeability and facies locations. In the first case, the facies locations are frozen and the porosity and permeability are perturbed; for the other case, the porosity and permeability distributions are fixed and the facies locations are perturbed. The results from a typical well (R03) are displayed in Figure 4.26.

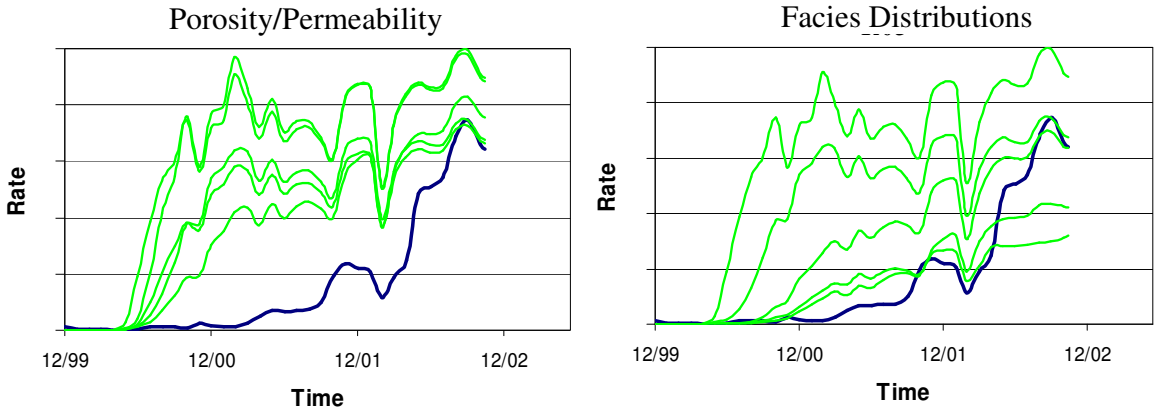

Figure 4.26: The comparison of matching potential for porosity/permeability and facies distributions for one well (R03). The green lines are for five realizations that are not conditioned to dynamic data, and the blue lines are for the field data.

Of all the properties examined, porosity/permeability caused the second largest range of response in the model production, yet from Figure 4.26, it is apparent that a match could

not be achieved by freezing the channel geometry and perturbing the porosity and permeability alone. However, by perturbing the channel locations, a match appears feasible. The behavior was not specific to this well; all 20 wells had similar responses. Therefore, this reservoir model is matched by perturbing only the facies geometry.

The locations of the channel facies are perturbed by changing the random seed in the outer loop and optimizing the sets of r_{Dk} values in the inner loop. When starting each inner loop, the initial r_{Dk} values are equal to 0.50. After three flow simulations in the inner loops, the overall objective function value is compared to the previous best overall objective value. If the new one has improved, the inner loop continues its optimization. If it has not improved, the inner loop is ended and the random seed is changed to start a new inner loop.

The MP-PPM is used with regions of constant r_{Dk} values. The field under consideration has been water flooded from almost its beginning, so streamline defined regions are both possible and appropriate. The regions are controlled by the production wells (one region per well), so 20 regions are used. As the model is perturbed and new realizations are generated, the facies distribution changes. Consequently, the drainage area of the production wells is different for different realizations, and the streamline-defined regions have to be updated. This is done one time at the beginning of each outer iteration. The region definitions for one realization are displayed in Figure 4.27.

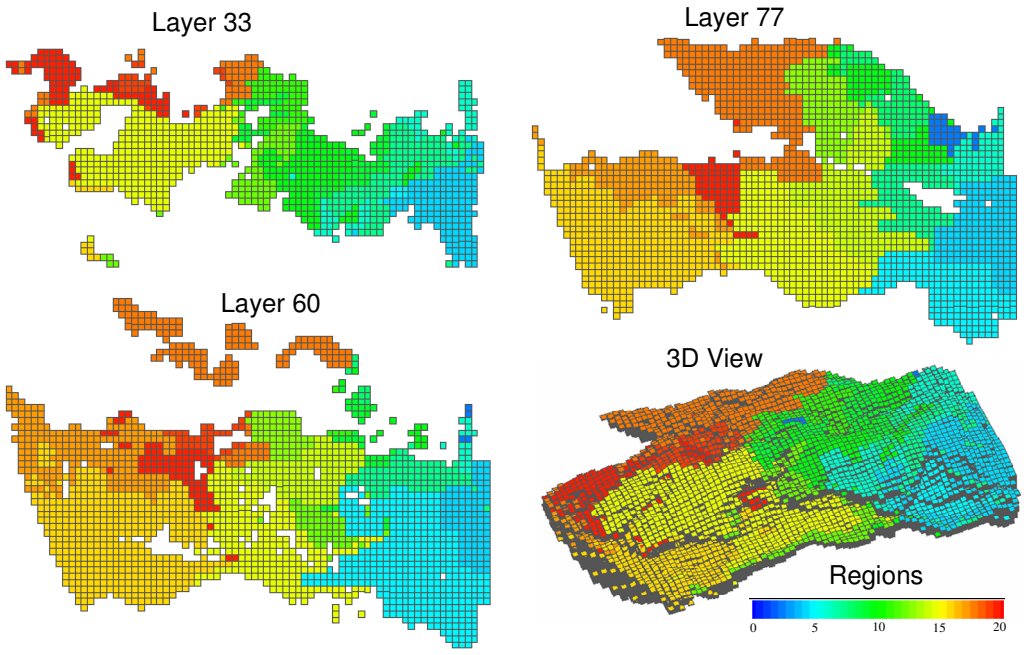

Figure 4.27: Streamlined defined regions for one realization in 3D and for three selected layers.

The streamline simulation is only used to define regions (i.e. it is not used to deduce oil and water rates). In the reservoir (and the finite difference model), wells come on line at different times over the 3 years; whereas in the streamline simulation model, all wells are included at the initial time. All wells must be included at the initial time to ensure that each region contains the gridblocks that are most consequential for flow to the well in that region. The production rates are averaged over the life of each well, and the streamline simulation rates are constrained to these averages. Because streamlines change very little when the well boundary conditions are constant and the streamline simulation is only used to define regions, the streamline simulation only needs to be run for a short duration (for this case ten days).

Geologic Model A – Results

In Model A, the deposition is considered to be a channelized system with a series of debris flow lobes. The lobes have poorer rock properties than the channels. There will

be four facies: channel, small higher quality lobes, large lower quality lobes, and background shale. The training image (Figure 4.28) contains these features and will be used to simulate the facies.

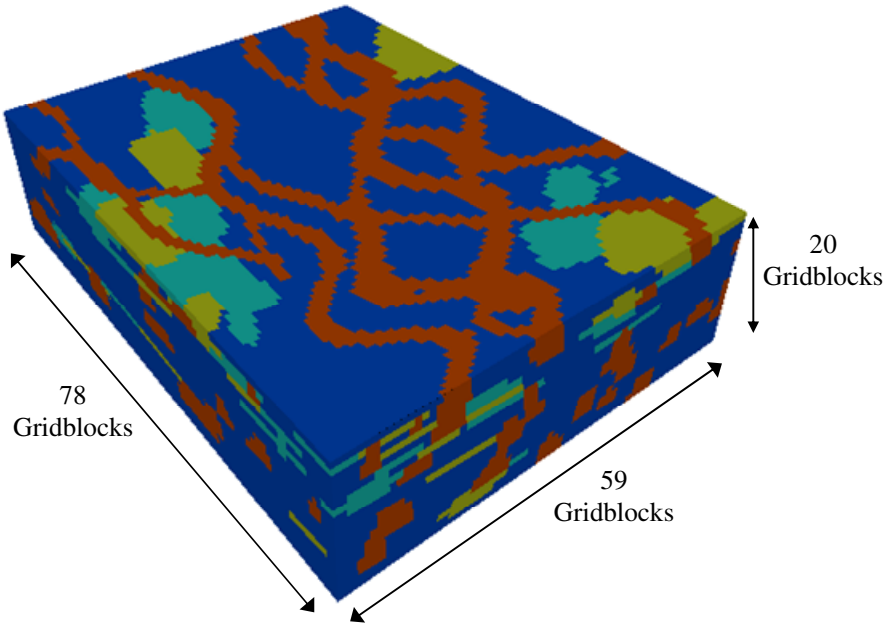

Figure 4.28: Four facies training image for WCA reservoir for Model A.

An object-based geomodeling technique is used to generate the training image. The image is 78 by 59 blocks areally and 20 layers in the depth. The proportion of channel is 0.31, small (yellow) lobes is 0.11, large (light blue) lobes equals 0.10, and shale is 0.48. For Model A, seismic information is not used in the geostatistical algorithm. Figure 4.29 displays the facies for two layers of a realization that is only conditioned to the static data (i.e. not history matched). The well locations for the wells that intersect the layer are also provided in the figure.

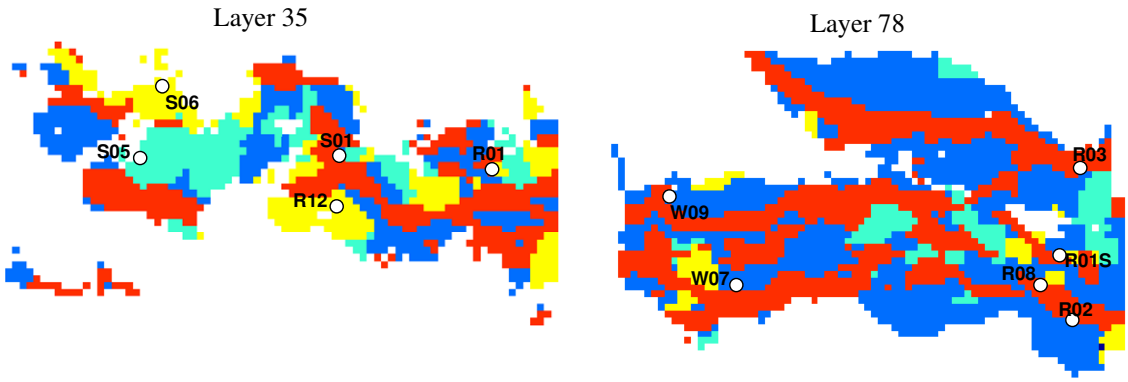

Figure 4.29: Facies distributions for initial realization.

The objective function value or mismatch is defined as the standardized least square error of the water rate data for the 20 production wells. Only water rates were included in the objective, but when the water rates matched, the oil and gas rates also matched because the wells are constrained to total produced liquid. The match requires only 27 flow simulations and six outer iterations. The objective function value for the initial realization is 132, and it is reduced to a value of 83. The match of the total rates is good for all three fluids (Figure 4.30).

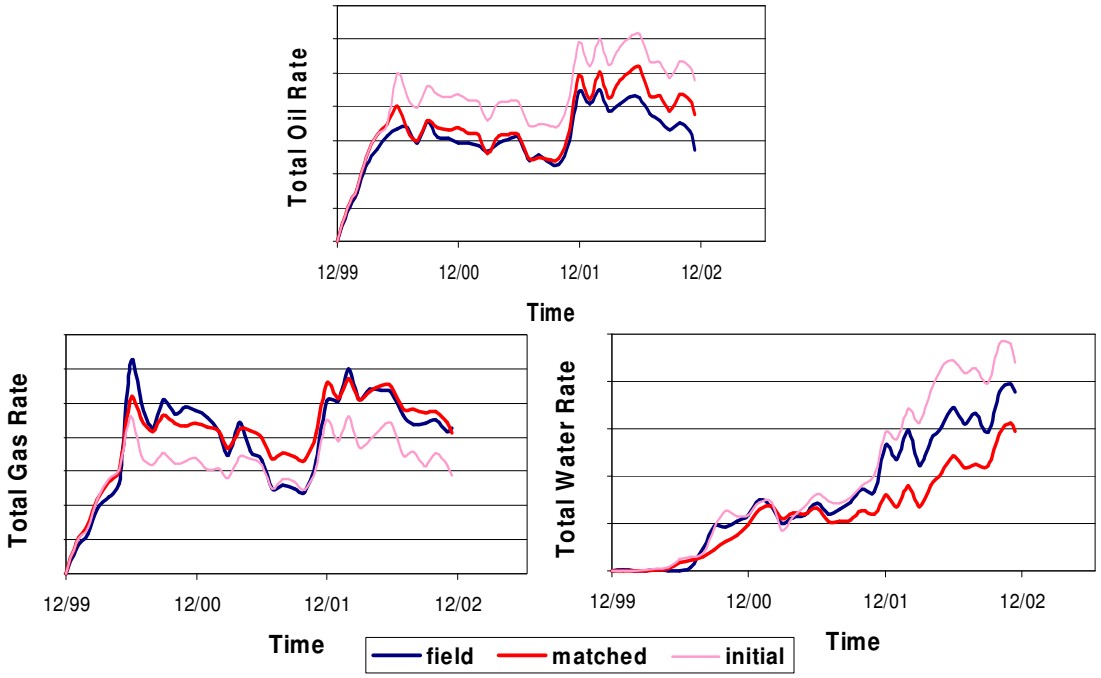

Figure 4.30: Total oil, gas, and water rates for the matched and field data.

In the final year, the oil rate is slightly over estimated, and the water rate is slightly under estimated, but overall the match is good. However, it is not sufficient only to match the total production rates; the individual well rates also need to be matched. Most of the wells (16) are matching to some degree, and only four wells are not matching. Plots are shown in Figure 4.31 of four representative wells (three matching wells and one that is not matching). The reference and matched oil, gas and water rates for all 20 wells are provided in the Appendix B.

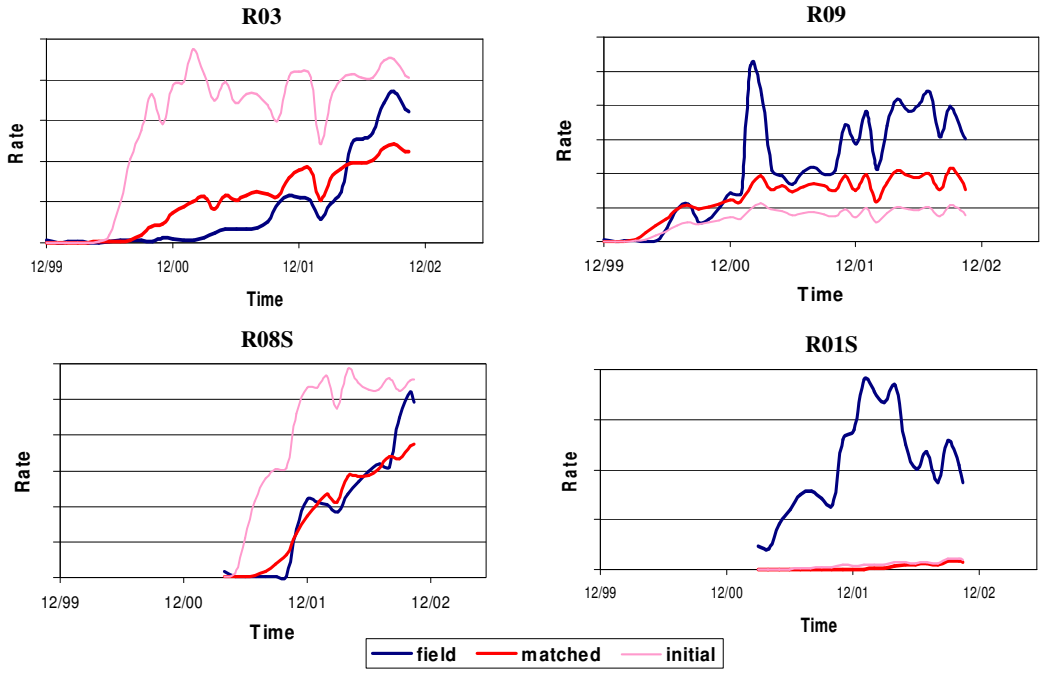

Figure 4.31: Water rates for initial, matched and field data for 4 wells for Model A1.

Well R08S has a high quality match, and wells R03 and R09 are much improved over the initial realization, and they are considered matched even though the wells are not exactly matching for the entire range. R01 is one of the four wells that are not matched. The reason that the four wells are not matching may be because a reservoir or geological parameter was frozen and should have been perturbed or because of some issue with the production data. Additional work is required to get these four wells to match *and* maintain a geologically consistent reservoir model.

The facies geometry for the initial model is displayed in Figure 4.29 for two layers, and the matched model facies are displayed for the same two layers in Figure 4.32. The locations of wells completed in the respective layers are displayed on the figures. The WCA reservoir model is complex, and the behavior of the production data cannot be fully explained by looking at only two layers, but part of the response from well R03 can be observed in the figure. In layer 78 of the initial model, a channel runs from the aquifer

part of the layer to alongside the R03 well location (Figure 4.29). This causes the initial rate at the well to be very high (Figure 4.31). In the matched model, the channel has moved (Figure 4.32). This caused the breakthrough times and water rates to be reduced.

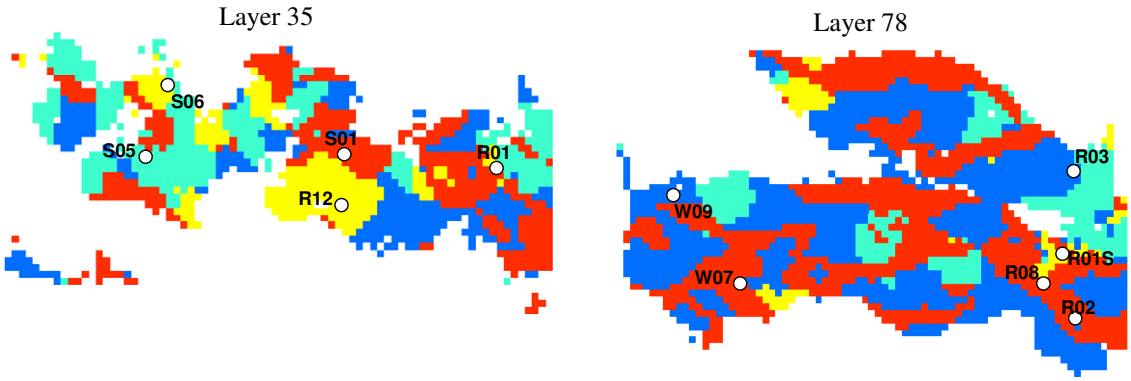

Figure 4.32: Facies distributions for matched realization of Model A1.

For some wells, the match might be further improved by perturbing other model parameters either jointly or in an hierarchical fashion (Suzuki, 2003). For example once the facies locations are determined, the porosity and permeability could have been perturbed. Since we expect the facies geometry to be the dominant factor, this was not done for this work. However, to show the possible improvement if history matching was continued by perturbing these parameters, five realizations of porosity and permeability are generated. The water rates for two wells are displayed in Figure 4.33. The red lines represent the history matched model obtained through perturbation of the facies geometry. The green lines show the response from five realizations generated by changing the porosity and permeability while freezing the facies from the history matched model. While it appears the R03 well could be improved, the R09 well most likely would not improve its match.

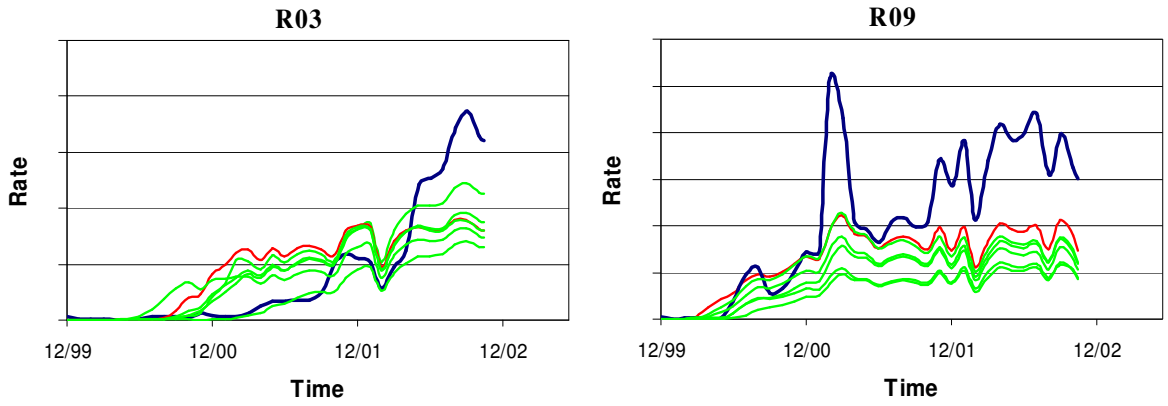

Figure 4.33: The possibility to improve the match for two wells (R03 and R09) by perturbing porosity and permeability. The red lines represent the history matched model; the green lines are for five realizations of porosity and permeability for the fixed facies, and the blue lines are for the field data.

By using the same geologic description (training image and well data) but a different initial realization, a second history matched model is generated (Model A2). The quality of the match is almost as good as Model A1. The overall mismatch is 91 so it is slightly higher than A1, but the initial mismatch is 141, so the improvement is about the same. The match was achieved in only 14 flow simulations. The same 16 wells are matched, yet the match for a few wells is not quite as good. The mismatch for one of the four previously non-matching wells is significantly higher accounting for most of the difference in the overall value. Plots for the same four wells given in Figure 4.32 are shown in Figure 4.34.

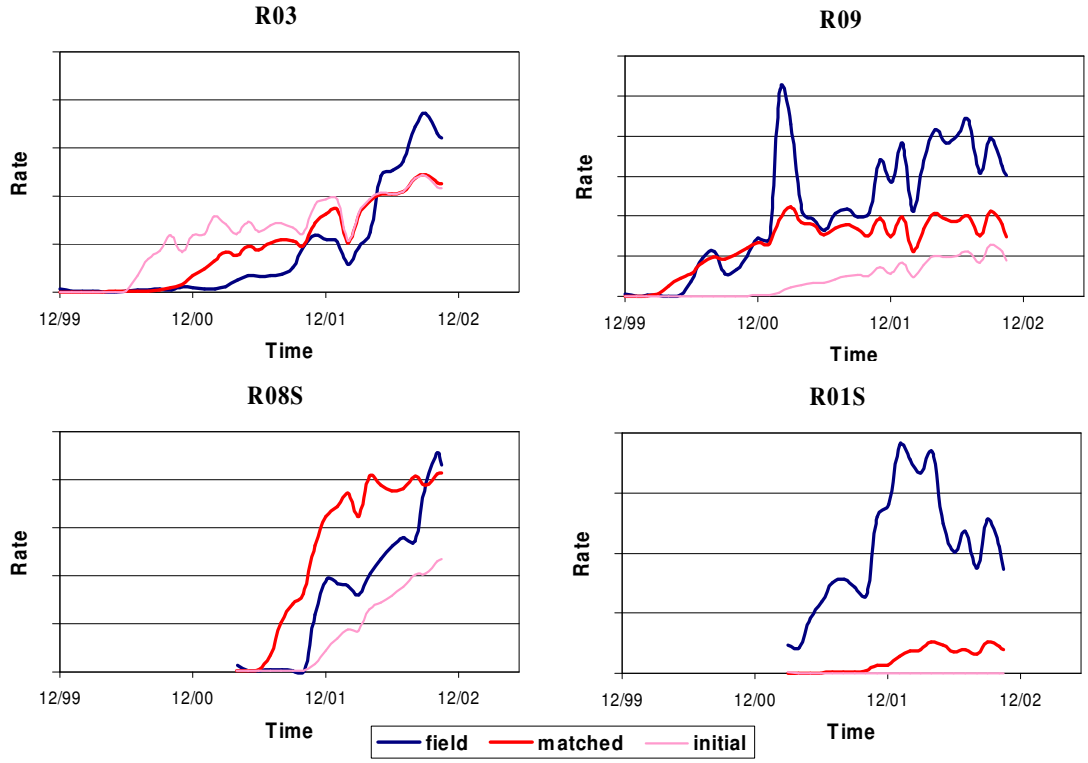

Figure 4.34: Water rates for initial, matched and field data for 4 wells for Model A2.

The quality of match for R09 and R03 is about the same as Model A1, but the match for well R08S is not as good. R01S is slightly better but not good enough to be considered matched. The same layers in Figure 4.32 are also displayed for the A2 model (Figure 4.35). In the initial there is not much channel around the R03 well, so the initial is already matching well and the history match only improves the quality of the match a small amount for this well.

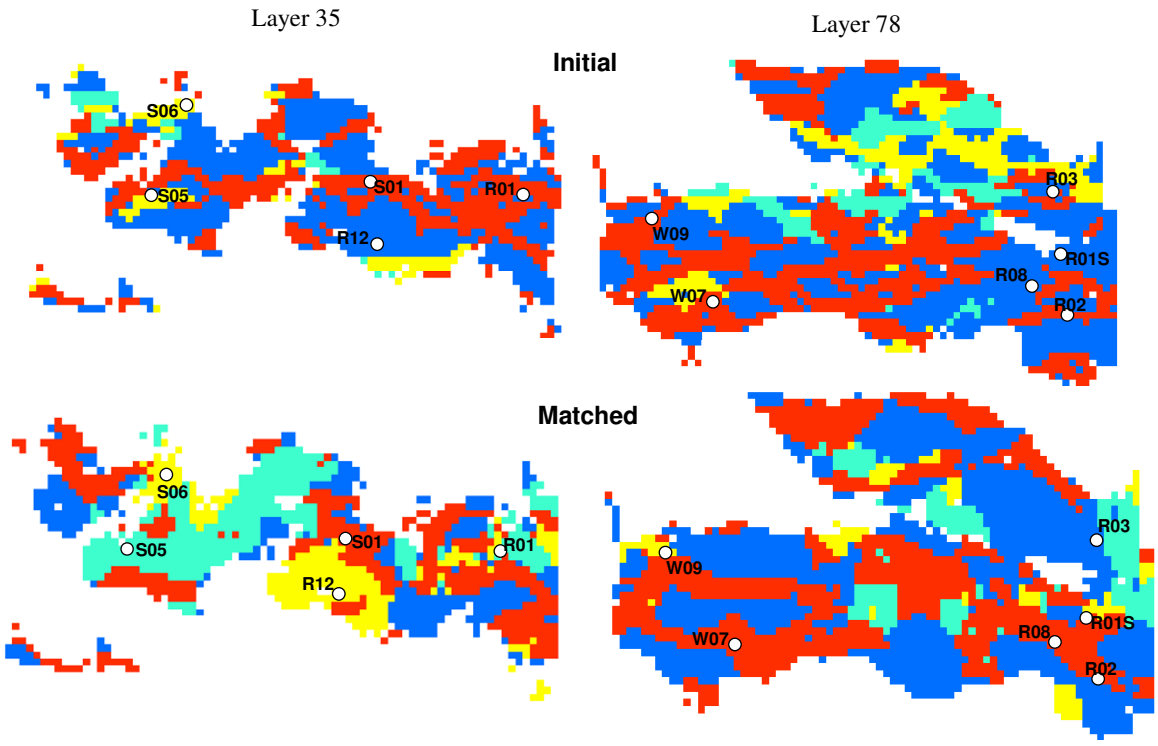

Figure 4.35: Facies distributions for initial and matched realization of Model A2.

A third history matched model was completed for the same training image; however, this time the facies geometry is not conditioned to the "hard" facies indicator data that come from well measurements. This allows more variability in the position of the channels, hence, the match is a little better for this case, especially at a couple of wells; for example, wells R04 (Figure 4.36) and S04 are matched much better. The match is only slightly better for the four wells previously discussed (R03, R09, R08S and R01S); however, the global mismatch has improved. The value for the global mismatch is 79 for model A3 while it was 83 and 91 for models A1 and A2, respectively.

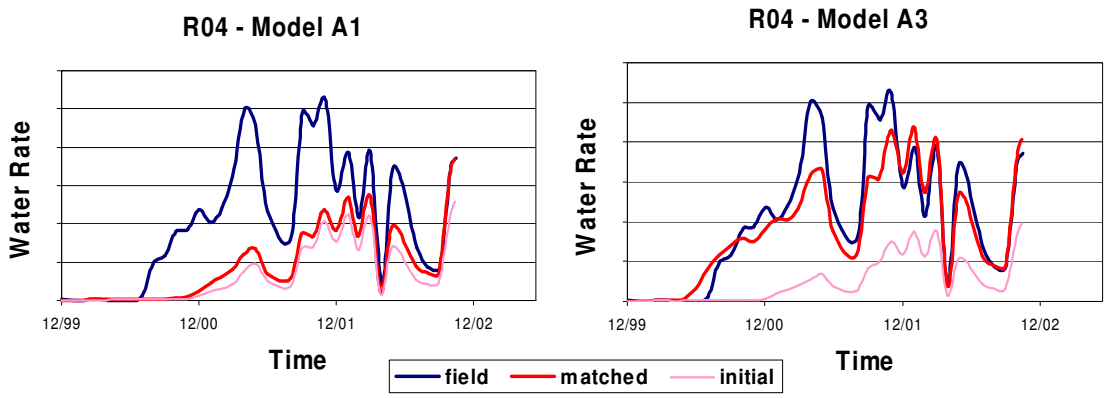

Figure 4.36: History match for well R04 when hard facies data is honored (A1) and when hard facies data is not honored (A3).

There may be conflict between the facies indicator data and production data, which allows the unconditional case to have a better match. Is it reasonable, however, to disregard all of the hard data to improve the match of the historical data? There are some issues with exactly honoring the hard well data directly in a reservoir simulation model. For one, important uncertainty exists for the facies interpretation from the well logs. Additionally, the scale of the measurements compared to the scale of the gridblocks raises the "missing scale" issue. The well data are measured on a very small scale (a few inches to a few feet), but they are assigned to a large gridblock (hundreds of feet). The information at the small scale may not be representative for the entire gridblock. Another issue particular to this example is the vertical averaging used. Each gridblock is a few (3-7) feet thick, but the well-log measurements are on the 0.5 foot scale, so there can be up to 15 measurements for one gridblock. For facies data, the measurement that is closest to the center of the gridblock is assigned to that gridblock. When the same facies occurs throughout the gridblock, this method is acceptable, but when multiple facies occur, a gridblock may be assigned a facies that is not representative of the block.

Although there are some concerns with the hard data, ignoring all of the hard data is probably not the best solution. When data is ignored, the model becomes less realistic and has less prediction power. A better solution would be to apply some probabilistic

honoring of the data. At each hard data location, a probability can be assigned to each facies. The probability would be based mostly on the well information, but other information such as its location in the reservoir could also be taken into account. Then, prior to the geostatistical simulation, the facies at the well locations could be determined using a sampling technique (e.g. Gibbs' sampling (Geman & Geman, 1984)). The probabilistic Gibbs' sampling technique is only proposed in the current work but results using it have not been obtained.

There could also be some issues with the production data (e.g. the errors and noise discussed in the History Matching Real Reservoirs sub-section). Therefore, sacrificing the hard data to improve the match may cause poorer predictive models. Having a slightly worse match, but honoring the hard data in a smart way should lead to more realistic, more believable, and more predictive models.

Geologic Model B – Results

The second major geologic model (Model B) also has four facies, but the reservoir facies for this model have all channel geometries. The training image used for Model B is provided in Figure 4.37. The red facies has the highest quality reservoir properties, and the blue facies is the background shale. The yellow and light blue facies have lower quality properties. The proportion of each facies for this training image is as follows: channel equals 30%, shale equals 52%, and the secondary and tertiary sands each have 9%. These proportions are very similar to Model A.

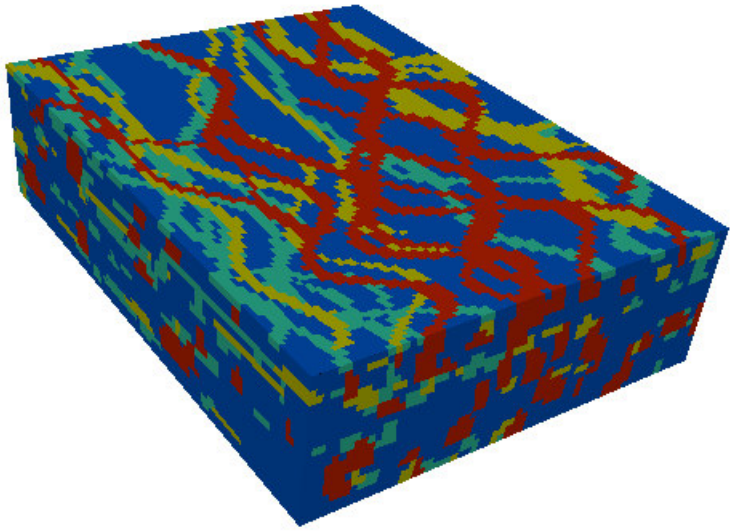

Figure 4.37: Four facies training image for WCA reservoir for Model B.

The other significant difference for Model B is that seismic information is used as soft conditioning for the facies simulation. A cube of seismic-derived v_{shale} was generated from seismic impedance by means of well to seismic calibration. A further transform was applied to generate the four (one per facies) facies probability cubes from the seismic v_{shale} cube so that data could be used in the facies simulation. Due to preferential well drilling, fluid effects and other reasons, the v_{shale} cube may be biased and may reproduce some artifacts inherent to the seismic data. Thus, the seismic v_{shale} cube is only a soft data used to model depositional facies. The soft data from seismic that has been transformed into a probability will be termed, P(A|C).

The development in Eq. (2.2) assumes that there is only one soft data, and for the previous cases, production data (P(A|D)) provided the soft data. However, the extension of Eq. (2.2) for multiple soft data is straightforward (Journel, 2002). Equation (2.1) can be rewritten as

$$\frac{x}{a} = \left(\frac{d_1}{a}\right)^{\tau_1} \cdot \left(\frac{d_2}{a}\right)^{\tau_2} \cdots \left(\frac{d_p}{a}\right)^{\tau_p} \tag{4.1}$$

where

$$x = \frac{1 - P(A \mid D_1 \ldots D_p)}{P(A \mid D_1 \ldots D_p)},\ a = \frac{1 - P(A)}{P(A)} \text{ and } d_i = \frac{1 - P(A \mid D_i)}{P(A \mid D_i)},\quad i = 1, 2 \cdots p\,.$$

D_i can be data of any type, and p is the total number of data types. Once again different tau, τ_i, values can be assigned to the various hard and soft data to account for any dependency among the data. For this example, only three data types are used. P(A|B), P(A|D) and P(A|C) are all combined to create a new probability, P(A|B,C,D) that is used for multiple-point geostatistical simulation. Using these three probabilities and unit tau values, Equation (2.2) becomes:

$$P(A \mid B, C, D) = \frac{a^2}{a^2 + bcd} \tag{4.2}.$$

One layer of the seismic probability map for the red channel facies is shown in Figure 4.38. The red colors indicate a high channel probability while blue indicates low channel probability. Seismic data is able to distinguish some areas with a much higher probability of channel. Seismic is only used to influence the location of the facies; the petrophysical properties are calculated with the same procedure used for Model A.

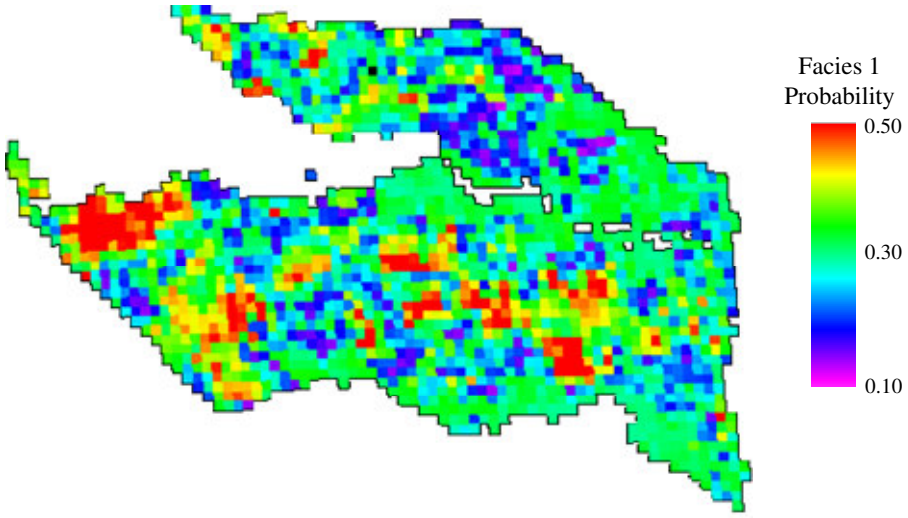

Figure 4.38: Channel probability derived from seismic data for one layer. Red is high probability and blue is low probability.

The facies locations and permeability for one layer are displayed in Figure 4.39. Due to the different training image, the channels are thinner than Model A, and the poorer quality sand facies have a less lobe-like structure.

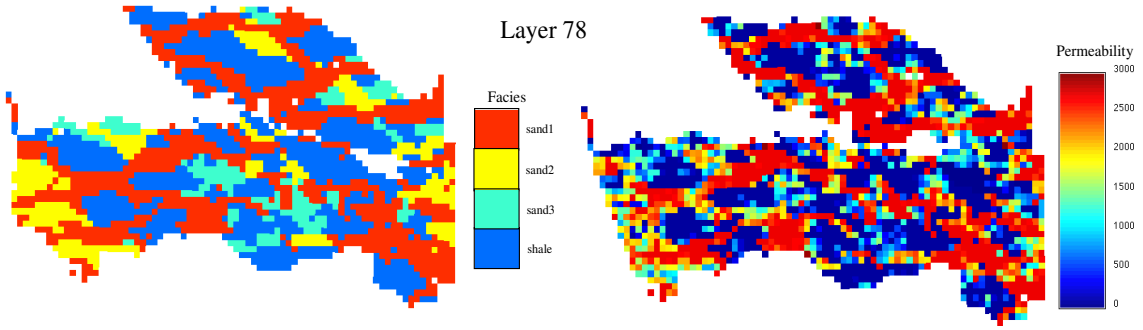

Figure 4.39: The facies map and permeability (md) map for a matched realization of model B are displayed.

By perturbing the channel locations, two equiprobable history matched models (B1 and B2) are created with this description. Each model started with a different initial model. Both models match satisfactory, but the matches are not as good as Model A. The overall mismatches are 98 and 97, respectively, and only 15 wells are considered matched. The

same four non-matching wells from Model A and one additional well are not satisfactorily matched for this model. The number of flow simulations needed to reach this match equals 26 and 32 for the two models. The production data at four wells for Model B1 is presented in Figure 4.40.

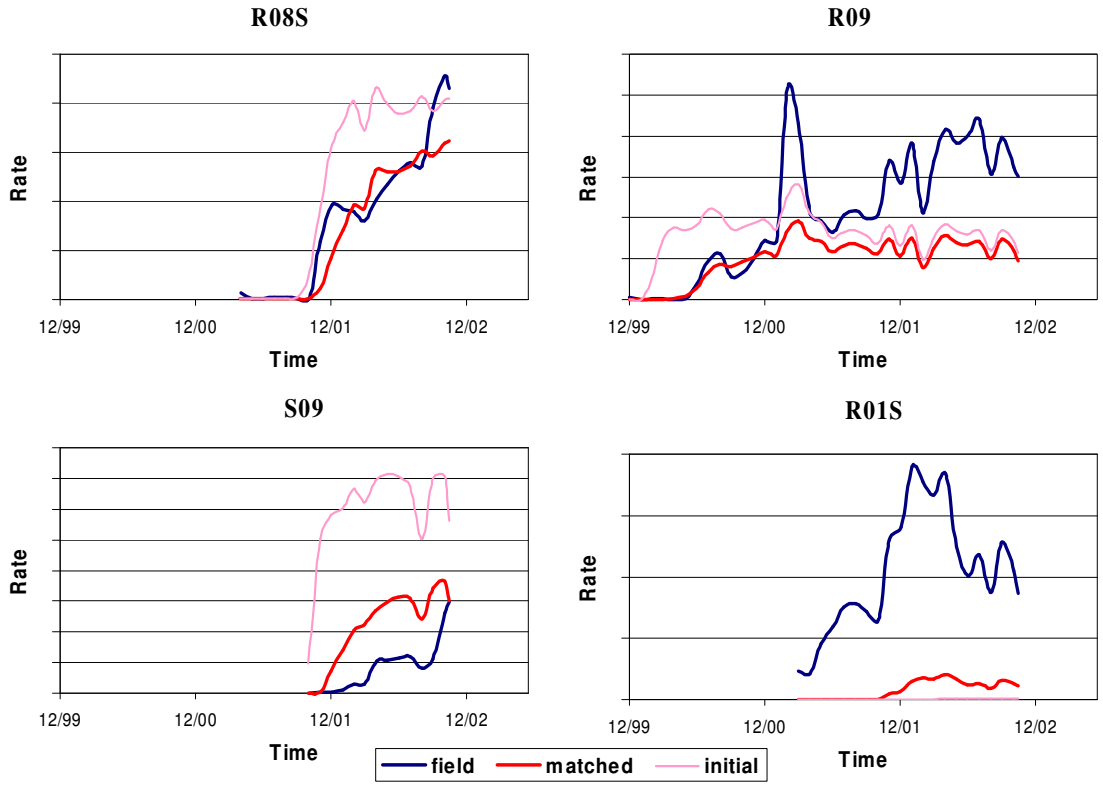

Figure 4.40: Water rates for initial, matched and field data for 4 wells for Model B1.

R08S again has one of the best matches. The breakthrough time improves for R09, but the production from the last year of the model continues to under produce the water measured in the field. S09 dramatically improves over the initial, but it is over producing field data, and again R01S is one of the wells that is not matching.

A third model (B3) is also history matched for this training image, but the seismic data is not used to condition the facies for this model. Model B3 matches equally well as B2 or B1 (mismatch equals 99) and requires approximately the same number of flow

simulations (30) to match. Additionally, in the matched realizations, the primary (red) sand facies are located in the same general positions (Figure 4.41) for the case with seismic (B1) and the case without seismic (B3).

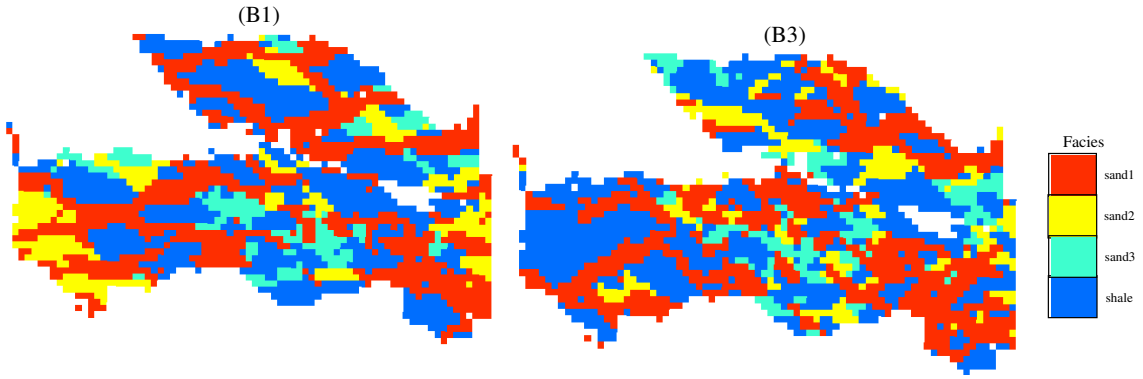

Figure 4.41: Facies map for history matched realizations when seismic is used to condition the facies (B1) and when seismic is not used to condition the facies (B3).

For this particular model with 28 wells and a fairly coarse grid, the well data provides significant information, and tightly controls the spatial distribution of the facies; therefore, the seismic becomes quite redundant. The redundancy could have been taken care of by assigning values lower than 1 to the tau-parameters in Eq. (4.1) using a tau-value estimation technique (Krishnan, 2004). Since all tau-parameters are equal to 1, the seismic data tend to reinforce the influence of the well data. According to model A3 (versus A1 and A2), to get a better match, the well data constraint should somehow be relaxed and definitely not be reinforced. This could be one reason that we obtain a worse match for model B.

Actually, the agreement between the seismic and hard data is quite similar and may have some positive aspects. Thirty facies realizations are generated that are not conditioned to production data. At each gridblock, the facies indicators for the thirty realizations are averaged into an e-type estimate. If channel occurs at a gridblock in most of the realizations, then channel will be highlighted in the e-type average. Figure 4.42 shows

the average facies of one layer for two situations: (A) conditioned to the hard data and the seismic data and (B) conditioned only to the hard data.

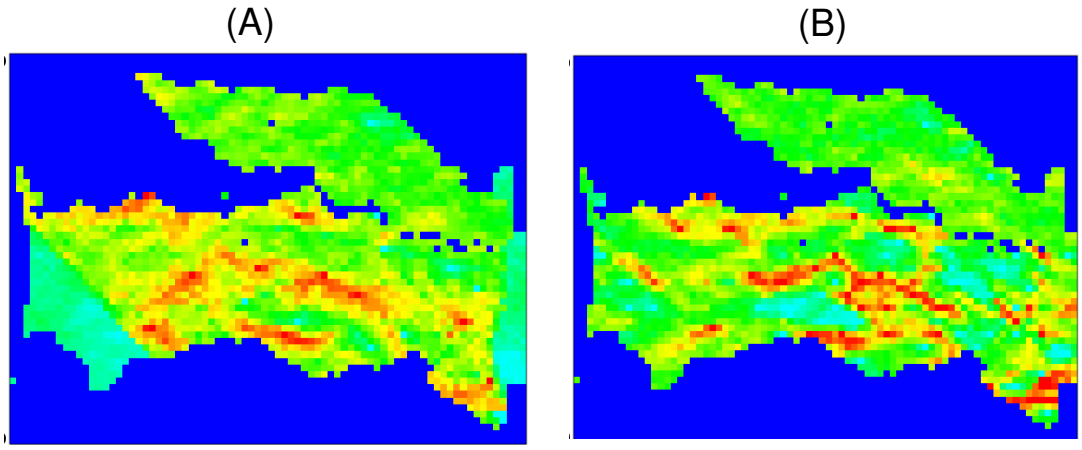

Figure 4.42: The average facies of thirty realizations for one layer. (A) conditioned to the hard data and the seismic data and (B) conditioned only to the hard data.

The channels appear in the same locations with the same frequency for both situations. In other reservoirs, conflict between seismic and well data often exists, so the consistency in the WCA data provides confidence in the static data for this reservoir. Additionally, it further supports the idea of honoring the hard data rather than having a slightly better production data match as discussed in Model A.

At this stage in the reservoir life, the seismic data information is mainly redundant with the numerous well data; however, earlier in the reservoir life when there were fewer wells, the seismic would have been more important for modeling the reservoir facies. Figure 4.43 has the averages from thirty realizations with and without seismic when only three wells are used for conditioning. Channel locations are definitely more visible for the model conditioned to seismic data (A); however, the channels are less defined than when all the hard data are included (Figure 4.42).

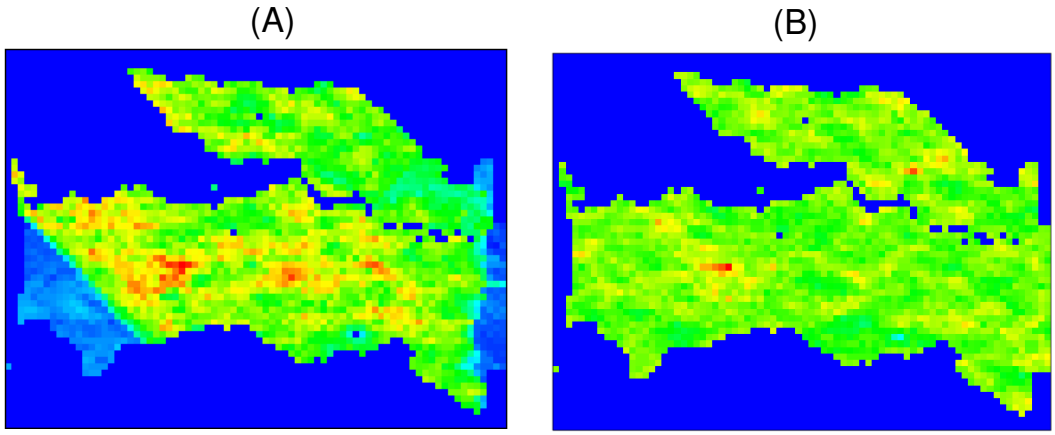

Figure 4.43: The average facies of thirty realizations for one layer. (A) conditioned to the hard data from three wells and the seismic data (B) and conditioned only to the hard data from three wells.

One of the main points of Model B is to history match a different geologic scenario so that some amount of geologic uncertainty is accounted for in the reservoir model. There are further details about this point in the next sub-section. The other main reason for Model B was to demonstrate that the MP-PPM is able to handle more than one soft datam. While this case may not have been the best application to showcase this ability, it did show that traditional soft data such as seismic facies probability can be combined with the production derived soft data when generating history matched reservoir models. The method is very general and is not limited to two soft data; other secondary information (from geology or other types of seismic information) can simultaneously condition the model. Equation 4.1 illustrates how the tau parameters might account for any redundancy or conflict in the data by assigning different weights to different soft data although the tau parameters were fixed to unity for all cases studied.

Validation and Uncertainty

Reservoir models are not created to see how well the model matches the historical production data. They are primarily built as a prediction tool; thus, the previously history matched models will be tested against “future” production. The WCA reservoir was

produced for 3 years, and then it was shut in for about 4 months. Once it was brought back on line, it produced for another 1 ½ years. The models are history matched to the first 3 years of production before the reservoir was shut in; afterward, the models are validated over the remaining 1 ½ years.

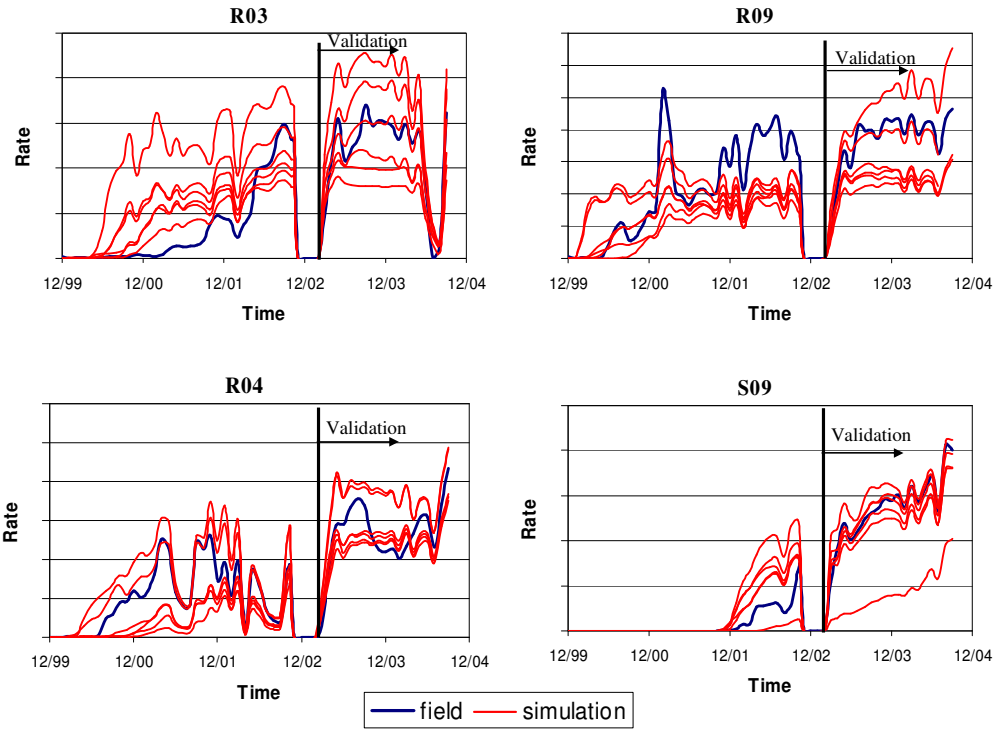

Figure 4.44: Water rates for validation of five history matched models compared with actual production rates.

The history match and validation for four wells is presented in Figure 4.44 for six models (A1, A2, A3, B1, B2 and B3). The oil, gas and water rates for all twenty wells for all six models are included in Appendix C. Flow in the validation period is still conditioned to the total fluid rate, so some of the peaks and valleys from the field data are captured in the simulated production. Nonetheless, the validation data is not included in the objective function, so the quality of the replication of the field data during the validation period is quite remarkable. The average mismatch for the validation is 102. This is the

average of the six models and it is standardized over time so a comparison can be made with the history portion of the model. The average of the six models for the history portion is 91. The validation is not quite as good as the matching period, but from model to model and from well to well the similarity varies considerable, and some wells are even closer matched in the validation period.

For all four wells, at least one realization is predicting very accurately, and all of the realizations (with the possible exception of one in S09) are good. The models do not all exactly match the prediction from the field data, but the six models provide a nice range of uncertainty. For most wells, the field data falls within that range. We did not choose six "good" models to show. These six were the only history matched models that had validation runs completed on them.

These six models represent some measure of uncertainty for the WCA reservoir, but by no means does this capture the full range of uncertainty. We have captured the uncertainty by using different training images and different seismic and well conditioning, but to capture a more realistic reservoir uncertainty, additional training images need to be incorporated and other properties need to be perturbed (e.g. fault transmissibility, relative permeability, etc.). Ten to twenty (or more) history matched realizations would be required to capture the most significant amount of uncertainty. Nevertheless, these six models do provide a much clearer understanding of the reservoir uncertainty than a single realization (or multiple realizations that only differ by the random seed) could.

Summary

The WCA reservoir has been history matched with two separate but related geologic scenarios. There were three history matched realizations for each scenario. High quality matches were obtained for most of the wells. Within each scenario, additional aspects such as conditioning to seismic information and conditioning to hard data were also studied. For this case, seismic brought little additional information, but this example

demonstrates that the method is able to use multiple soft data. When hard data are not honored, a better match could be achieved, but sacrificing hard data information to improve the match probably leads to less reliable models. After the history match period, the wells are simulated for an additional one and a half years. The agreement between the field data and the simulation models during this period is extremely good. Table 4.3 summarizes the results for the WCA reservoir.

Table 4.3: A summary of history matching results for the WCA reservoir.

Model Name	Model Description	Number of Flow Sim.	Initial Error	Matched Error	Validation Error
A1	TI1, Hard cond. data, No seismic	27	132	83	91
A2	TI1, Hard cond. data, No seismic	14	141	91	112
A3	TI1, No hard cond. data, No seismic	39	133	79	97
B1	TI2, Hard cond. data, Seismic	26	154	98	116
B2	TI2, Hard cond. data, Seismic	32	146	97	102
B3	TI2, Hard cond. data, No seismic	10	140	99	96
	Averages	25	141	91	102

For the three years of history match, the average error of the six modes equal 91, and it ranges from 79 to 99. We are able to reduce the mismatch by about 40% from the initial models, and we are able to achieve this reduction on average with less than thirty flow simulations. The amount of mismatch for the validation runs is very similar to the amount for the history matched period, which provides some confidence that the models are not overmatched.

The main conclusion from the study is that real reservoirs can be history matched by perturbing large scale properties such as the locations of channels. With the MP-PPM, a match can be achieved while honoring the prior geologic information from the reservoir, and when reservoir models are history matched and honor the geologic description, the "future" production from the models can be very accurate.

"An investment in knowledge
pays the best interest."

- Benjamin Franklin

Chapter 5

Perturbing Proportions

The probability perturbation method is a tool to assist history matching reservoir models, and it can be tailored to particular problems. For the case study in the previous chapter, a quality match was achieved simply by perturbing the location of facies; however, for other cases, perturbations of this type are not enough, and other properties in the reservoir model will have to be changed. Chapter 5 demonstrates one way the method is extended to include other properties. Instead of perturbing only the location of facies, both the proportion and locations of facies are perturbed in a consistent manner. The perturbation of these two properties is coupled through the perturbations parameters, r_{Dk}, thereby maintaining the efficient parallel optimization of before. Additionally, different regions of the reservoir can be perturbed by different amounts. The new method is first tested successfully on a 2D synthetic example. Then it is applied to a real North Sea reservoir; this reservoir contains low permeability calcite nodules that significantly influence flow paths in the reservoir. By regionally perturbing the locations and proportions of these nodules, the field production data (pressure and rate data) is more accurately reproduced.

5.1 Limitations of 'Random Seed' Perturbations

Geostatistical simulation techniques require a random number series. In sequential simulation, these numbers determine the path used to visit the nodes and the random drawing process. A unique number (random seed) defines a different series of random numbers (given an efficient random number generator). If all the parameters in the geostatistical algorithm are kept constant and only the random seed is changed, multiple "equiprobable" realizations can be created. There is some amount of variability among the different realizations; for example, Figure 5.1 displays three equiprobable realizations for a continuous and a non-continuous case. For each case, the parameters of the algorithm and the algorithm used to generate them are kept constant. The spatial variability of the continuous property and the position of each facies are now uniquely determined by a single random seed.

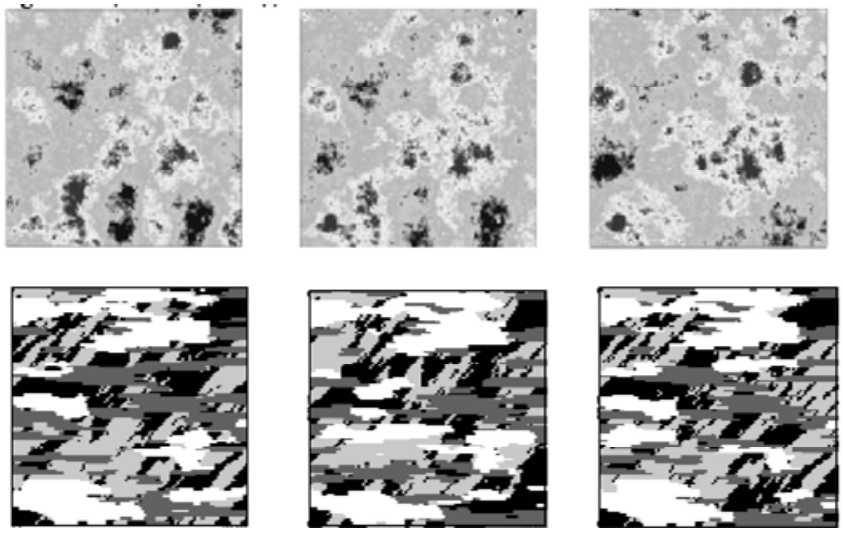

Figure 5.1: Three equiprobable realizations for continuous case (top) and multiple facies case (bottom).

By changing the random seed, a realization can be generated that matches some defined criteria better than any other realization. Essentially, this is what is done in the examples from the previous chapters. History matches are achieved by moving the facies locations

around. However, for other examples, an adequate match may not be achieved. Often by only perturbing the random seed, not enough change is created to be able to achieve a history match. The random seed uncertainty allows only minor changes in the model, while the other two types of uncertainty (reservoir parameters and geologic scenarios) allow the models to change by much larger amounts.

As discussed in the Uncertainty section of Chapter 1, the uncertainty on geologic scenario is not easily perturbed in a history matching sense. A scenario is intrinsically discrete in nature, and there are only a few. Therefore, a scenario should be considered as an important part of the sensitivity study rather than a parameter to be perturbed in a history matching task. Alternatively, reservoir parameters can be perturbed within their range of uncertainty, and there are a wide variety of parameters that can be perturbed. Traditional history matching parameters such as relative permeability and well productivity index can be included. Likewise, geologically consistent history matching parameters can also be perturbed. Continuous variables such as variogram parameters (range, azimuth, etc.) and the mean of permeability or porosity are good examples. For large-scale parameters, work has been done to include such things as the preferential direction and proportion of facies (Kim and Caers, 2003) and the fracture density (Suzuki *et al.*, 2005).

Figure 5.2 shows the production data for an example with geobodies. In the case on the left, only the random seed is perturbed, and in the other case, both the random seed and the proportion of those bodies are perturbed. This example illustrates that for this well the breakthrough times could not be matched by relying on the random seed uncertainty alone.

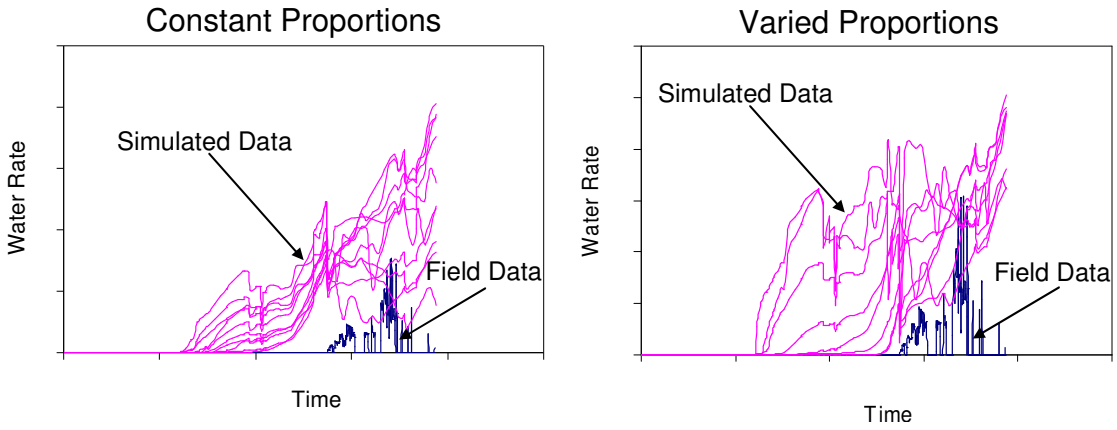

Figure 5.2: Sensitivity of two parameters and their affect on production data.

All of the reservoir properties can be perturbed along with (or in place of) the random seed. The one aspect that we will explore in this chapter is the uncertainty associated with the proportion of facies bodies.

5.2 Perturbing Facies Proportions

One aspect of the geologic model that is important for a large number of applications is the local proportion of facies. The local proportion (LP) is the amount of a certain facies in a region expressed as a fraction or percentage of the total region. Because of limited measurements and imprecise knowledge of the geology, there is often a large range of uncertainty for the local proportion of geologic bodies. Moreover, the LP of these facies bodies can be significantly different for various parts of the reservoir. Therefore, when history matching these types of reservoir models, the LP must be perturbed by varying amounts in different reservoir regions. The following paragraphs describe how facies proportions can be perturbed in conjunction with the probability perturbation method.

If the production data in a region closely matches the historical field data, only small changes should be made to the local proportion (LP), however for a poorly matched region, larger changes should be made. One could introduce the LP of each region as an additional history matching parameter; however, this would increase the dimensionality of the optimization problem. Instead, the perturbation of local proportions is coupled to set of the parameters, $\{r_{D1}, \ldots, r_{DK}\}$. When the production match is good, r_{Dk} decreases, hence, LP should remain similar as well. When the production does not match well, r_{Dk} increases, then LP should change considerably. Using Equation (5.1) the perturbation in LP is coupled to r_{Dk} as follows:

$$LP_k^{NEW} = LP_k^{OLD} + i_k \bullet (r_{Dk}) \bullet Fc \quad \text{for } k = 1, ..., K \tag{5.1}$$

where K is the total number of regions. Fc is a user-defined constant that characterizes the amount of change allowed each step. Since the values of r_{Dk} range from 0 to 1, when r_{Dk} equals 1, LP^{OLD} is either increased or decreased by an amount equal to Fc. From sensitivity studies, good values for Fc range from 1% to 10% and are usually around 4%. Too much or too little change per iteration may decrease the efficiency of the method.

By linking the perturbation of LP_k to the parameter r_{Dk}, the dimensionality of the optimization problem does not increase.

The indicator term, i_k, determines whether the LP should increase or decrease and is defined as follows:

$$i_k = \begin{cases} 1 & \text{if increase in local proportion is desired} \\ -1 & \text{if decrease in local proportion is desired} \end{cases} \quad \text{for } k = 1, ..., K. \qquad (5.2)$$

For simple cases where the relationship between the amount of facies bodies and the production data is well understood, i_k can be determined prior to running flow simulation. Consider an example where a set of high permeability bodies are located between an injector and a producer. If the model is over predicting the water cut at the producer, then the proportion of bodies needs to be reduced, and if it is under predicting, the amount of bodies needs to be increased. However, one can imagine that in some cases the relationship between facies proportion and production data is not as trivial. For example if a reservoir has channels that run both North-South and East-West, increasing the LP may increase or decrease water production from a production well located North of an injector. Likewise if facies bodies have low permeability, they may inhibit (i.e decrease) flow, or they may channelize (i.e. increase) flow between a producer/injector pair. Hence, increasing the LP can lead to an increase or decrease in flow. In these cases a different approach must be used (Kim & Caers, 2003). The probability perturbation method is still utilized, but since it is not known whether the LP should be increased or decreased, one simulation is completed where the LP is increased and one simulation is ran with a decreased LP. Then the simulation that best matches the field data is selected to continue the perturbation algorithm.

Since LP is constantly perturbed, a mechanism is designed to constrain the new realization to the perturbed local proportion (LP). In order to achieve a given LP in each region, we expand on the servo-system idea of Strebelle (2002). The servo-system of

Strebelle is used to achieve a given LP for the entire reservoir. This is done by adjusting (steering) the initial probability model, $P(A|B)^{OLD}$, at each unsampled gridblock. The adjusting amount is proportional to the difference between the simulated proportion, LP^{SIM}, and target proportion, LP^{NEW}.

$$P(A \mid B)^{NEW} = P(A \mid B)^{OLD} + \mu (LP^{NEW} - LP^{SIM}) \tag{5.3}$$

A user-defined servo-system parameter, μ, defines how precisely the LP will be reproduced. Values of μ that are high cause the generated realizations to honor the LP more closely than low values. To accommodate regions, the target LP may be different for different regions (Eq. 5.1), and the simulated LP will be calculated per region.

$$P(A \mid B)^{NEW} = P(A \mid B)^{OLD} + \mu (LP_k^{NEW} - LP_k^{SIM}) \tag{5.4}$$

For the following examples, the servo-system parameter was set high so that the specified regional LP will be honored in the new geostatistical realization. Therefore, the resulting realizations will have different proportions for different parts of the reservoir model.

5.3 2D Synthetic Example

The synthetic example used to illustrate perturbing local proportions is a 100 by 100 two-dimensional model. It is a two facies model with the high permeability facies equal to 1000 md and the low permeability facies equal to 100 md. The overall high-permeability proportion for the reference is 30%, but it ranges from 15% to 45% for four distinct regions (Figure 5.3).

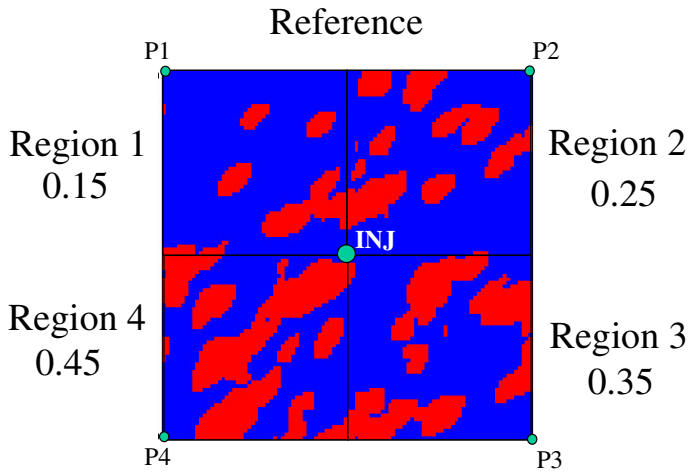

Figure 5.3: Reference model for proportion perturbing example. The regional reference proportions and well locations are shown.

There are four producing wells corresponding to the four regions and one injecting well at the center of the model (standard five-spot). Water cut is to be matched at the four producers, and the reference water cuts are displayed in Figure 5.4.

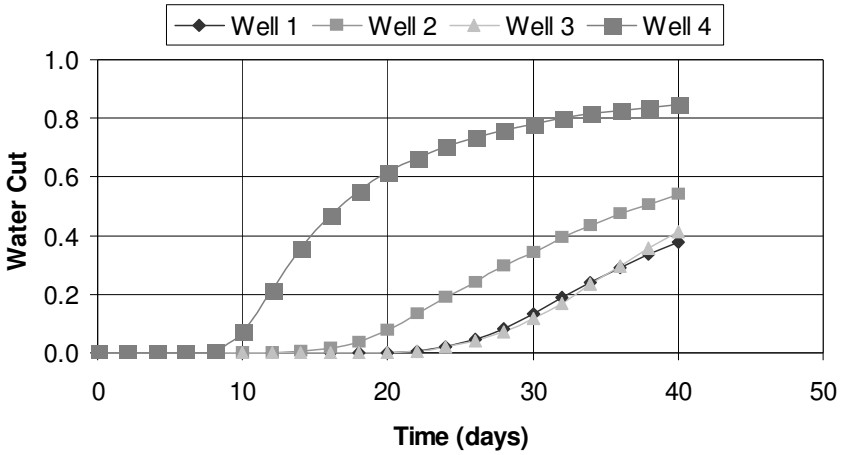

Figure 5.4: Reference water cuts for four wells.

The water cut for Region 4 is much greater than the other three regions. This is not solely due to the increased net to gross ratio. The facies are configured in such a manner that there is a well-connected high permeability path between the injector and producer (Figure 5.3), which increases the water cut for the well in this region. Likewise, Region 2 has the third highest proportion of high permeability facies, but it has the second highest water cut. In addition, Region 3 has a water cut that is almost the same as Region 1, even though the LP for Region 3 is 35% and for Region 1 it is 15%. This is again due to the configuration of facies bodies in each region. The high permeability facies for Region 3 are mostly on the edge of the region, and there are few high permeability facies bodies in the direct path between the injector and producer, P3. Conversely, almost all of the high permeability bodies for Region 1 are between the injector and the producer, P1. Since about the same amount of high permeability facies is between the injector and the two producers, the water cut from these two wells is about the same. Similarly, the water cut in Region 2 is higher than in Region 3 because Region 2 has more high permeability facies in the direct path from the injector to the producer.

The proposed method where both the proportion and the location of facies are regionally perturbed is compared with two cases where only facies locations are perturbed: SP-PPM

and MP-PPM. Ten different realizations were completed for the three cases, and all realizations were matched to an error tolerance of 0.02. Figure 5.5 shows the initial realization and matched realization along with the local proportions for each matched region for Realization 7.

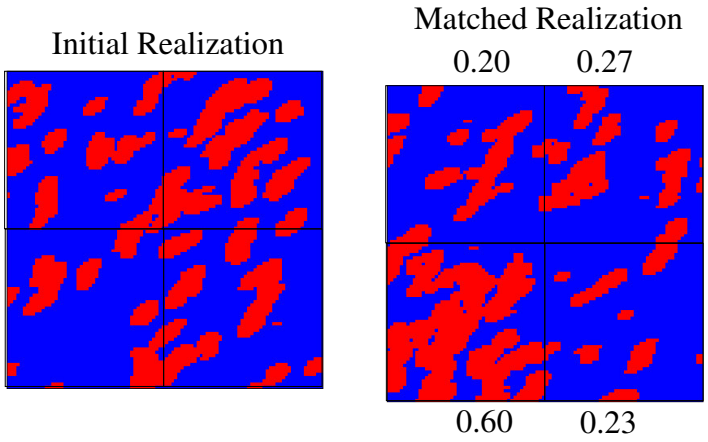

Figure 5.5: Initial and matched realizations including matched proportions for Realization 7.

It is visually apparent that Region 4 (lower left) has a larger proportion of the high permeability facies, but it is difficult to distinguish differences in the other three regions. Figure 5.6 shows how well the water cut matches for the history matched realization and the initial realization.

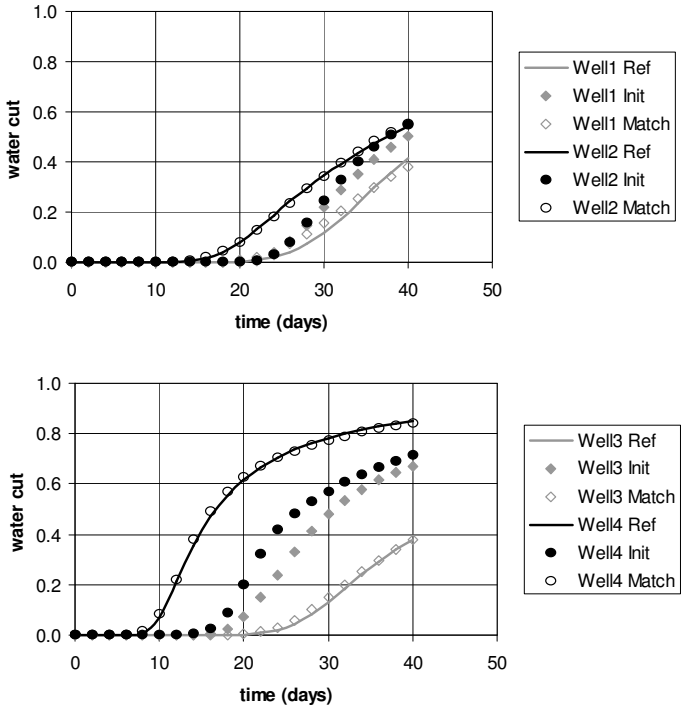

Figure 5.6: Reference water cut, initial water cut, and matched water cut.

All wells show a good match to the reference production data. Table 5.1 shows the number of outer iterations to converge for both methods and all ten realizations. The last columns show the average number of iterations for each case. The entry ">100" states that the realization could not converge in 100 iterations, and for averaging purposes, a value of 101 is used.

Table 5.1: The number of outer iterations for three cases to match. MP-PPM-wP perturbs both facies and net to gross per region. MP-PPM perturbs facies only per region and SP-PPM perturbs facies for only one region.

	Number of outer iterations to converge										Average	Average
Realization	1	2	3	4	5	6	7	8	9	10	Iteration	Flow Sim.
MP-wP	8	14	6	14	27	22	14	3	19	9	13.6	95.2
MP-PPM	7	2	6	51	35	16	14	27	51	33	24.2	169.4
SP-PPM	30	>100	43	55	>100	>100	19	46	>100	30	62.7	438.9

When both facies proportions and locations are regionally perturbed (MP-wP), history matching converges much faster. On average MP-wP takes ten fewer iterations than the MP-PPM and almost fifty fewer than the SP-PPM. The amount of high permeability facies in a region has a large impact on water cut from that region, and the current method directly perturbs that amount per region. The other perturbation methods do find realizations that have different amounts of high permeability facies in different regions, but those method can only blindly search for the correct proportions. Searching for realizations becomes much slower than making the desired changes explicitly.

This example uses a small model with relatively simple geology, but even for this example, some of the realizations from the comparison cases could not converge. For large reservoir models with complicated geology, achieving a history match will be even more difficult. Regionally perturbing parts of the geologic model such as the local proportion will allow the probability perturbation history matching technique to be practical for some field scale applications.

Next, we check if the method reproduces the LP of the reference. Table 5.2 displays the percent of sand for each region in the matched model and the global sand proportion. The next to last row displays the average proportion for the ten realizations and the last row has the reference proportions. Although the initial model for all realizations start with the same proportion in each region, the history matched models have vastly different proportions per region.

Table 5.2: Percentage of sand in each region for matched realizations for the proposed method (MP-wP).

		Percent of Sand in Matched Model				
		Region 1	Region 2	Region 3	Region 4	Global
Realization Number	1	25.9	32.8	27.4	58.8	36.3
	2	18.4	23.4	26.5	54.9	30.8
	3	21.9	27.0	29.7	60.0	34.7
	4	23.1	26.2	14.9	60.0	31.0
	5	12.1	13.6	10.0	60.0	23.9
	6	22.7	21.8	14.2	60.0	29.7
	7	19.5	27.1	23.1	60.0	32.4
	8	28.8	19.2	24.9	52.5	31.3
	9	17.4	15.5	23.9	60.0	29.2
	10	16.1	13.9	11.5	58.9	25.1
Average NTG		20.6	22.0	20.6	58.5	30.4
Reference NTG		15.0	25.0	35.0	45.0	30.0

Table 5.2 shows that on average the method does not reproduce the regional proportions. Regions 1 and 4 over-predict the LP while Region 3 under-predicts it. The problem of reproducing the LP per region is not a consequence of the current method used to perturb the LP. For the case where the LP was not perturbed (MP-PPM), the regional LP of the reference was not reproduced either. For the MP-PPM the average matched LP were similar in value to the MP-wP method (Table 5.3).

Table 5.3: Percentage of sand in each region for matched realizations for the regional facies only method (MP-PPM).

		Percent of Sand in Matched Model				
		Region 1	Region 2	Region 3	Region 4	Global
Realization Number	1	12.0	12.1	30.4	60.2	28.7
	2	7.4	26.1	18.7	55.4	26.9
	3	13.8	16.9	10.7	53.4	23.7
	4	20.6	16.3	8.8	65.5	27.8
	5	24.6	15.9	18.5	48.9	27.0
	6	27.3	16.8	14.6	54.1	28.2
	7	36.1	36.0	30.6	48.3	37.7
	8	14.1	25.1	10.4	58.9	27.1
	9	18.0	25.6	28.6	46.1	29.6
	10	18.9	40.3	31.3	64.7	38.8
Average NTG		19.3	23.1	20.2	55.6	29.5
Reference NTG		15.0	25.0	35.0	45.0	30.0

The fact that the proportion is not reproduced regionally is caused by two similar but distinct reasons both relating to the location of the facies. Recall that the water cut from the reference model was strongly influenced by how much of the high permeability facies was in the direct path between the injector and producers. Therefore, the average LP obtained by history matching is more likely representing the LP in the direct path between the injector and producers, conceptually represented in Figure 5.7.

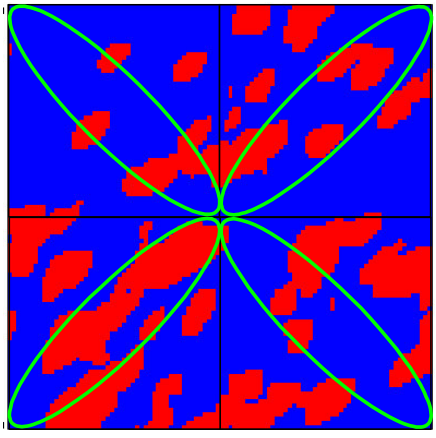

Figure 5.7: Reference model highlighting the direct path between the injector and producers.

The proportion of facies outside these areas will have only a small effect on flow, but they will significantly change the regional calculation of LP. The second reason for the mismatch of LP is a function of the connectivity of the facies. To better understand this, consider two cases. The first has a small proportion of high permeability facies, but they are connected such that they form a direct link from the injector to the producer. The other has a higher proportion of facies, but they do not connect to each other at all. The first case will have earlier breakthrough and higher water cuts than the second; however if only the net to gross ratio information was available, one might expect the second to have higher water cut.

The proportions and locations of facies are two dependent parameters; hence, multiple combinations of these two parameters can produce a history match. Since both properties tend to have high uncertainty, neither can be determined exactly. Therefore, when reservoir predictions are needed, multiple matched realizations with different LP and facies distributions can be used to better understand and better define the uncertainty in future production.

An interesting outcome of this method is that the global proportion is predicted unbiasedly (i.e on average it matches the reference). This occurred for both cases presented even though the correct proportion was never specified in either method. Moreover, this result was not confined to these two examples. Six additional sensitivity cases (including different convergence criteria, different initialization and different random numbers) confirmed the unbiasedness of the global proportion, but the local proportion was never exactly reproduced.

An intuitive explanation is provided for the unbiased prediction of the global proportions. Because of the containment boundaries of the simulation model, all of the injected water either is produced through a well or remains flowing in the model. Darcy's Law describes this type of flow problem; thus, the flow of fluid is constrained by the permeability of the system. Since a specific amount of fluid is being injected and a

specific amount of fluid is being produced, the permeability averaged over the entire reservoir model (a direct function of the proportion of facies) needs to be relatively constant in each history-matched realization for the injection and production data to match.

An algorithm has been developed to regionally perturb both facies locations and proportions without creating the geological discontinuities. For the synthetic case, this resulted in a more efficient history match than perturbing facies locations alone. In the next section, the method will be implemented on a real reservoir.

5.4 North Sea Case Study

Production data from a North Sea hydrocarbon reservoir is used to improve the description of the reservoir's simulation model. In the reservoir, a number of lenticular calcite bodies exist that have very low permeability. It is likely that these bodies inhibit vertical fluid flow, and it is possible that they could funnel flow in certain horizontal layers. The location and proportion of the bodies are generally unknown. This work attempts to stochastically model these bodies, and then perturb their local proportion and position so that the reservoir model more accurately reproduces the field production data.

Geologic Description

The reservoir has a moderately faulted layered structure. The faults are primarily vertical and are generally not fully sealing. There are four segments separated by the faults; however, there is pressure communication areally among all of the segments. Figure 5.8 shows the top of the structure and the general location of the four segments. The reservoir is at a depth of 2500 m and is about 230 m thick.

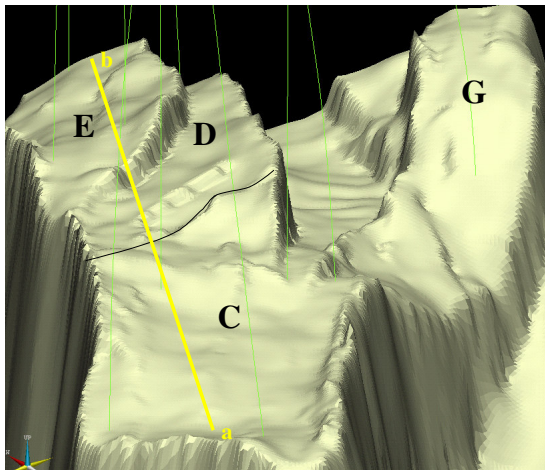

Figure 5.8: Top of the structure is shown along with the four segments. The yellow line a-b corresponds to the subsequent cross section map (Figure 5.9).

Vertically, there are five major formations. An impermeable shale layer (Horizon 2) isolates the top formation from the bottom three formations. Although the furthest right

fault in Figure 5.9 appears to have open communication between Horizon 1 and Horizon 3, the fault shear is likely sealing as there is no indication from production data that communication exists between the two layers.

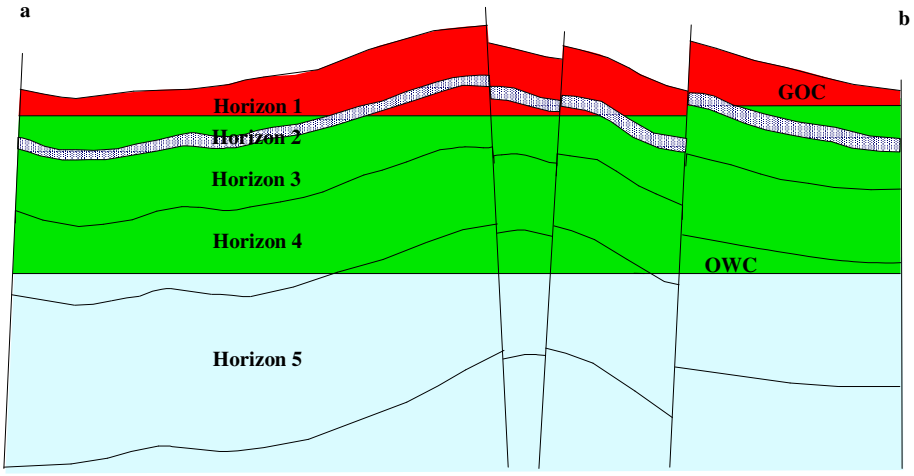

Figure 5.9: Vertical cross section showing five horizons, oil-water contact, and gas-oil contact.

A gas cap exists in the reservoir and makes up a large portion of the top formation. Below the shale layer, there is only a small amount of free gas present. The water zone accounts for the principle amount of Horizon 5 (Figure 5.9).

Twenty-two wells delineate the reservoir. There are 8 injection wells that vertically intersect the reservoir. Five injectors have been strictly under water injection, while the other three have injected both water and gas. Most of the 14 production wells are either horizontal or deviated (10 wells). Figure 5.10 shows where the wells intersect the top of the reservoir. Wells with letter I are injection wells; wells with P are producers. Wells I-1, I-3 and I-4 have injected both water and gas. Figure 5.10 is rotated clockwise 90º compared to Figure 5.8.

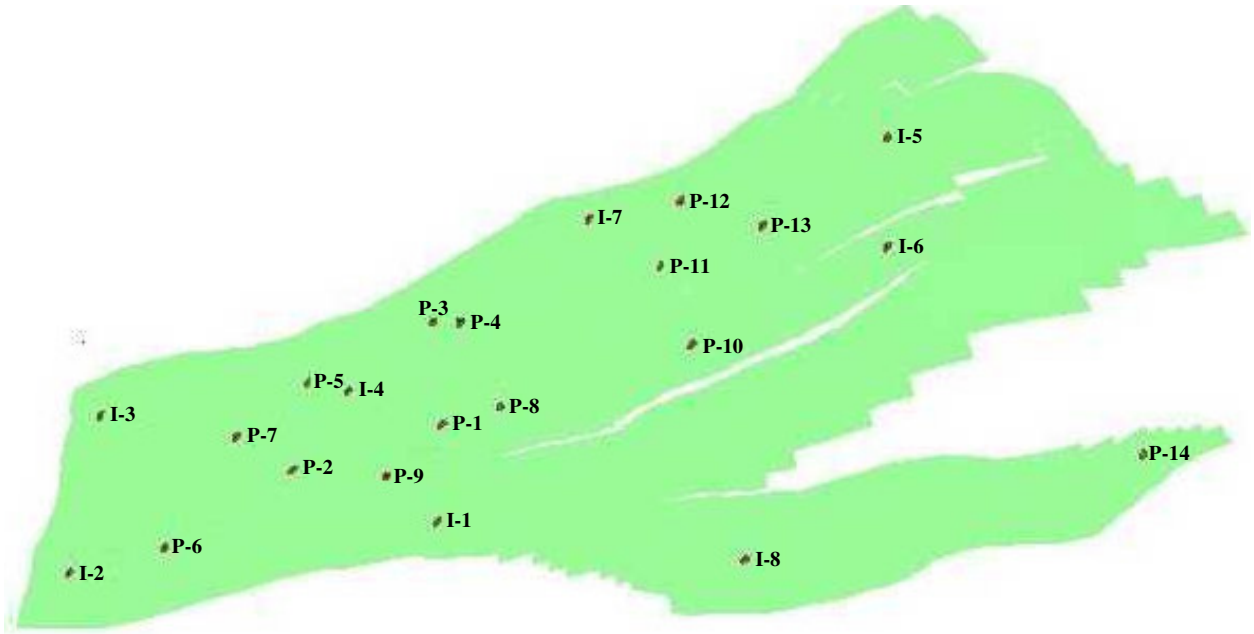

Figure 5.10: Locations where wells cross the top of the reservoir.

A significant number of very low permeability nodules are found in the reservoir. They were created by the diagenesis of calcite and tend to have a thin lenticular shape. They are at most a meter thick and typically have an areal extent of a few meters to tens of meters. Where clusters of these bodies are found, they can have a large detrimental affect on vertical permeability. By inhibiting vertical flow, they may also cause channeling of flow in the horizontal direction. Figure 5.11 shows three separate zones in which bodies are thought to occur with enough quantity to alter vertical flow. Observed well log and core data suggests that the lower section of Horizon 3, the top of Horizon 4, and the bottom of Horizon 4 have a large amount of calcite bodies.

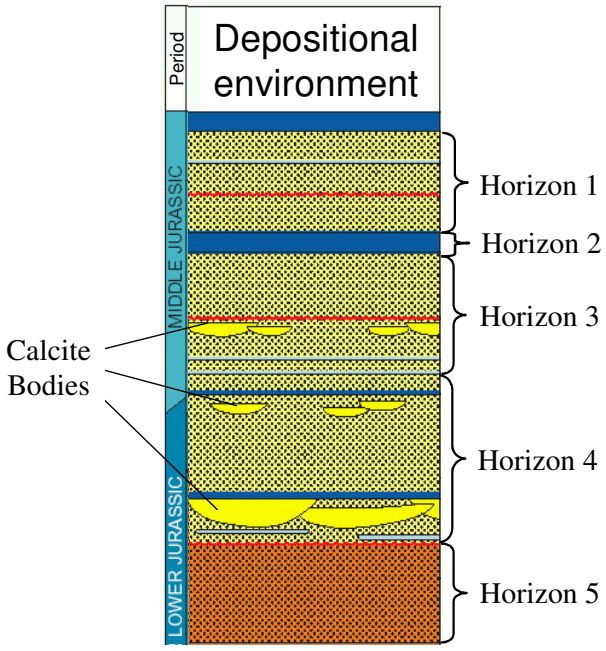

Figure 5.11: Vertical locations of calcite bodies.

The main goal of this work is to produce a simulation model that matches the historical production data and honors the geology of the reservoir. It is apparent that the low permeability calcite bodies can have an influence on fluid flow and are to be included in the model. Well information is too sparse to adequately constrain the proportion of bodies, and because they are so thin, seismic is unable to resolve them. Since the location and proportion of the bodies is highly uncertain, they provide an important model parameter to modify during history matching. The probability perturbation method where facies are perturbed locally is well suited to solve this type of history matching problem.

Reservoir Simulation Model

A non-history-matched (NHM) reservoir model was provided by a major oil company, and most of the data that is discussed in the following section comes from that model. The reservoir model is a structured stratigraphic model with 39 cells in the x-direction, 98 cells in the y-direction, and 41 cells in the z-direction. The z direction has been refined

from a previous model that has 26 z-direction blocks. The 41 grid layers delineate five horizons. There are just over 150,000 total gridblocks in the model and about half of them are active. The active gridblocks are displayed in Figure 5.12. The vertically refined layers appear as thick black layers in the figure.

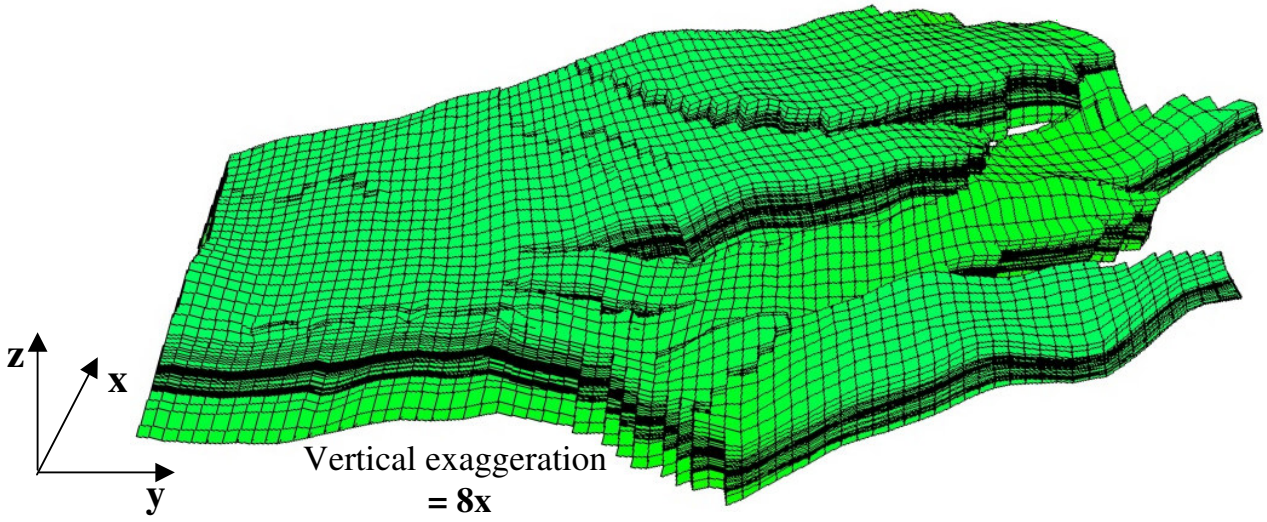

Figure 5.12: Grid for the reservoir model.

There are numerous faults in the model, and Figure 5.12 reveals that some of the displacements are rather significant. Most of the faults tend to have reduced permeability, though none of them are impermeable over the entire fault. Areally, there are four separate segments, and they are depicted in Figure 5.13. The individual segments correspond to different fault blocks; however, horizontally none of the segments are entirely isolated because none of the faults are completely sealing. Vertically, there is isolation because the second horizon is an impermeable shale layer that separates the top horizon from the lower three horizons (Figure 5.9).

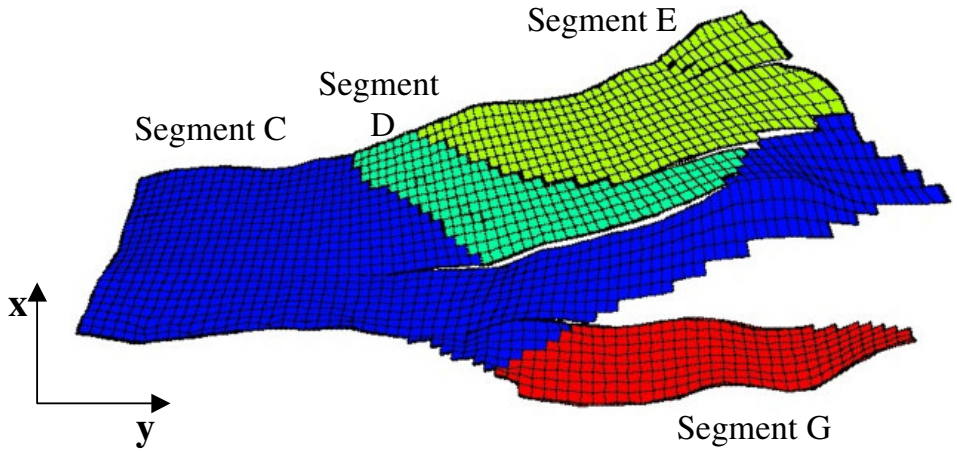

Figure 5.13: Four areal segments.

There is an aquifer present and the water contact in the model varies from 2575 to 2694 meters in depth depending on the segment and whether above or below the shale barrier. There is also a gas cap in the reservoir, and above the shale barrier, it is rather large. However, below the shale layer, where most of the oil reserves are, there is only a small amount of free gas present.

The porosity and horizontal permeability initially provided in the model are shown in Figure 5.14.

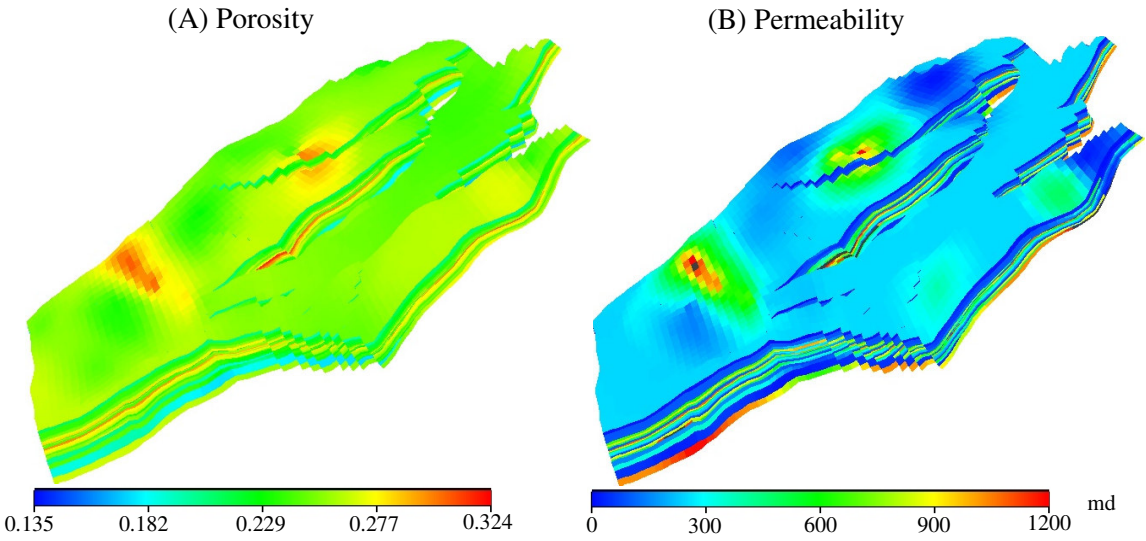

Figure 5.14: Porosity and permeability for initial simulation model.

The porosity is calculated using a kriging algorithm. From these two graphics, it is apparent that a linear transformation is used to calculate the permeability from the porosity. The x and y direction permeability are equal, but the z direction permeability has been reduced for different layers and different areas of the reservoir. Porosity and permeability vary considerably layer to layer, but within each layer little variation exists.

The initial pressure in the reservoir was around 270 bars. There are 18 separate relative permeability regions where the values are calculated from Corey type curves with different exponential coefficients and end points. The 18 water-oil relative permeability curves are displayed in Figure 5.15.

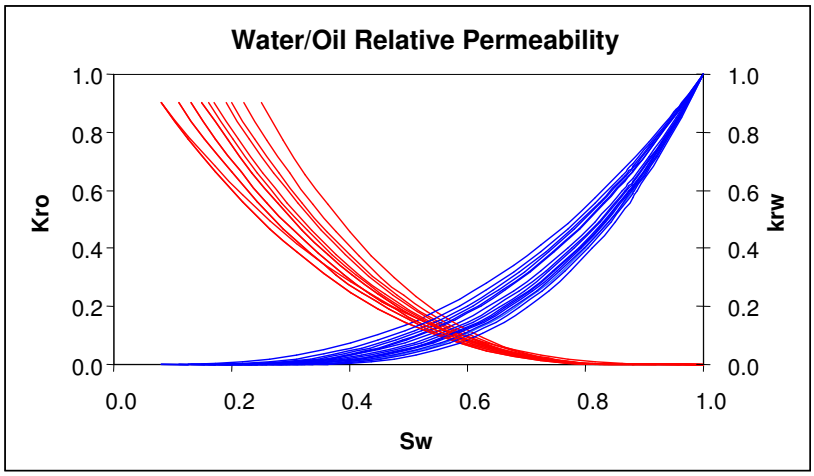

Figure 5.15: Relative permeability curves.

In addition to the faults, there are also calcite bodies in the reservoir that influence fluid flow. Since the bodies are relatively thin, they are given no vertical thickness in the simulation model, and they are modeled as z direction transmissibility barriers. Gridblock containing a body or a cluster of bodies will get a reduced or zero transmissibility. So for this case, LP is not a volumetric value of proportion, but rather the proportion of gridblocks that have a reduced vertical permeability due to the presence of the calcite bodies. For example if the LP is 30%, this does not mean that 30% of the volume in a region is calcite. It means that 30% of the gridblocks in the region have reduced vertical permeability.

Production Data

The flow data used to match is the water rate. In the model, wells have fixed liquid rates; hence if water rates are correct, oil rates are also correct as are the water cuts. The available pressure data comes from RFT pressures taken at 12 wells when they were first drilled. For each well there are numerous measurements taken throughout the reservoir column. The objective function is simply defined as the square difference between the simulated production and measured production data. Depending on when a well came on-line and went off-line there is between 17-44 months of production for each well.

Equal weight is given to each point of data, not to each well, so theoretically, a well that is producing for 40 months will have twice the influence of one that has been producing only 20. Likewise, a well with RFT measurements taken over 200 m will have twice the influence as a well where only 100 m were measured; however, most wells had similar distance intervals measured. The pressure data is weighted such that it comprises about 40% of the total objective, and water rate data account for the other 60%.

The water rates from the six production wells in the D and E segments are displayed in Figure 5.16 as the solid blue lines. The water rates from the provided non-history-matched (NHM) model are also shown in the figure as red lines. Although continuous lines are shown, the data used to match is actually monthly data. Figure 5.16 only shows the water rates for six wells, but oil, water and gas rates for all of the producers along with injection rates for all of the injectors are given in Appendix D.

Of the fourteen production wells, two are of primary concern, P-3 and P-4. For well P-3 the water breaks through to the producer too early and too much water is produced in the model. Although the well is mainly in the D segment, the water that is produced primarily comes from the E segment. It crosses the fault along the bottom where the fault is not sealing, and due to the pressure depletion caused by P-3, water travels up to the well.

Conversely, for well P-4 the water breaks through too late in the simulation model. This well is also in the D segment, but most of the produced water at this well comes from the aquifer portion of the D segment.

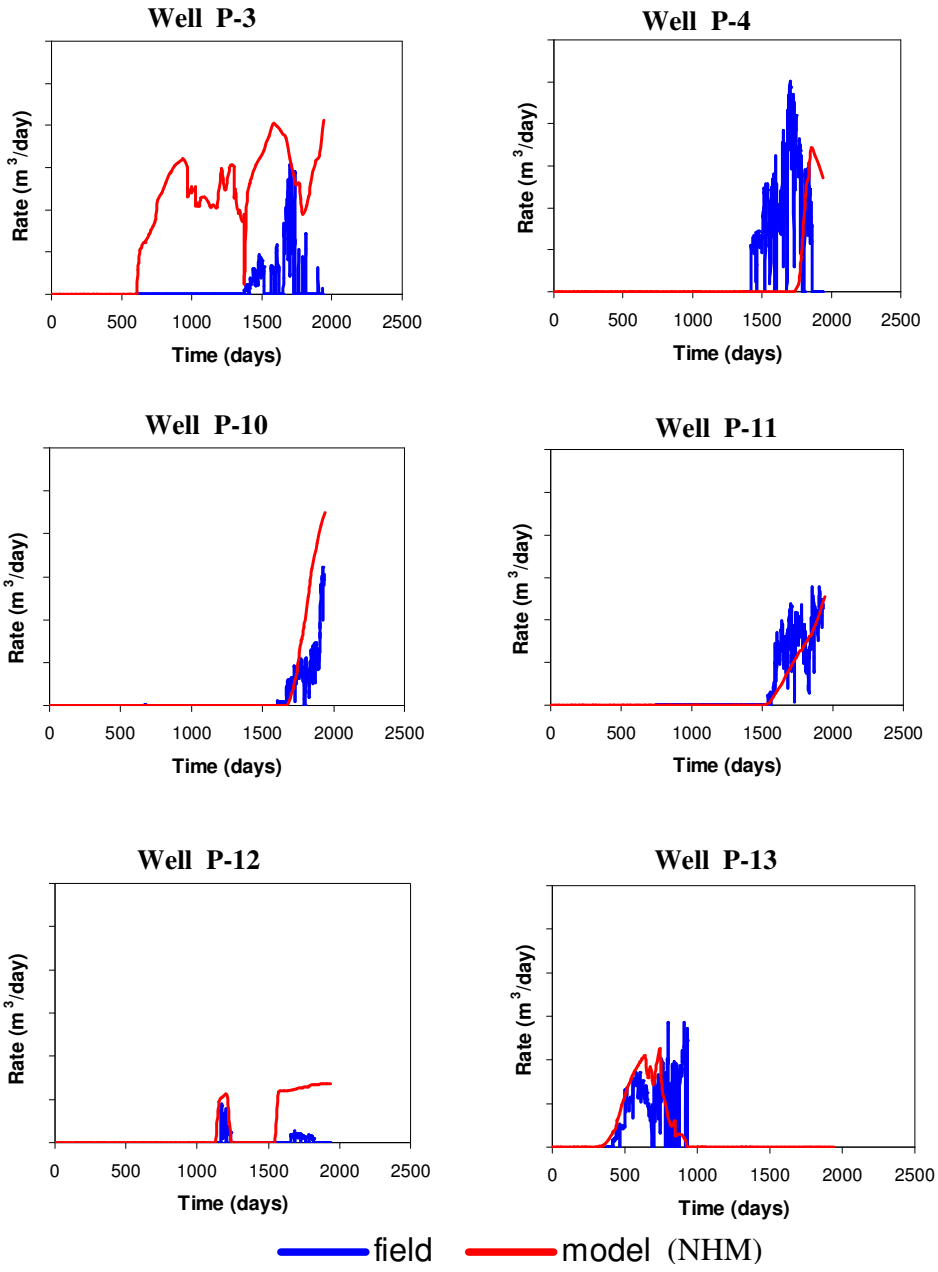

Figure 5.16: Water rate for the six production wells in the D and E segments and the modeled water rate from the NHM simulation

Pressure versus depth data is available for four wells in the D and E segments (Figure 5.17). Pressure from the NHM model is also plotted on the graphs. The original match is good for wells P-13 and I-7, but it is not as good for wells I-5 and I-6.

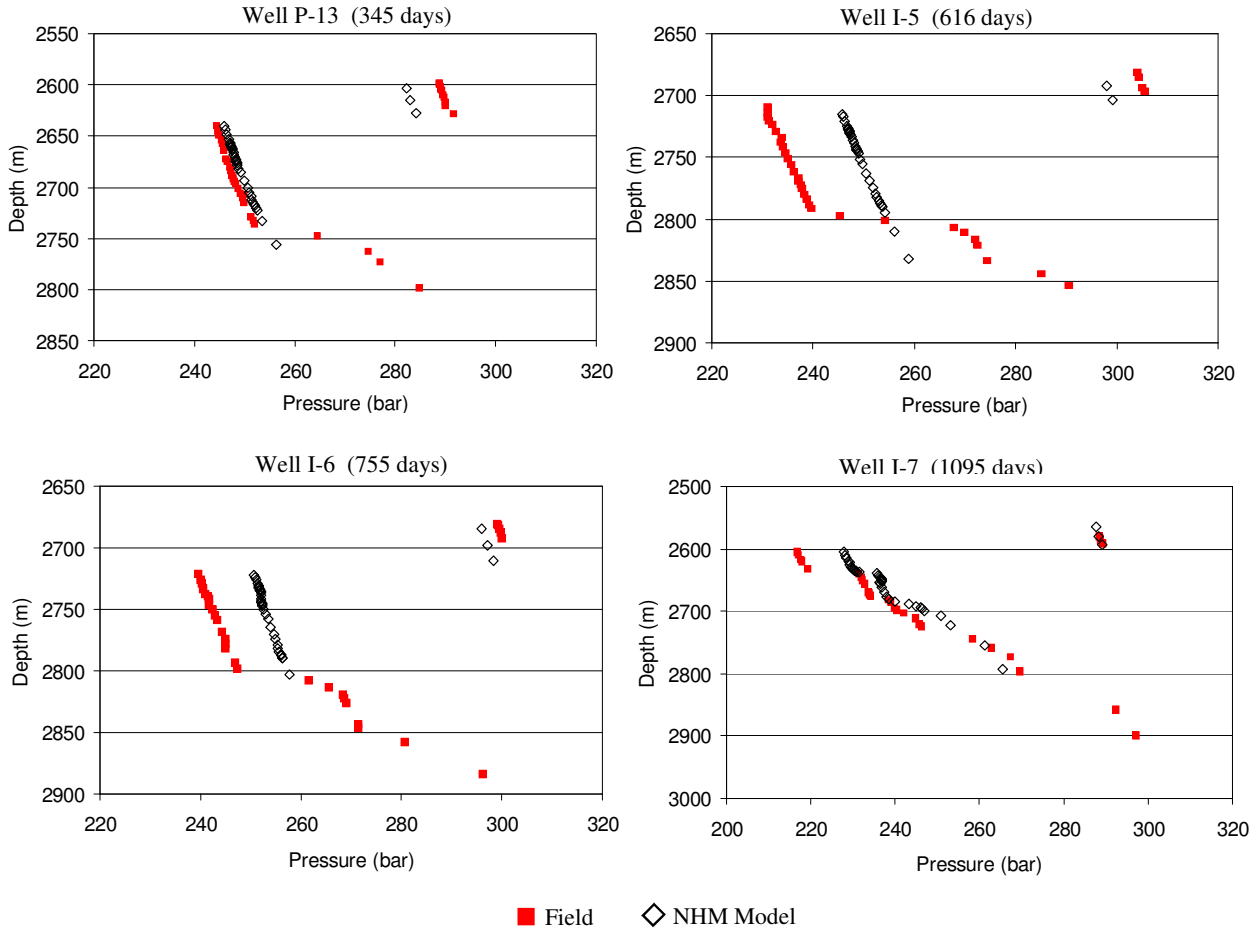

Figure 5.17: Pressure for four wells.

The pressures in Figure 5.17 above 300 bars that are at shallow depths are pressures in Horizon 1. These pressures are higher than the rest of the horizons, and they remain higher throughout the production period. This clearly demonstrates that Horizon 1 is isolated from the lower three formations.

Reservoir Regions

To perform history matching with the probability perturbation method, a method for defining regions in the reservoir is required. Streamlines are well suited for the job because they directly show the flow paths by which fluid enters a production well. These paths identify the gridblocks that, if changed, will have an obvious impact on production from a well. All blocks hit by the set of streamlines entering a well define the "drainage zone" for that well. The various drainage zones will define the geometry of the regions used for history matching. As the model is perturbed and new realizations are generated, the facies distribution will change. Consequently, the drainage area of the production wells will be different for different realizations, and the streamline-defined regions will have to be updated. The region definitions for one realization are displayed, Figure 5.18.

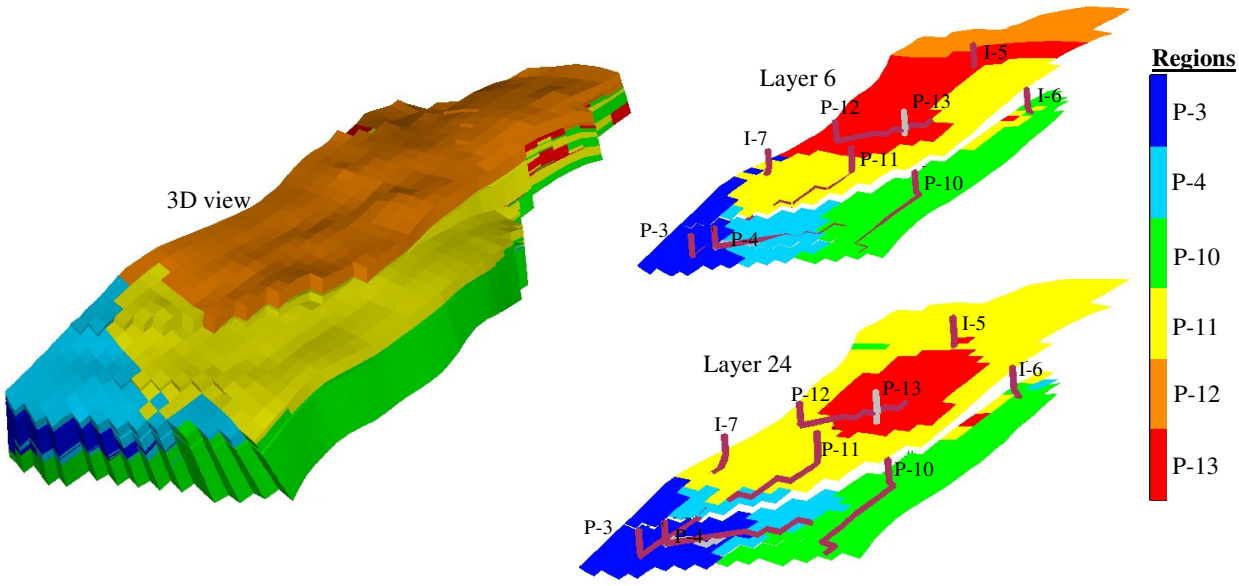

Figure 5.18: Streamline defined regions in 3D and for two selected layers.

The streamline simulation, 3DSL, (StreamSim, 2002) is needed only to define regions (i.e. find the main gridblocks that influence flow at a well); not for deducing water and oil rates. Thus, the actual streamline simulation model is different than the finite difference model. The grid dimensions, porosity and permeability are the same for both models, but

the streamline model is incompressible; its gas phase is not treated, and simplified relative permeability curves (x-curves) are used in the streamline model. Nonetheless, the biggest difference between the streamline and finite difference models is how the wells are modeled. In the reservoir (and the finite difference model), wells come on line at different times over the 5 ½ years; however in the streamline simulation model, all wells are included at the initial time. This is necessary to ensure that each region contains gridblocks that are most consequential for flow to the well in that region. In the NHM model, the wells have fixed liquid rates that are constant over the life of the wells, so those values are used in the streamline model (e.g. P-3 = 6000 m^3/day, P-10 = 7500 m^3/day, I-7 = 13000 m^3/day). Because streamlines change very little when the well boundary conditions are constant and the streamline simulation is only used to define regions, the streamline simulation only needs to be run for a short duration (10 days).

Geostatistics Model

Calcite concretions occur in a number of reservoirs, but perturbing calcite bodies in a stochastic and geologically consistent manner is difficult. A sequential simulation technique (truncated Gaussian) was used to model these types of deposits for a Wyoming reservoir (White, *et al.*, 2003). However, this technique created bodies with more areal continuity than we expected for the North Sea reservoir. Therefore we used a multiple point geostatistics technique, snesim, (Strebelle, 2002) to generate the calcite bodies. The wells were used to hard condition the calcite bodes at the well locations. The training image used for the current work is shown in Figure 5.19, and it only needs to be 2D because these bodies are modeled with no z dimension. 30% of the cells in this training image are low permeability bodies.

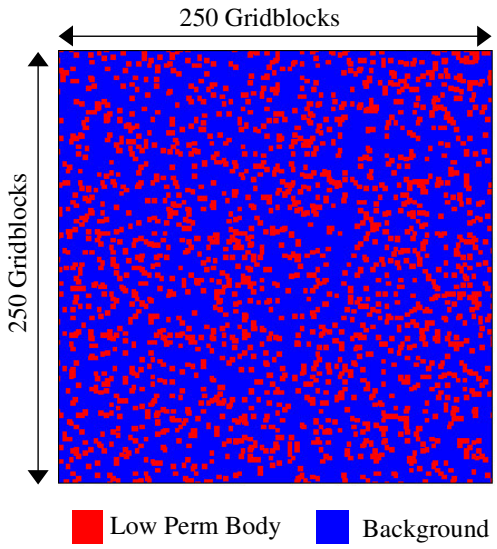

Figure 5.19: Training image.

Information on the size and the distribution of the calcite bodies is limited. Although the size and shape of the bodies is not precisely understood, we are assuming that where clusters of bodies occur, their effects are over a minimum area of 0.04 km^2. On average the gridblocks have a length around 100 m in the x and y directions, so the size of the bodies in the training image is typically two gridblocks square. The most important characteristic of these bodies is their local clustering. If a large amount of bodies cluster, they form a vertical flow barrier and a horizontal flow conduit. The training image in Figure 5.19 does not contain the information on clustering; this information is included by regional variation of the proportion of bodies in the geostatistical simulation that is calculated from Eq. (5.1).

Once the bodies are generated, they are mapped to the simulation grid where they are applied as transmissibility barriers. Where bodies occur, the transmissibility is set equal to zero. The locations and proportions of bodies are perturbed until a history match is achieved. All other aspects of the simulation model remain the same as the original NHM simulation model. In the pre-history matching process, the porosity and permeability, relative permeability and fault transmissibility were all examined to see if a

history match could be achieved by perturbing those parameters. Figure 5.20 the range of uncertainty for these parameters for well P-3.

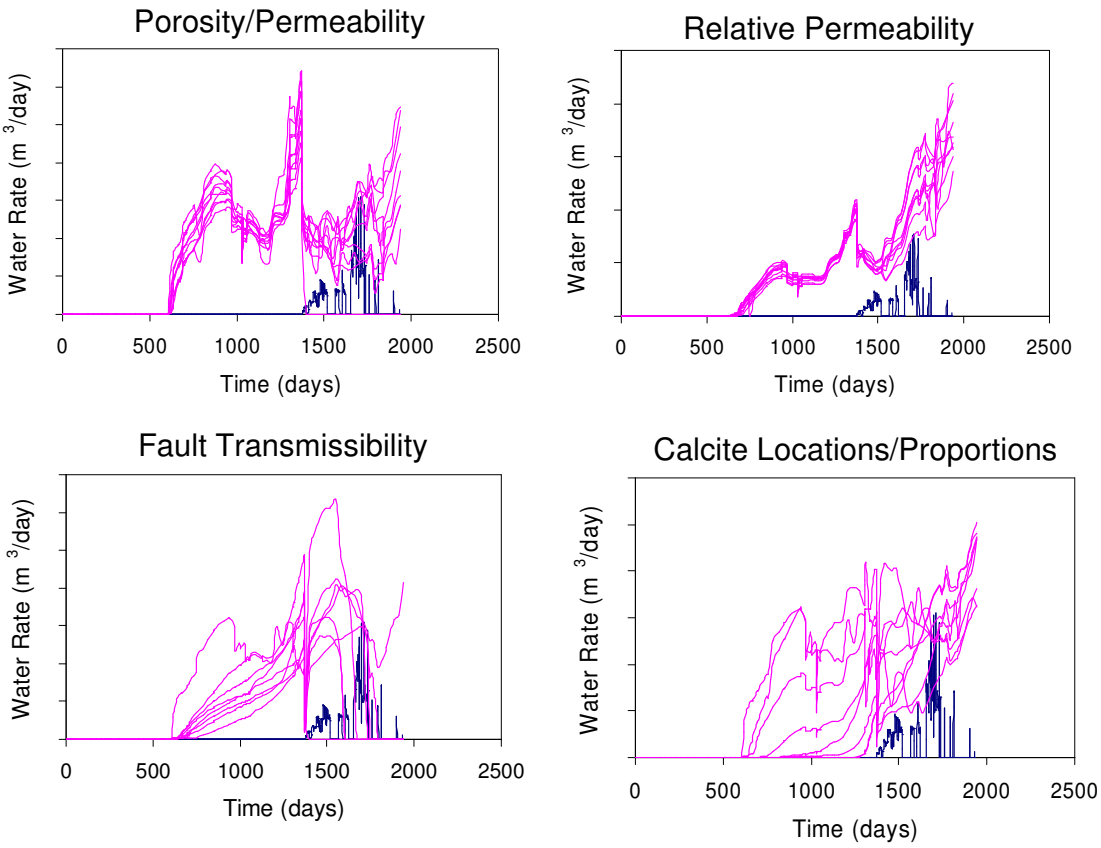

Figure 5.20: The range of production data for four different parameters tested in the pre-history matching process (sensitivity analysis).

While there are some differences in the range of these parameters, none of them would be able to achieve a match by only perturbing those properties, which is why the calcite bodies are perturbed to history match.

Although some calcite bodies are present throughout the entire reservoir, there are three vertically defined zones that are thought to have a higher concentration of bodies as observed from well log and core data. Zone 1 is located near the bottom of Horizon 3 and accounts for layers 11-17 in the simulation model. Zone 2 is at the top of Horizon 4

and corresponds to layers 18-25, and the third zone is at the bottom of Horizon 4, layers 32-37. Figure 5.21 shows where the zones are in the simulation model. Initially bodies will only be located in these three zones in the history matching process.

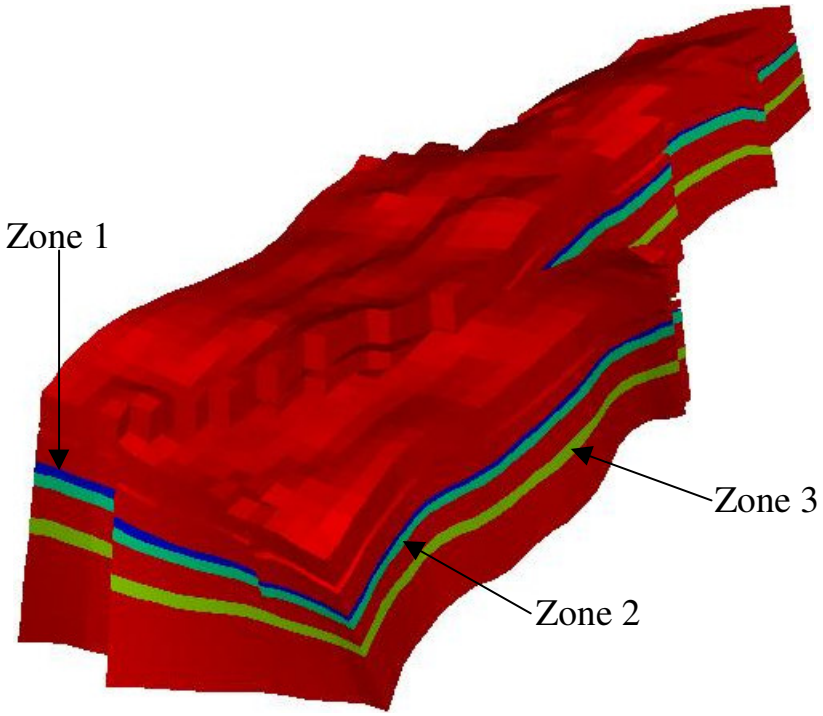

Figure 5.21: Three zones that are thought to have a significant proportion of calcite bodies.

DE Segment

During the first phase of this study, only a part of the full field model is considered. A reduced model was selected in an attempt to speed up the development of the history matching method. The full field model took approximately 1-½ hours to complete on the available computers, whereas the reduced model took around ½ hour. More time could be focused on creating and improving the history matching algorithm and less time waiting for simulations to finish. Once the method is understood on the smaller model, it is expanded to the full model. This initial study area is termed the DE segment, and it contains 6 production wells and 3 injection wells (Figure 5.22).

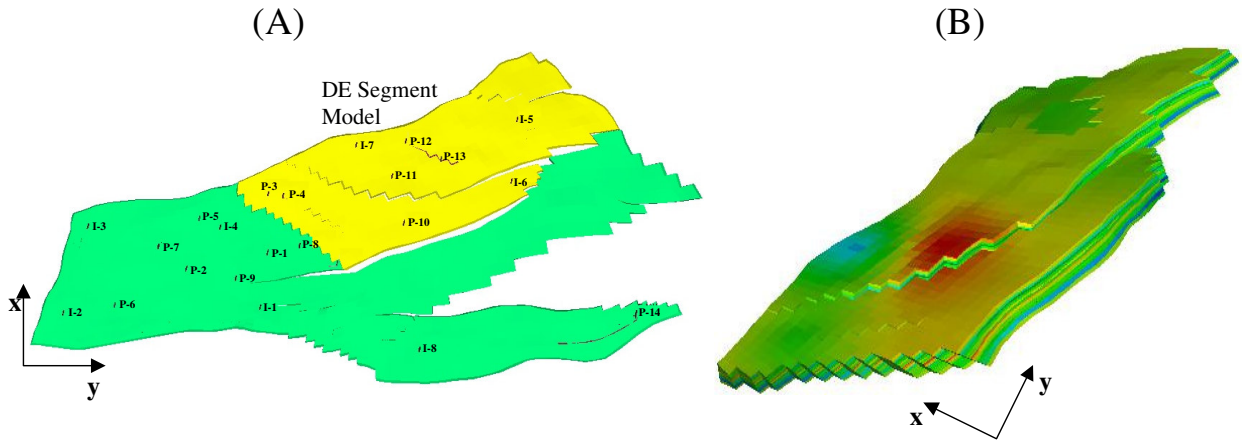

Figure 5.22: (A) Plan view and (B) 3D view of DE segment model.

To achieve the reduction in modeling time, a flux boundary model was created for the smaller section of the reservoir (i.e. the DE segment). During one full field simulation, flows across the boundary of the DE segment are written to a file. This file is read during the reduced field runs to reproduce the appropriate boundary conditions (Schlumberger, 2003). Perturbations can be performed on the reduced model until a history match is achieved for the wells inside the smaller section.

Results – DE Segment

For the first history matching run, bodies are included in all three vertically described zones and all six streamline-defined regions. The location of the bodies as well as the local proportion (LP) is perturbed per region. The LP are perturbed with basically the same method used for the synthetic case. If the model is over predicting the water rate for a region, the LP for that region is increased and vice versa. This is slightly different than the original development because the bodies for the synthetic case are high permeability, and in the current case, they are low permeability bodies. The initial LP for all regions is 30%, and it is maintained between 10% and 60%. If a perturbation causes the LP to go outside these bounds, then the LP is reset to the boundary value. The specifications are presented in Table 5.4 for all history matching runs that are performed on the DE segment.

Table 5.4: Specifications for history matching runs.

Run #	Run Specifications	Figures
NHM	• provided model – non-history matched • bodies only placed deterministically from well log information	Figures 11, 12
Run 1	• bodies in grid layers 11-25, 32-37 • range of proportions 10% - 60% • local proportion (LP) perturbed based on over/under predicting field data • initial r_{Dk} = 0.50; initial LP = 30% • sensitivities include (a) bodies only in grid layers 11-25 (b) bodies only in grid layers 32-37	Figure 15
Run 2	• bodies in grid layers 26-28 - Only Region P-3 • fixed proportion = 90% • no perturbations	Figure 17
Run 2a	• bodies in grid layers 32-37 - Only Region P-3 • fixed proportion = 90% • no perturbations	Figure 18A
Run 2b	• bodies in grid layers 26-28 - Only Region P-3 • fixed proportion = 60% • no perturbations	Figure 18B
Run 3	• bodies in grid layers 11-37 • range of proportions 1% - 99% • local proportion (LP) perturbed based on increasing/decreasing and keeping the best matched in each region for 3rd flow simulation • initial r_{Dk} = 0.75; initial LP = 80% for region P-3 & 20% for other regions	Figures 20, 22

For Run 1 some improvements, (e.g. well P-4), are observed; however, most of the water rates either did not improve over the NHM model (Figure 5.16), or they got worse (Figure 5.23). The breakthrough time for well P-3 continues to be over two years too early and well P-11 does not match as well as the NHM model. Additional cases (not shown) are completed where the bodies are only placed in Zone 3, or only placed in Zones 1 and 2. These cases show some minor regional improvements, but the breakthrough time for well P-3 remains a problem; thus, a different approach is attempted.

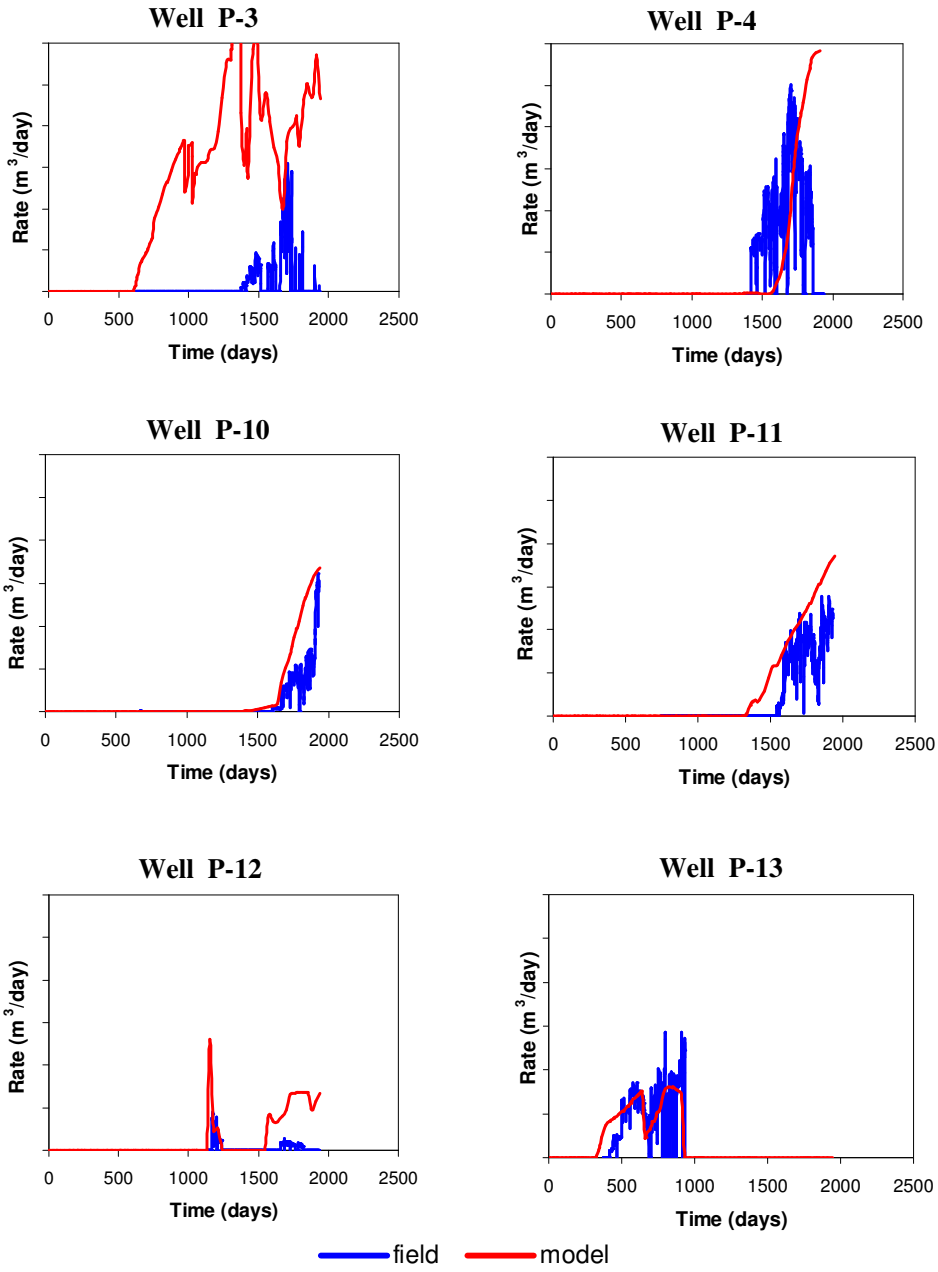

Figure 5.23: Water production from field and model .

The vertically impermeable bodies are placed with a proportion of 90% in the three layers that are directly below the lowest perforations for well P-3 (Figure 5.24).

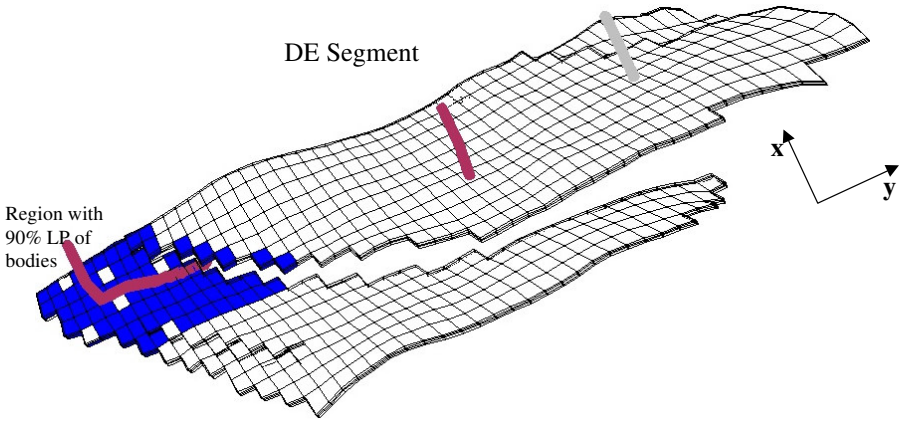

Figure 5.24: Area where calcite bodies are placed (Layers 26-28).

Recall that this does not mean there is 90% calcite in this region, only that 90% of cells in this region have their vertical permeability affected by the bodies. The bodies are only put in region P-3; the other five regions have no bodies. This is called Run 2 in Table 5.4. The water breakthrough time for well P-3 is reasonable and the match for well P-4 improved as well (Figure 5.25). The match for the other four wells did not change significantly.

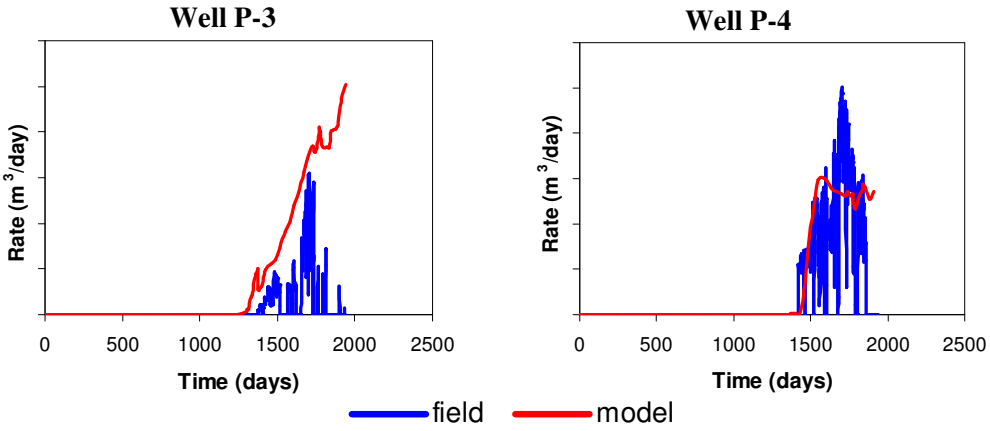

Figure 5.25: Water rates from wells P-3 and P-4 for case with 90% proportion of bodies in the three layers directly below well P-3.

Bodies were not included in these grid layers (layers 26-28) in Run 1. Grid layers 26-28 are in Horizon 4, but they are between the two vertically defined zones (Zone 2 and Zone

3) where calcite bodies were included in Run 1. To check if the large proportion of bodies and not their location caused the improved match, a similar case (Run 2a) is completed where 90% proportion is placed in the original Zone 3, layers 32-37 (Figure 5.26A). This, however, did not have the same beneficial effect as the Run 2 (Figure 5.26C). To test if the location (and not the high proportion of bodies) is the reason for the improved match, a case (Run 2b) where only 60% proportion of bodies is put into the three layers 26-28 (Figure 5.26B). Although the match is better than Run 2a, it does not show the same improvement (especially regarding breakthrough time) as Run 2 (Figure 5.26C). These two examples demonstrate that it is both the locations and proportions of bodies that kept the model from matching the data in Run 1 (Figure 5.23).

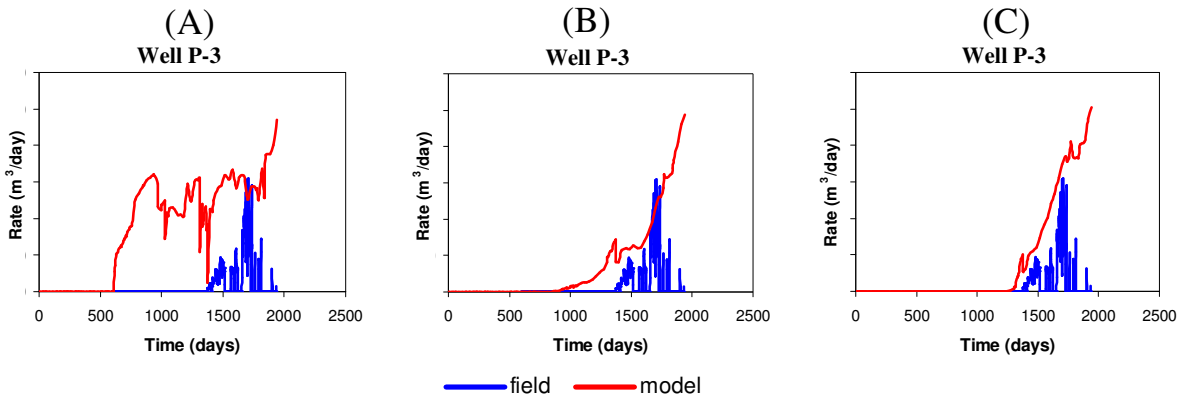

Figure 5.26: (A) 90% LP in Zone 3 of Region P-3 (B) 60% LP in layers directly below P-3 and (C) reference plot from Error! Reference source not found. **(90% in layers directly below horizontal well P-3)**

The match of production data shown in Figure 5.25 is acceptable; however it is not geologically realistic to have the situation illustrated in Figure 5.24: in a small area, 90% of the gridblocks have very low vertical transmissibility; elsewhere, none of the gridblocks are altered. Thus, the conceptual geologic model must be altered so the bodies can be stochastically modeled in a more geologically realistic manner. Although the calcite bodies are present to some extent in most of the reservoir, it was thought that they only existed with enough quantity to affect flow in the three previously defined zones.

The preceding analysis (Run 2) demonstrates that, at least in the region around well P-3, bodies are also a factor in the middle of Horizon 4. Therefore, bodies will be allowed in layers 26-31 (between Zones 2 and 3) in the next history matching runs. Consequently, all gridblocks in layers 11 to 37 will be able to have their z transmissibility reduced due to the low-permeability calcite bodies.

A further modification to the algorithm is necessary with regard to how the LP is updated. In Run 1 the LP of the calcite bodies are perturbed as follows. If the water rate for the model is greater than the observed rate, the proportion of bodies is increased in an attempt to reduce flow to the well, and if the model is under predicting the field data, the proportion of low-permeability bodies is decreased. This works for some wells such as P-3, but for others such as P-10, this logic actually makes the match worse. Because P-10 is completed in the upper layers, an increase in vertical flow barriers allows less injected water to travel down by gravity. In this situation, the presence of calcite bodies "funnels" the water to the well and causes an earlier breakthrough (Figure 5.27). The LP may need to be increased for some regions and decreased for others, and it is not immediately known from production data, which one will lead to a better match.

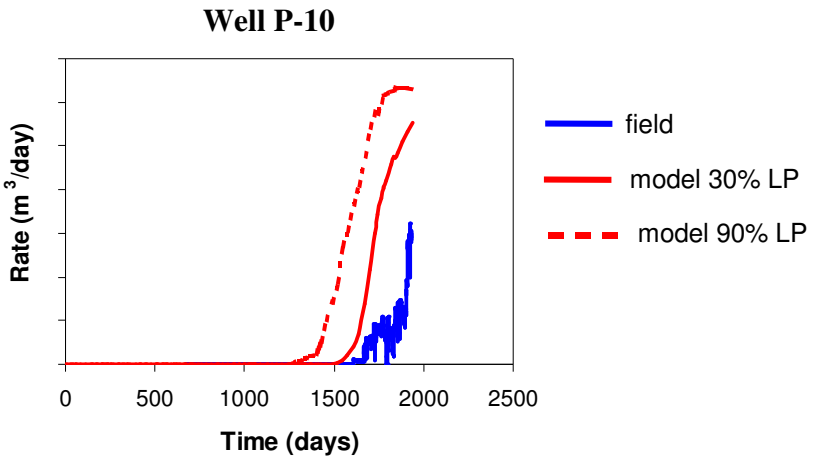

Figure 5.27: Comparing water rate for models with 30% LP and 90% LP.

Therefore, a different method for perturbing LP, similar to one proposed by Kim and Caers (2003) is used. First, the LP for all regions is increased (i.e. i_k in Eq. 1 is +1 for all k) and flow simulation is run. Then, the LP is decreased for all regions (i_k = -1 for all k) and flow simulation is again completed. The regional objectives are compared for the two cases. Some regions may have a better match for an increased LP, while other regions will match better for the decreased LP. On this basis it is decided whether to increase or decrease the LP in a region (i.e. whether i_k should be -1 or +1 for each k). Finally, a new model is created that contains the various regional increased and decreased LPs, and a third simulation is run. The global mismatch will be calculated from this third simulation and compared with the best previous mismatch in the optimization routine.

Other changes to the method allowed a larger space of possible LP values to be searched and the values to be searched in a faster manner. First, the LP was allowed to vary more during each perturbation. Fc is set to 10%, hence the average LP perturbation is around 5%. Second, a set of higher initial r_{Dk} values started each inner iteration (0.75 instead of 0.50). The Simple 2D Example section in Chapter 3 showed that increasing the initial r_{Dk} values increases the initial convergence rate. Third, the search space of the local proportion was expanded from 10% - 60% to 1% - 99%. Finally, the initial LP for region P-3 is 80%, and it is 20% for the other five regions.

Figure 5.28 shows the water rate match achieved for the six wells, and it is rather good for all wells. The initial model is also displayed in the figure. This is not the NHM model; it is the first generated realization before any history matching perturbations are completed.

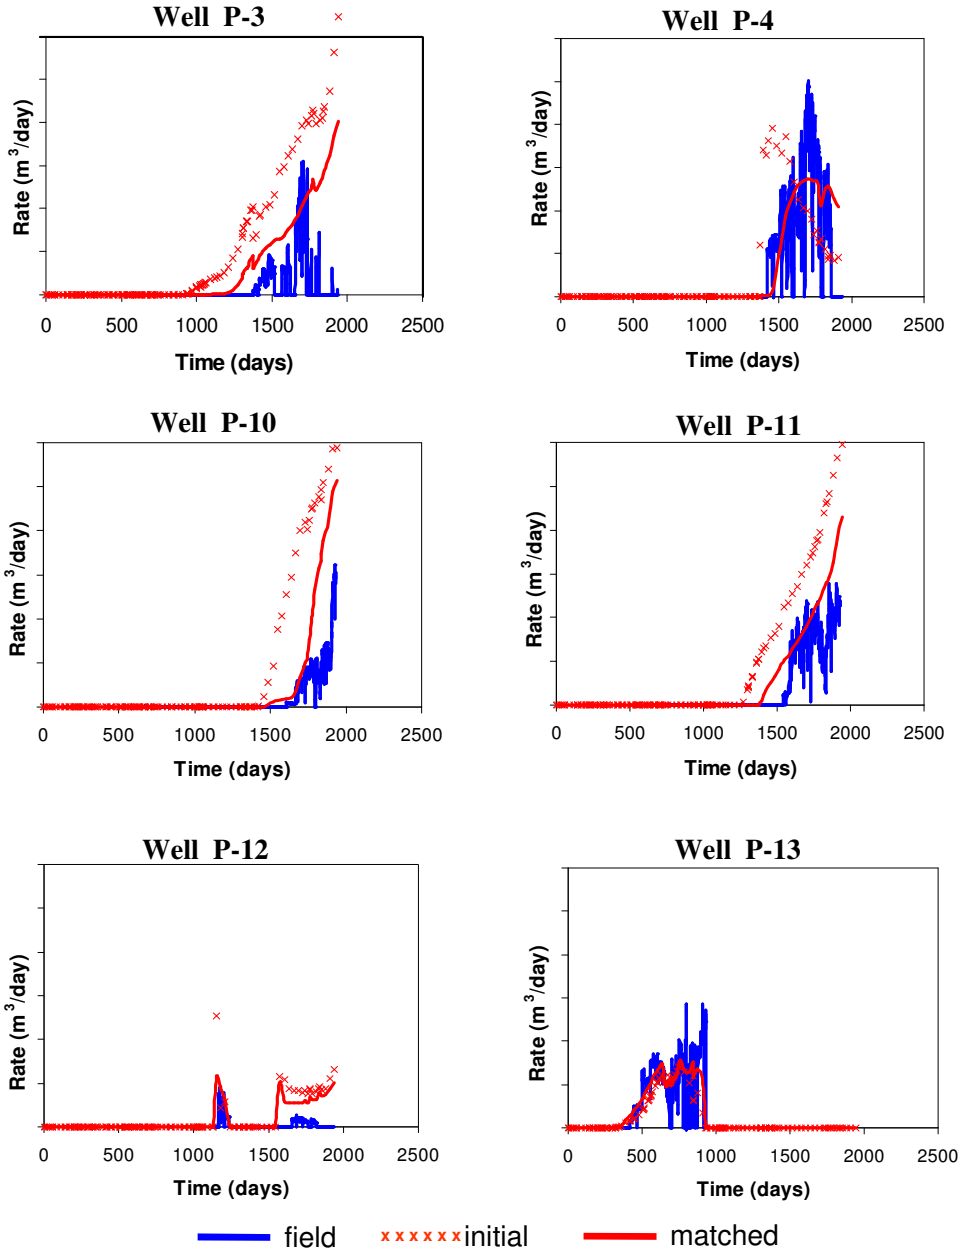

Figure 5.28: Match of simulation model data to field production data for DE segment.

For wells P-3 and P-4, the initial model already displays a vast improvement over the NHM model (Figure 5.16), and the perturbations only further improve the match but to a

lesser degree. This demonstrates a point alluded to earlier: if the parameterization of the model is ineffective or incorrect, it will be impossible to achieve a history match (unless the geologic data is invalidated altogether), however, once the correct parameterization is complete, history matching is straightforward. For this reservoir, once bodies were placed in layers 26-31 and with high proportion in region P-3, history matching went smoothly.

This particular history matching run took only 44 flow simulations to achieve the outcome, and another history matching run that matched almost as well took only 59 simulations to converge to its best solution. The final proportions for the six regions are 53% for P-3 region, 12% for both P-4 and P-10 regions, 1% for P-11 region, 13% for P-12 region, and 22% for P-13 region. The LP for region P-3 was expected to be higher as demonstrated by Figure 5.26; however, it ended up at a much more reasonable value below 60%. This is partially due to the fact that we are searching for a global match and the lower proportion of bodies in P-3 may help other wells match better. Although the breakthrough time is a little earlier for this run than Run 2, the amount of water produced stays lower longer for this run, which helps the match quality.

The locations of the cells where the bodies influence vertical flow are shown in Figure 5.29. While most of the model has relatively few altered gridblocks, the region around P-3 (bottom left of grid) has a slightly higher proportion of bodies.

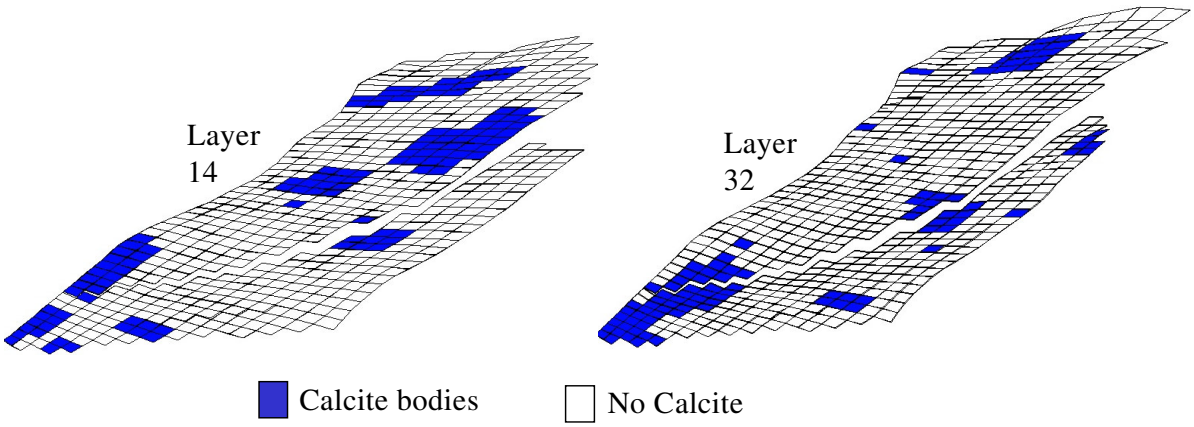

Figure 5.29: Location of calcite bodies in two layers

An adequate pressure match is also achieved for the model (Run 3), but the match is not exact for all wells (Figure 5.30). Well P-13 matches well for both the NHM model and the history matched model. Both wells I-5 and I-6 improve their match over the NHM model. Well I-6 is very closely matched while I-5 only improves its match to about half way between the NHM model and the field data. The pressure match for well 1-7 is worse for the "matched" model than the NHM model, and pressures in the matched model are much higher than the observed pressure. The increase in pressure may have been due to the boundary fluxes used for DE segment model because in the full field model this increase in pressure at well I-7 is no longer observed.

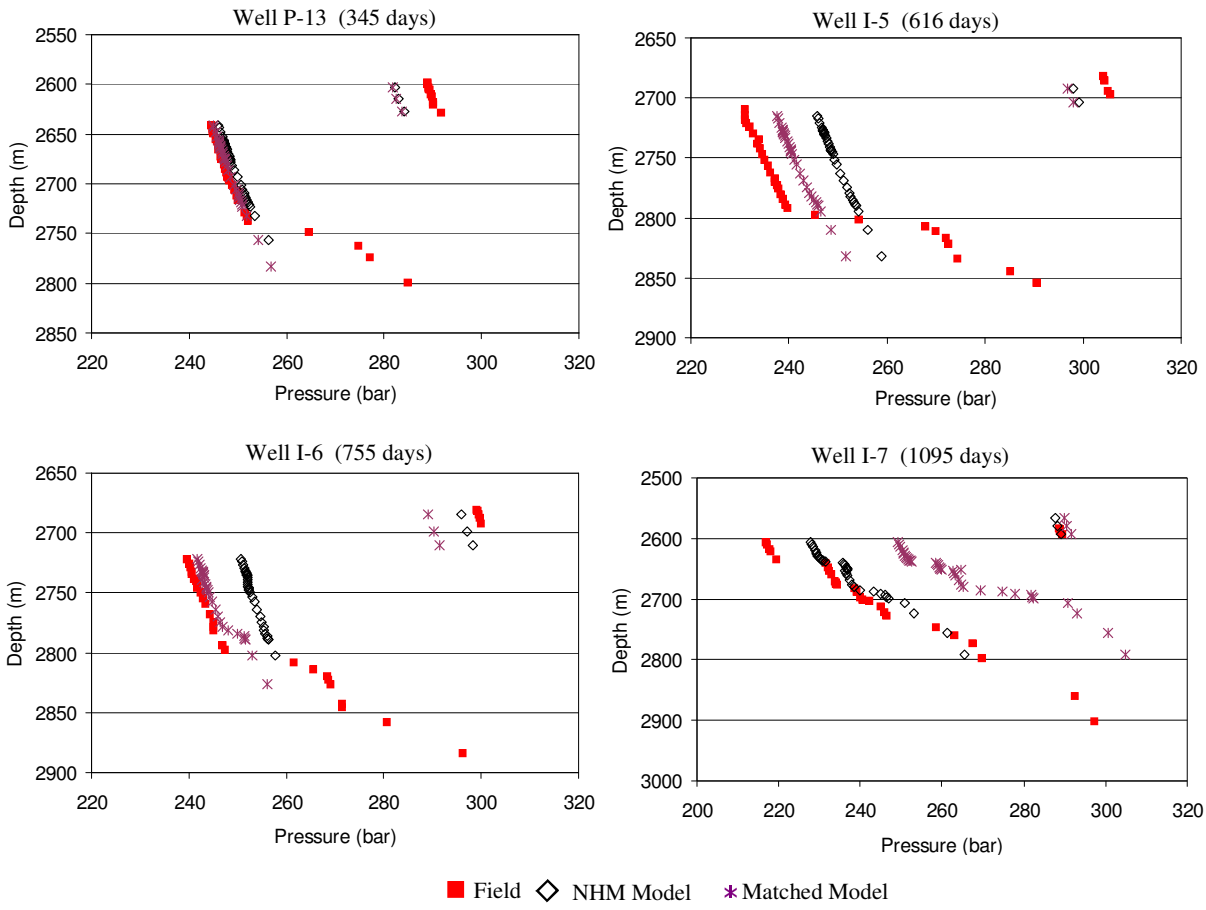

Figure 5.30: Pressure match for four wells for the DE segment runs.

Results – Full Field

By working on the DE segment, we are able to better understand how to achieve a match for this reservoir. We will now take the specifications from Run 3, and apply them to the full field model. Everything remains the same except that the full field model contains fourteen production wells; hence fourteen regions are created, fourteen perturbations parameters are defined and fourteen regional objective functions are optimized. The regions for one model are displayed in Figure 5.31. The vertical allocation of regions is shown in the top figure while two different layers provide the areal geometry of the regions.

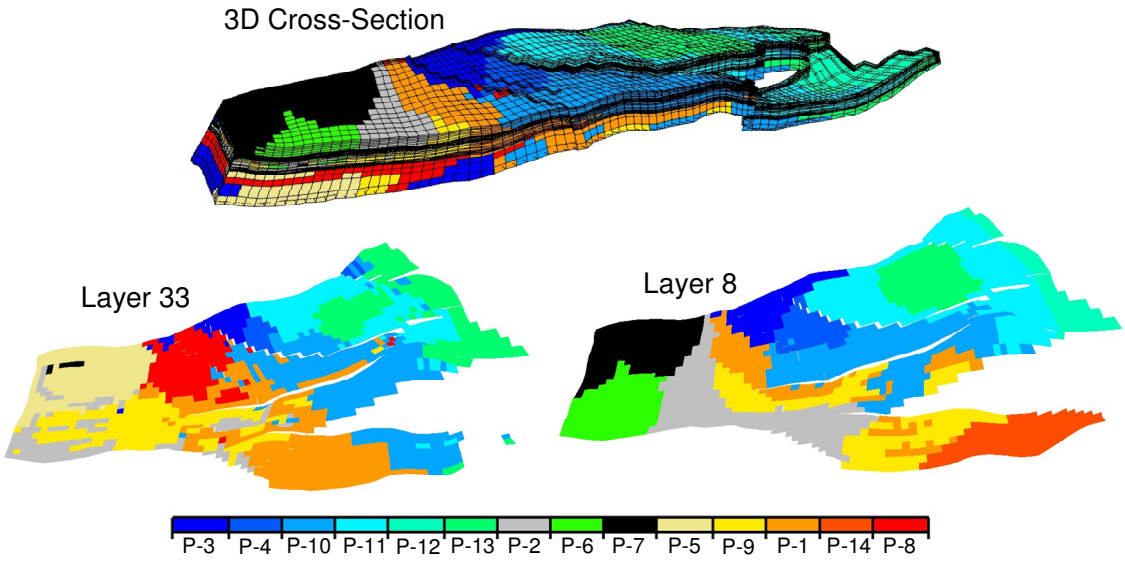

Figure 5.31: Regional definitions of one realization for full field model.

The initial proportions are different for the fourteen regions and are provided in Table 5.5. The initial proportions were manually selected; for the six DE segment wells, the initial proportion originated from the history match obtained on DE segment alone, and for the eight new wells, the initial values were determined by running ten "pre-history matching" flow simulations with varying proportions.

Table 5.5: Initial proportions used in the simulation.

Well Name	Initial Proportion	Well Name	Initial Proportion
P-1	30%	P-8	30%
P-2	30%	P-9	30%
P-3	60%	P-10	20%
P-4	30%	P-11	5%
P-5	20%	P-12	20%
P-6	30%	P-13	20%
P-7	50%	P-14	30%

For the fourteen wells, a quality match is achieved in sixty flow simulations. The goodness of fit for six wells is displayed in Figure 5.32. The remaining wells have fine matches; the oil, gas and water rates as well as the pressures for the field, matched model,

initial model and NHM model are in Appendix D. The light blue lines represent the model without bodies (NHM model). The pink line is the initial model with bodies, but there locations and proportions have not been optimized. For most wells, simply including the bodies improves the match considerably. The red lines show the response from the history matched models.

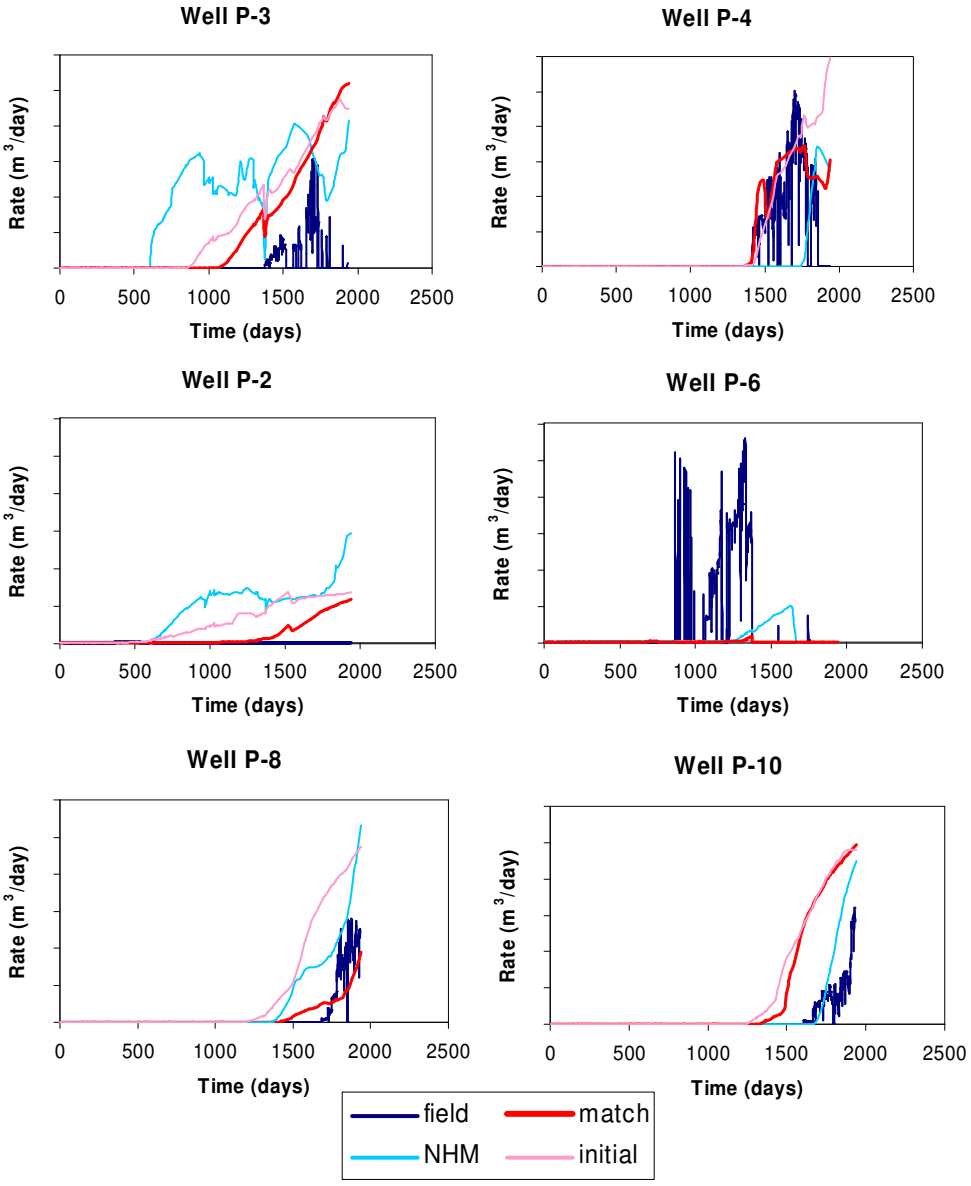

Figure 5.32: Water rates from the field and three models are displayed for six wells in the full field run.

The match for well P-4 is very good; actually, it is matched better for this run than any of the DE segment runs. However, the P-3 well is not matching as well as the DE segments runs. The breakthrough time is a little earlier, and the main problem is that the late time water rate is much too high. Since there are more wells, less weight or significance is

applied to matching this one well. The well P-2 is not producing any water over its life, and although the matched model does not capture this behavior perfectly, it significantly improved over the NHM model and the initial model. The time of breakthrough has been increased almost three years, and the amount of water produced is much less. The worst matched well is P-6, and it does not improve during the history match. Obviously, some "ineffective parameterization" is occurring with respect to the well. This was discussed with the reservoir engineers for the field; however, we currently do not know what is causing water to be produced at this well. The final two wells in the figure show one well where the match got a little worse (P-10) and one well where the match improved to a very nice match (P-8).

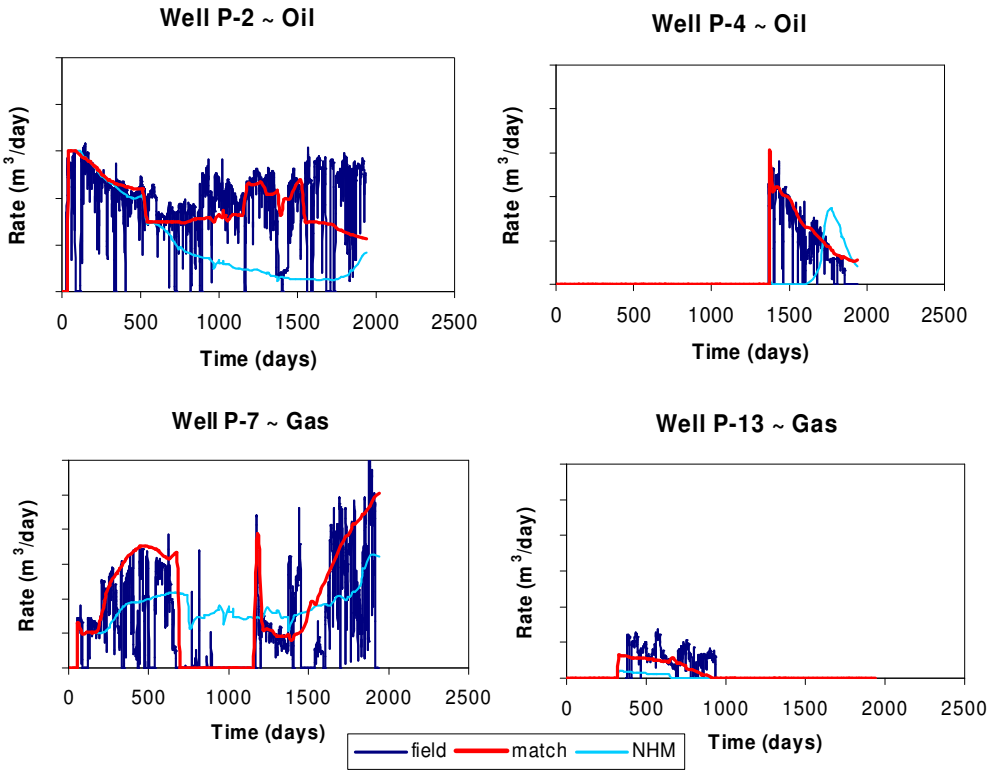

Figure 5.33: Oil rates for two wells and gas rates for two different wells from the full field run.

Although the oil and gas rates are not included in the objective function calculations, the match for these fluids also improves. Figure 5.33 show the oil rates for two wells and the gas rates for two different wells. (All well rates and pressures can be seen in Appendix D.) The oil rate match improved greatly for wells P-2 and P-4; the gas rates also improved but not as much.

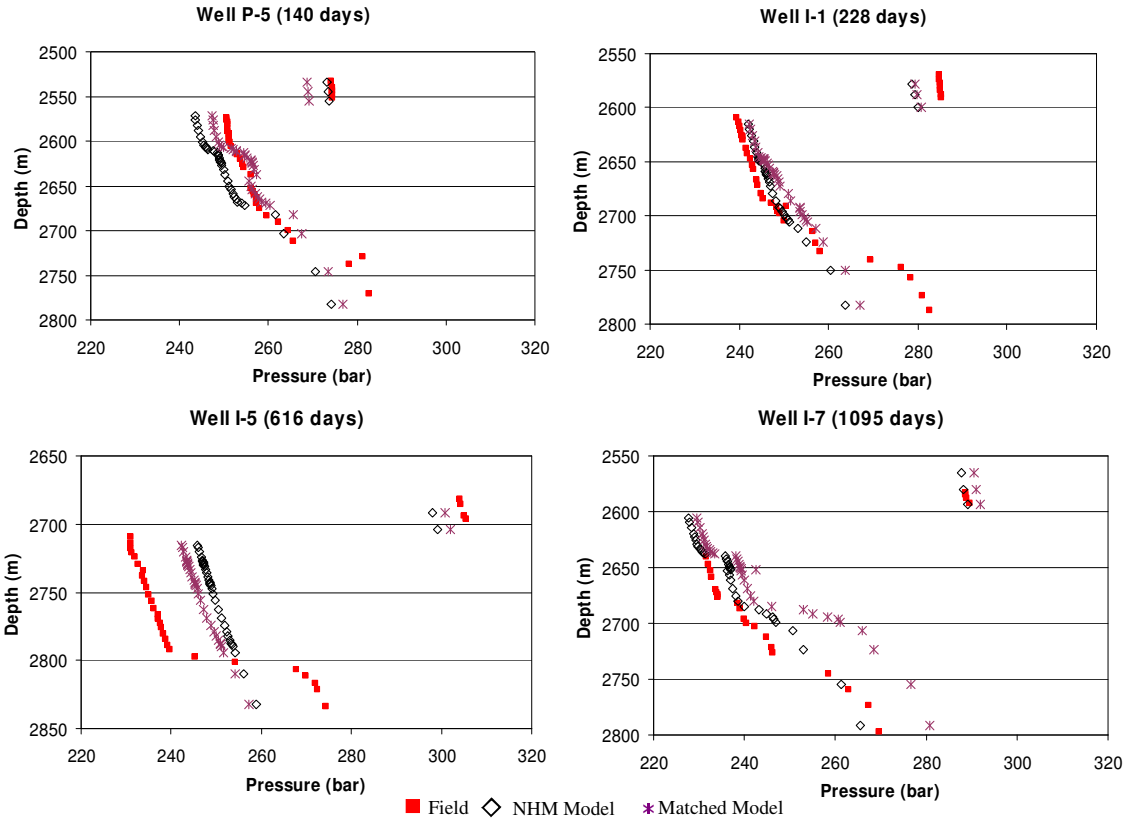

Figure 5.34: The pressure data for four wells from the field, the non-history matched (NHM) model and the matched model for the full field run.

The pressure data is included in the total mismatch value; however, no parameters are perturbed in the model based on the mismatch from pressure data, so there are no regional objective calculations for pressure. The pressures for four wells are presented in Figure 5.34. Wells that were initially matching pressure maintained their match (wells P-5 and I-1). Although the match is not perfect, well I-7 is matched much better for this

run than for the DE segment runs (Figure 5.30). For the few wells that were not matching in the initial model, the match only improved a small amount (well I-5).

The final proportions did not change significantly from their initial values (Table 5.6). Five regions had an increase in proportion, and the rest decreased from their initial values. The proportion of bodies in the region for well P-3 increased to 69%, the highest of any of the regions, and the proportion for the region around well P-11 decreased to 1%, but the rest of the wells were very consistent and stayed between 15% and 36%.

Table 5.6: The final regional proportions in the matched realization.

Well Name	Initial Proportion	Well Name	Initial Proportion
P-1	25%	P-8	24%
P-2	29%	P-9	36%
P-3	69%	P-10	16%
P-4	26%	P-11	1%
P-5	24%	P-12	21%
P-6	26%	P-13	29%
P-7	20%	P-14	15%

Although the proportions did not change much from the initial model, the match improved significantly. In addition to optimizing on the proportion of bodies, the method also allocates the bodies in the best locations. Well P-9 illustrates this and provides a better understanding of how bodies redirect flow in the model to improve the match. The well is completed in the hydrocarbon portions of Horizons 3 and 4 (below the shale layer and above the water table). When no bodies are included (NHM model), the well is producing too much water (Figure 5.35).

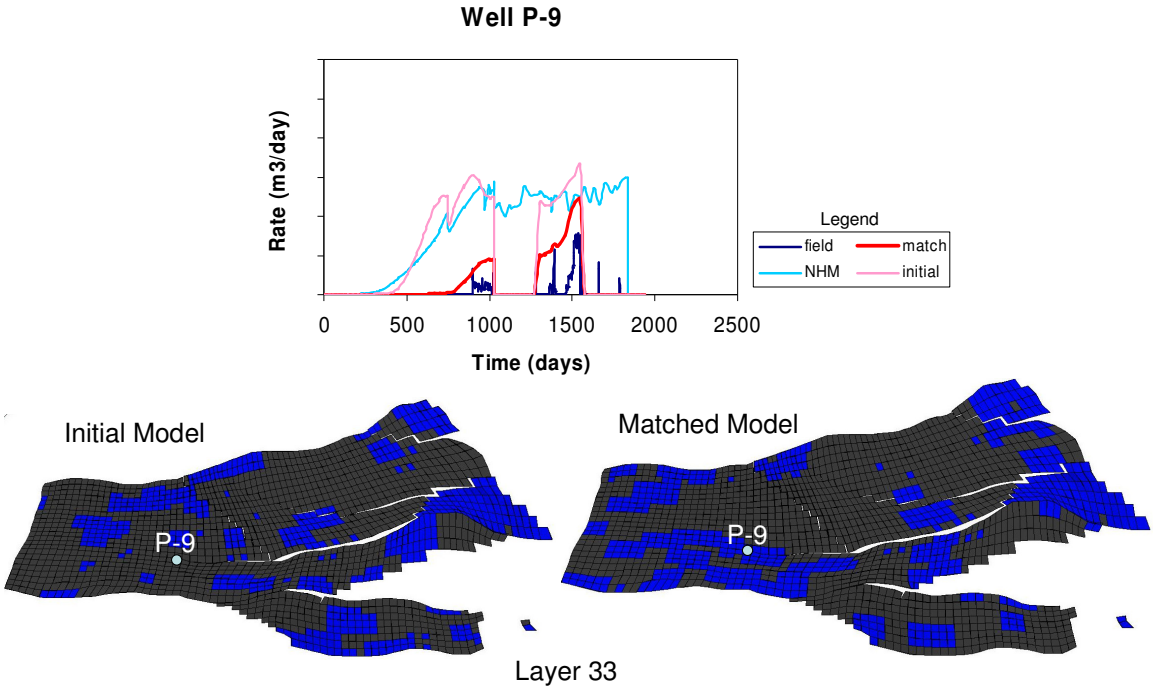

Figure 5.35: Water rates for well P-9 and the location of calcite bodies in layer 33 for the initial model and the matched model.

Even when bodies are included in 30 % of the gridblocks for the initial model, the well continues to produce much more water than the field. For the matched model, the proportion of bodies has increased to 36%; but more importantly, the location of the bodies has changed by perturbing with the PPM, and this caused the match to improve. Figure 5.35 shows body locations for one layer (layer 33) of the initial and matched models. This layer is located below the lowest perforations in the well and above the water table. The location of well P-9 is highlighted, and the approximate region definition for this well at this layer can be seen in Figure 5.31. Since this layer is located between the aquifer and the well completions, having bodies in this layer affects water flow to the well; however, in the initial model, there are few bodies located around the P-9 well. In the matched model, there are significantly more bodies for this layer, and the water rates are more accurately reproduced.

Summary

This chapter shows how the probability perturbation method can be extended to include uncertainty from parameters in addition to random seed uncertainty. Specifically, we looked at perturbing facies proportions along with facies locations. A synthetic example demonstrated that a match can be achieved much faster if both proportions and locations are perturbed. The North Sea case study demonstrated how the method works for a real field case. By including calcite bodies in a stochastic manner, the rates and pressures from the field are better matched in the model.

In this chapter we studied extending the PPM to allow for perturbing of proportions; however, extending the method to other properties such as fracture intensity (Voelker, 2004; Suzuki *et al.*, 2005) and fault displacements (Suzuki and Caers, 2004) has been completed. Furthermore, the flexibility of the PPM would allow additional properties such as aquifer strength, pore volumes and many others to be included in the framework.

"The end is near, the end is near,
Have no fear, the end is near."

- Dr. Suess

Chapter 6

Conclusions and Future Work

6.1 Conclusions

A newly developed history matching technique has been presented. With this technique, reservoir models are constructed that match production data and honor other critical reservoir data such as geological continuity. In present-day reservoir modeling, geostatistical methods are used to calculate many of the reservoir properties for a flow simulation model. Regardless of which geostatistical methods are used, they are stochastic and hence require probabilities or probability distributions. In this work, a match is obtained by perturbing the underlying probabilities and not the properties directly. The probabilities are perturbed by introducing a new probability that is iteratively conditioned to the production data. By perturbing in this manner, the geology, seismic and well data are maintained and a history match is achieved.

A multi-parameter development allows different parts of the reservoir model to be perturbed by different amounts, which dramatically increases the efficiency of the

method. In the multi-parameter method, again only the probabilities in the geostatistical algorithms are perturbed; therefore by construction, these regional perturbations do not result in discontinuities at the region borders. When simulating a reservoir property at a particular grid node, the method searches for the previously simulated grid blocks with a search neighborhood that can cross the region boundaries. When simulating a gridblock in one region, the gridblock properties in other regions are used to determine the probabilities, and this ensures that any history matching related artifacts are avoided.

The multi-parameter development requires a method to define the regions for the entire reservoir model. Streamline-defined regions work well for cases that are dominated by water injection, and unless there is a prevailing geologic reason to use geologic or user defined regions, using streamlines is probably the best choice. If access to a streamline simulator is not available, Voronoi polygons or kriging can be used, but flow based methods (such as streamlines) are preferred for efficiency reasons. Voronoi polygons or kriging could be used for cases with only primary production.

One of the most novel aspects of this history matching method is that large scale parameters can be perturbed. It is often the large scale features that have the greatest influence on fluid flow in a reservoir; thus, production data in a model often is most sensitive to perturbing these parameters. The main concentration of this dissertation is to look specifically at varying facies distributions and proportions, and a number of both synthetic and real examples illustrate how this can be effectively completed; however, the PPM is not limited to facies models or even to large scale parameters.

As demonstrated by a few synthetic examples, other large scale properties such as fractures and even continuous properties such as permeability can be perturbed just as easily with this method. Often the geological scenario (e.g. type of deposition) is an important part of the prior, but due to its intrinsically discrete nature, the geological scenario is not considered a perturbation parameter. However, for the West Coast African reservoir, we accounted for the geological scenario uncertainty by using multiple

training images and various data conditioning situations to generate multiple history matched models.

In the second chapter, the theoretical foundation of the probability perturbation method is succinctly described, and its relation to Bayesian inversion methods is presented. While the method is theoretically sound, its greatest strength may be its practicality. The method is successfully applied to two real field cases. These were both complex reservoirs with many layers, many fault blocks, and more than twenty wells. In the preliminary models, the production data was poorly matching, yet in both cases, once the pre-history matching is completed, the method is able to significantly improve the match in relatively few (< 50) flow simulations.

One reason for the success in practice may have been identifying that history matching is inherently a two-step process. The second step of history matching is actually the easier one: history matching by means of parameters optimization. The time spent here is mostly CPU, and there is much less manual interaction. The PPM is essentially a second-step technique where large scale parameters can be perturbed in an efficient and geologically consistent manner. The first step (labeled in this work as "pre-history matching") is always a manual expert-driven step that must be completed before any second step (automatic or manual) history matching method can be employed (i.e it is not unique to the PPM).

The pre-history matching process ensures that all aspects of the reservoir that may influence flow are included in the model. Also during pre-history matching, one determines the parameters to be perturbed and their maximum range of perturbation. In the Bayesian context discussed in Chapter 2, this first step consists of determining the "prior" model. To obtain a match, both geologic data and engineering data must be included in the prior. The geologic portion of the prior includes reservoir structure, porosity or permeability distributions, variograms or training images, and many others.

The engineering data in the prior could be fluid contacts, relative permeability, aquifer strength or presence, k_v/k_h ratios and many others.

When history matching real reservoirs, there is almost always something initially missing from the prior. It could be the maximum range of perturbation is too small or an aspect of the reservoir simply is not included such as the east-end aquifer in the WCA reservoir example. With many automatic history matching techniques, the prior is fixed before any flow simulations are completed. These techniques work well for synthetic cases where the prior is known, but for real cases where the prior is not known, often a match cannot be achieved.

To determine the prior for real field applications, a large number of flow simulations are required. Typically many more flow simulations are needed to define the prior than are needed for the second step (parameter optimization). This fact often leads people who perform traditional history matching to question the value of automatic history matching; however, automatic history matching techniques are important for the second step because they ensure that the geology and other reservoir data are honored in the final models.

In this dissertation we acknowledge that this first step must be done before any automatic history matching method will work, and we provide some ideas and tools to assist in the first step (e.g. run five to ten flow simulations by perturbing the parameter of interest, and then see if the production from these runs spans the field data). Furthermore, by identifying the two-step nature of history matching, a basis is provided for additional automatic history matching algorithms to be developed for the second step.

The motivation of the dissertation is to create reservoir models with improved prediction power. Matching production and honoring all other information about the reservoir including geologic continuity leads to better predictive models than matching production data alone. A number of examples substantiate this hypothesis. The first real case study (the WCA reservoir) showed that the “future” production response from matched models

can very accurately capture the field response, and by employing a number of matched realizations, an apparently realistic range of uncertainty is provided. The Stanford V example takes this idea further and compares the current method with a "traditional" history matching method that uses different permeability multipliers for different regions. In this case, the difference between the realistic geological continuity of the PPM and the artifact discontinuity of the traditional method is observed. In addition, the PPM obtains improved predictions of future reservoir performance.

6.2 Recommendations for Future Work

Reservoir Modeling

What are the future requirements from history matching? Before attempting to answer that question, let's digress and ask (and answer) a more fundamental question. What are the future requirements from reservoir modeling? As stated in the introduction, the goals of history matching are very tightly coupled with the goals of building reservoir models: a model that accurately predicts future responses from the reservoir. However, we can never build a "perfect" model. A perfect model would be one that accurately reproduces all future responses from the reservoir. From my viewpoint, there are three primary reasons that this cannot be accomplished.

(1) Physics and Numerics – Although fluid flow in porous media can be described very accurately under controlled experimental conditions and for some simplified cases (packed glass beads), the physical processes occurring in the reservoir are much too complicated to capture rigorously. Equations must be simplified and discretized to solve, and the models are further approximated and coarsened for numerical reasons (time constraints). Many simplifications in the physics can be improved upon (e.g. well model and rock compaction), and these are probably the areas where the largest gains in the physical modeling may be achieved.

(2) Human Factor – Another reason models never reproduce field productions is that future production is not only a function of reservoir response but also affected by human decisions and mistakes. The production response is influenced by when and how the wells are drilled, what surface facilities are installed and how proficiently the field is operated. Typically, these types of issues are not accounted for in the reservoir modeling stages. Furthermore, the historical production that is being matched is also affected by the same human factors. Some aspects such as the installed facilities are well known, but other aspects such as formation damage and measurement error in the water and gas production

are not known as well. Therefore, we could be matching production data that does not represent true reservoir response, which would cause predictions to be inaccurate.

(3) Reservoir Property Uncertainty – The last reason for inaccurate predictions is that there is significant uncertainty in the reservoir properties. Uncertainty exists in all measurements taken from the reservoir. This includes non-sales production data (as previously mentioned), but also petrophysical property measurements, seismic measurements/interpretations, and geologic interpretations are all uncertain. Additionally, there are scale issues for many properties in the reservoir. Measurements are taken on a scale that is orders of magnitude smaller than the size of the block they are applied to. There are scale issues for properties we typically think of: porosity and permeability, but there are also issues for other parameters such as relative permeability. Measurements are averaged over a few well cores, but there is usually only one or a few curves specified for the entire reservoir model. Because properties cannot be known exactly, and they are specified for a much larger volume than they were measure on, the reservoir production can not be known precisely either.

The three stated problems can never be fully solved, so the future response will never be perfectly known (i.e. we can never build a "perfect" model). Therefore, instead of building one deterministic model, multiple models can be used to define a range of possible future production. Currently, this is usually done by building a single model, history matching it, and then running the model with different conditions. The different conditions are typically obtained by changing the well constraints (production scenarios), while the properties in the model (e.g. porosity, relative permeability) are usually frozen. Although defining some range for the future predictions is beneficial, we should also continue to improve upon the others areas of concern.

A large amount of research is being done to improve the numerical representation of physics (e.g. gridding and solvers) and that will lead to continuously better models, but the improvements are relatively minor. Better algorithms and faster computers will allow finer and finer gridding, but without some fundamental change in the way the reservoir is modeled (in a way that I cannot envision, but someone else may), we will continue to be limited by the assumptions made in discretizing and simplifying the equations. Research on aspects such as rock compaction and the well model will continue to improve the physical representation of the reservoir in the model.

There is potential in the second area (human factor) because humans do have a substantial impact on what happens in the reservoir and because little attention has been devoted to this aspect. Human behavior is very difficult to characterize (and not something engineers are particularly well suited to analyze), which is probably why so little effort has been spent on this aspect. Nonetheless, it should not be completely ignored because there is an opportunity to better understand the reservoir and improve prediction accuracy. For example, we can explore what component of the production data comes from reservoir response and what is influenced by human interactions. If the two could be separated, one could history match only on the reservoir response part. Then on the prediction runs, one could add back in the human influence. This could dramatically improve ones understanding of the reservoir.

Possibly the simplest way to have more reliable models is to improve the manner in which uncertainty is handled when reservoir modeling. By developing methods to better define uncertainty, more reliable models can be provided. The idea is to generate multiple models that all honor all pieces of information and capture all types of uncertainty: random seed, reservoir property, and geologic scenario (see Section 1.2) and the human factor, part of which accounts for the measurement data uncertainty. To capture fully the reservoir uncertainty, a large set of history matched models may be required. Additionally, the process would be dynamic over the reservoir life; as more

data becomes available, some models will drop out of the set and other new models would become part of the set. This is a daunting task, but the ultimate goal should be clearly stated, so that we know what we are working towards, what research is required to get there, and how ones current research fits into this goal.

History Matching

History matching is an important part of building reservoir models, and possible future improvements for history matching are discussed in this section. There has been a large push recently to automate many aspects of the history matching process; however, it can probably never be fully automated (nor should it be). Before history matching any real reservoir, we must ensure that the prior model contains all aspects of the reservoir that influence flow. This is a manual step that requires an engineer to run flow simulations and to inspect the results. There can be tools to assist this step, but the nature of the models requires it to be a manual process. Once this step is completed, however, I think the rest of history matching can be largely automated.

The ultimate goal of the automated (second) step would be to have a tool to generate multiple history matched realizations that all honor the static and dynamic data and account for all types of uncertainty. Obviously, this type of tool cannot be built immediately or even in a years worth of research, but we really should be aiming for this goal. At first this goal may seem overwhelming; therefore, it has been divided into three achievable stages to be discussed.

At the completion of the first stage, the tool would be able to match a single realization by perturbing multiple properties. Some properties that do not influence each other will be perturbed sequentially (e.g. reservoir structure and petrophysical properties). Other properties are best perturbed jointly. For example, Suzuki (2003) showed that facies and porosity/permeability need to be perturbed at the same time. History matching would proceed in an hierarchical manner starting with the most influential set of parameters, followed by the next most influential set of parameters. This can be continued until all

important parameters are perturbed. If the match is good enough, the algorithm is stopped, and if it is not, the method would loop back to the initial property and repeat the process.

During the pre-history matching step, the user will manually determine which properties to perturb and in what order to perturb them. The method will need to be able to perturb reservoir properties from the largest (reservoir structure) scale (e.g. layering, fault locations and fluid contacts), the middle (regional) scale (e.g. net to gross ratio and facies distributions and proportions), and the finest (gridblock) scale (e.g. permeability and porosity) as well as other more traditional history matching parameters such as relative permeability, k_v/k_h ratio and well productivity index all in a geologically consistent manner.

From the geoscientist side, tools are becoming more readily available to model and perturb all of these kinds of properties (Earth Decision Sciences, 2003). Ongoing research at Stanford's Center for Reservoir Forecasting provides additional tools to perturb many other properties such as faults and fractures (Voelker, 2003; Suzuki, 2004; Rojas, 2005; Suzuki *et al.*, 2005). If the user provides the properties to perturb and their distribution, the technology available today is close to being able to achieve this first step; however, for many of these properties determining the distribution and range of uncertainty is still under research. Therefore this is something that must be studied along with building a tool that can perturb all possible properties in a looped fashion.

The next stage is to incorporate a method into the tool that can determine which properties influence production most. Once a model is available that has all of the required elements of the reservoir, one would simply provide the range of uncertainty for all the properties and the program would determine which properties have the greatest affect on production data (within their range of uncertainty). Then the program would use the looped hierarchical algorithm to perturb those properties and achieve a history matched model. There are tools that have been devised to determine the most sensitive

parameters; for example, experimental design (Damsleth *et al.*, 1992; White *et al.*, 2001), sensitivity coefficients (Schlumberger, 2003) and possibly artificial neural networks as well as others. However, including these techniques in one package that can be automated requires significant work.

The final stage would be to allow multiple geological scenarios to be input and to have the program generate multiple history matched realizations that could span all levels of uncertainty. For a given geologic scenario and a set of properties, multiple history matched models would be generated to span the random seed uncertainty. Then the set of properties would be changed and again multiple models would be matched to account for the uncertainty from the properties. Finally, new geologic scenarios would be introduced, and again, various properties would be perturbed, and multiple realizations would be generated. All of these models could be ran under various fluid rate constraints (as is currently done for a single model) to account for flow data uncertainty. This should then capture the essential uncertainty in the reservoir.

Directly including the production data information into the reservoir model (i.e. non-iterative technique) may be seen as the ultimate goal of including dynamic data in reservoir models. However, production data has a non-linear relationship with the model response through the model properties, and only if the relationship could be sufficiently linearized would a direct method be effective. Linearizing this relationship appears impossible, and, most likely, iterative methods will always be required.

Given this description of how I see the future of history matching developing, I want to conclude by stating how the current work fits into the big picture. This dissertation provides a small but important piece by primarily studying how to efficiently perturb large scale properties and in particular, facies.

"Extra, et cetera, addendum, codicil;
whatever you call it, here it is."

- on anonymous gravestone

Appendices

A. Files and Fortran Code required for PPM

The files displayed here are for a 3D multi-facies reservoir model that uses streamline defined regions. The first sections are for the Fortran code and the last two sections are parameter files for the MPS geostatistics simulation (snesim) and the flow simulation (eclipse). Other files (not shown) that are needed for PPM include a file with the reference production data to be read by the objective function calculation routine and executables for the geostatistics and flow simulations and any files that they require.

Main program (deform.f90)

```
!*************************************************************
!
!     Program: deform (main program)
!
!     Purpose: This program is used to history match reservoir
!           models and honor other data by perturbing the
!           probabilities used in the geostatistical routines.
!           The program uses a two-loop development.  The inner
!           loop is completed by brentpara.f and the outer loop
!           is completed in this routine.  There are 9 separate
!           routines.
!
!*************************************************************
```

```
program main

real t(4), tmin(4), test, minmin, funz(4)
integer iter, seed

nreg=4            !number of regions
imax=20           !maximum numer of iterations
opttol= 1         !global convergence tolerance
ii=0              !iteration count

open(35,file='outer.out')     !output file
open(33,file='deform.dbg')    !output file

open(50,file='rand.dat')   !File with large, odd random numbers

! Initial guess

call system('snesim')!exe program to perform MPS to modle facies

call perm()       !converts facies info. to perm. data

call objfun(funz) !obj. fun. routine - runs flow simulation
                  !and calculates the mismatch

minmin = sum(funz)
write(35,*) ii, minmin
write(33,*) ii, funz, minmin
test=minmin

call system('copy test.rsm test_init.rsm')
call system('copy cosnesim.out snesimsoft.out')

do while (ii.le.imax.and.minmin.ge.opttol)      !starts outer loop
 ii=ii+1

 read(50,*) seed
 call chnpar(seed,'cosnesim.par')   !changes random seed in MPS run

 !performs inner loop,return best match, test, and best rD values, tmin
 call brentpara(0.0,0.75,1.0,0.05,tmin,nreg,test)

 write(35,*) ii, test
 write(33,*) ii, tmin, test

 if (test < minmin) then

  call system('copy cosnesim_opt.out snesimsoft.out') !saves best runs
  call system('copy well.dat  wellopt.out')
```

```
  write(33,*) 'kept this outer iteration', ii
  minmin=test

 else
  test=minmin

 end if

end do

close(50)
END
```

Inner loop (Brentpara.f)

```
!*****************************************************************
!
!  PROGRAM: chnpar
!
!  PURPOSE:  Solves npara 1D optimization problems in parallel
!
!*****************************************************************

      SUBROUTINE brentpara(ax,bx,cx,tol,xmin,npara,brent)


      INTEGER ITMAX
      REAL brent,ax,bx,cx,tol,CGOLD,ZEPS
      PARAMETER (ITMAX=7,CGOLD=.3819660,ZEPS=1.0e-10)
      INTEGER iter,npara

      real a(npara),b(npara),e(npara),fu(npara),fv(npara),fw(npara)
      real   fx(npara),u(npara),w(npara),x(npara),xm(npara),v(npara)
      real  tol1(npara),tol2(npara),xmin(npara)

      real d,etemp,p,q,r


!
! Initialize all npara parameters
!

      do np = 1, npara
       a(np)=min(ax,cx)
       b(np)=max(ax,cx)
       v(np)=bx
       w(np)=v(np)
       x(np)=v(np)
       e(np)=0.
```

```
      end do
      brent_old = brent

!
! Here we call the objective function with all npara parameters
! stord in x
!

      iiter=1

      call head_objfun(x,fx) !x=rD

      brent = sum(fx)

      write(33,*) iiter, fx, x
      iiter=iiter+1

      call regs()

      call system('copy cosnesim.out cosnesim_opt.out')
      call system('copy test.rsm well.dat')

      do np = 1, npara
       fv(np)=fx(np)
       fw(np)=fx(np)
      end do

!
! First search for a new set of parameters
!
      iswitch = 0
      iter = 1
      do while (iter.lt.ITMAX.and.iswitch.eq.0)


      imdone = 0
      do np = 1, npara
       xm(np)=0.5*(a(np)+b(np))
         tol1(np)=tol*abs(x(np))+ZEPS
         tol2(np)=2.*tol1(np)

         if(abs(x(np)-xm(np)).le.(tol2(np)-.5*(b(np)-a(np)))) then
            imdone = imdone+1
       end if
        end do
      if (imdone == npara) goto 3

      do np = 1, npara

        if(abs(e(np)).gt.tol1(np)) then
```

```
          r=(x(np)-w(np))*(fx(np)-fv(np))
          q=(x(np)-v(np))*(fx(np)-fw(np))
          p=(x(np)-v(np))*q-(x(np)-w(np))*r
          q=2.*(q-r)
          if(q.gt.0.) p=-p
          q=abs(q)
          etemp=e(np)
          e(np)=d
          if(abs(p).ge.abs(.5*q*etemp).or.p.le.q*(a(np)-x(np))
     +            .or.p.ge.q*(b(np)-x(np)))   goto 1
          d=p/q
          u(np)=x(np)+d
          if(u(np)-a(np).lt.tol2(np)
     +      .or. b(np)-u(np).lt.tol2(np))  d=sign(tol1(np),xm(np)-
x(np))
          goto 2
        endif
1       if(x(np).ge.xm(np)) then
          e(np)=a(np)-x(np)
        else
          e(np)=b(np)-x(np)
        endif
        d=CGOLD*e(np)
2       if(abs(d).ge.tol1(np)) then
          u(np)=x(np)+d
        else
          u(np)=x(np)+sign(tol1(np),d)
        endif

      end do

!
! Evaluate the function in these parameters
!

      call head_objfun(u,fu)

      write(33,*)  iiter,fu, u
      iiter=iiter+1


!
! Now decide, for each parameter, on the next 3 values used for
! parabolic interpolation
!

      do np = 1, npara

        if(fu(np).le.fx(np)) then
          if(u(np).ge.x(np)) then
```

```
            a(np)=x(np)
          else
            b(np)=x(np)
          endif
          v(np)=w(np)
          fv(np)=fw(np)
          w(np)=x(np)
          fw(np)=fx(np)
            x(np)=u(np)
          fx(np)=fu(np)

        else
          if(u(np).lt.x(np)) then
            a(np)=u(np)
          else
            b(np)=u(np)
          endif
          if(fu(np).le.fw(np) .or. w(np).eq.x(np)) then
            v(np)=w(np)
            fv(np)=fw(np)
            w(np)=u(np)
            fw(np)=fu(np)
        else if(fu(np).le.fv(np).or.v(np).eq.x(np).or.v(np).eq.w(np))then
            v(np)=u(np)
            fv(np)=fu(np)
          endif
         endif

        end do

                  brent1 = sum(fu)
            if (brent1 < brent) then
          call system('copy cosnesim.out cosnesim_opt.out')
            call system('copy test.rsm well.dat')
            write(33,*) 'kept this inner iteration', iiter-1
            brent=brent1
            end if

!
!           check to see if we found one better than the previous best
!           in first three runs of inner loop,
!           if yes - continue, if no - end inner loop
!

                  if(iter.eq.2) then
                        if(brent.gt.brent_old) then
                              iswitch=1
                        end if
                  end if
```

```
      iter = iter + 1
       end do


3     do np = 1, npara
      xmin(np) = x(np)
      end do
      return
      END
```

Objection Function Calculation (objfun.f90)

```
!**********************************************************************
!
!  PROGRAM: head_objfun
!
!  PURPOSE: Calls programs that get ready to run flow simulation
!
!**********************************************************************

subroutine head_objfun(rd,obj)

real obj(4), rd(4)

call createsoft(rd)
call system('cosnesim')
call perm()
call objfun(obj)

END

!**********************************************************************
!
!  PROGRAM: objfun
!
!  PURPOSE: Runs flow simulation and calculate mismatch for each region
!
!**********************************************************************

subroutine objfun(obj)

real obj(4)
nreg = 4 !numbre of regions

call system('$eclipse -ver 2001a_2 -file test')

call exwellinfo() ! extracts prod. data from simulation output files

!calculation of the objective function for each region
```

```
open(20,file='ref1.out')
open(40,file='well1.out')
open(21,file='ref2.out')
open(41,file='well2.out')
open(22,file='ref3.out')
open(42,file='well3.out')
open(23,file='ref4.out')
open(43,file='well4.out')

do i=1,nreg
obj(i) = 0
end do

do i=1,40
read(40,*) pres
read(20,*) truepres
obj(1)=obj(1)+(truepres-pres)**2
end do

do i=1,40
read(41,*) pres2
read(21,*) truepres2
obj(2)=obj(2)+(truepres2-pres2)**2
end do

do i=1,40
read(42,*) pres3
read(22,*) truepres3
obj(3)=obj(3)+(truepres3-pres3)**2
end do

do i=1,40
read(43,*) pres4
read(23,*) truepres4
obj(4)=obj(4)+(truepres4-pres4)**2
end do

obj(1)=sqrt(obj(1))
obj(2)=sqrt(obj(2))
obj(3)=sqrt(obj(3))
obj(4)=sqrt(obj(4))

close(20)
close(21)
close(22)
close(23)
close(40)
close(41)
close(42)
```

```
close(43)

end
```

Creates P(A|D) soft data (creatsoft.f)

```
!**********************************************************************
!
!  PROGRAM: createsoft
!
!  PURPOSE:Creates the P(A|D) soft data from rD and previous best model
!
!**********************************************************************

      subroutine createsoft(t)

      real t(4), val(3)
      nreg=4

        open(15,file='snesimsoft.out')    !previous best model
        read(15,*)
        read(15,*)
        read(15,*)
        read(15,*)

        open(16,file='regions.dat')       !file with region geometries
        read(16,*)
        read(16,*)
        read(16,*)

        open(45,file='softdata.dat')            !output file
        write(45,*) 'soft data'
        write(45,*) '3'
        write(45,*) 'zero'
        write(45,*) 'one'
        write(45,*) 'two'

            pa1=0.45          !prior probabilites [P(A)]
            pa2=0.35
            pa3=0.20

! Calculation of P(A|D) for three facies in 3D

        nx = 100
        ny = 100
        nz = 10
        do j = 1, nz
        do i = 1, ny
        do k = 1, nx
                read(15,*) dum
```

```
                read(16,*) reg
                  do ijk = 1,nreg
                  if (reg == ijk) then
                        if (dum == 0) then
                              val(1) = pa1*t(ijk)+1-t(ijk)
                              val(2) = pa2*t(ijk)
                              val(3) = pa3*t(ijk)

                        else if (dum == 1) then
                              val(1) = pa1*t(ijk)
                              val(2) = pa2*t(ijk)+1-t(ijk)
                              val(3) = pa3*t(ijk)

                        else
                        val(1) = pa1*t(ijk)
                        val(2) = pa2*t(ijk)
                        val(3) = pa3*t(ijk)+1-t(ijk)
                  end if
                  end if
            end do
                write(45,*) val

        end do
        end do
        end do

      close(45)
      close(15)
      close(16)
        END
```

Extract simulated production data (exwellinfo.f90)

```
!*********************************************************************
!
!  PROGRAM: exwellinfo.f90
!
!  PURPOSE: Extracts prod. infomation from simulation output to be used
!                 by objfun.f90 in calculating the mismatch.
!
!*********************************************************************

subroutine exwellinfo()

open(4,file='test.rsm')
open(5,file='well1.out')
open(6,file='well2.out')
open(7,file='well3.out')
open(8,file='well4.out')
```

```
do i=1,10
 read(4,*)
end do

do i=1,40
 read(4,'(27x,f10.8,3x,f10.8,3x,f10.8,3x,f10.8)') temp1, temp2, temp3,
temp4
 write(5,*) temp1
 write(6,*) temp2
 write(7,*) temp3
 write(8,*) temp4
end do

close(5)
close(6)
close(7)
close(8)
close(4)

END
```

Convert facies to permeability (perm.f90)

```
!**********************************************************************
!
!  PROGRAM: perm
!
!  PURPOSE: Calculating reservoir properties from facies is important.
!     Here it is done by assigning different perm. to each facies, but
!     more sophisticated techniques can be (and are) used.
!
!**********************************************************************

      subroutine perm()

      open(10,file='cosnesim.out')
      read(10,*)
      read(10,*)
      read(10,*)
      read(10,*)

      open(20,file='perm.out')
      write(20,'(a5)') 'PERMX'

      nx = 100
      ny = 100
      nz = 10
```

```
      do k = 1, nz
      do j = 1, ny
      do i = 1, nx
            read(10,*) idum
            if(idum .eq. 0) then
                  write(20,*) '    10'
            else if(idum .eq. 1) then
                  write(20,*) '    1000'
            else
                  write(20,*) '    0.01'
            end if
      end do
      end do
      end do
                  write(20,*) '/'

      close(10)
      close(20)

      END
```

New random seed (chnpar.f90)

```
!**********************************************************************
!
!  PROGRAM: chnpar
!
!  PURPOSE:  Modify the input gslib parameter with a new random seed
!
!**********************************************************************

      subroutine chnpar(seed,parfile)

      integer seed
      character a*70, parfile*12

! Modify parfile with new seed

    open(5,file=parfile)
    open(6,file='temp')
      do i = 1, 34
        read(5,'(a70)') a
        lb=len(trim(a))
        if(lb.le.17) then
          write(6,'(a70)') a
        else if (a(lb-17:lb).eq.'random number seed') then
          write(6,*)seed,'                          -random number seed'
        else
          write(6,'(a70)') a
```

```
          end if
        end do

   close(5)
   close(6)

    call system('copy temp '//parfile)
      call system('del temp ')

      return
    end
```

Region Geometries (regs.f90)

```
!**********************************************************************
!
!  PROGRAM: regs
!
!  PURPOSE: Runs streamline simulation and calc. regions from output
!
!**********************************************************************

      subroutine regs()

      integer idum(-20:70,-20:85,-10:25),idu(-20:70,-20:85,-10:25)

      nx = 100
      ny = 100
      nz = 10

!runs the streamline simulation

      call system('3dsl.bat')

      open(10,file='3dsl.wir') !sim. output that is read
      do i=1,5
       read(10,*)
      end do

      open(20,file='regions.dat')   !file where regions are written
      write(20,'(a5)')'regions'
      write(20,*) '1'
      write(20,*) 'regs'

!reading the output file

      do k = 1, nz
      do j = 1, ny
      do i = 1, nx
```

```
         xx=i
         yy=ny-j+1
         zz=nz-k+1
             read(10,'(7x,i2)') idum(xx,yy,zz)
       end do
       end do
       end do

!if no streamlines go through a particular block, this interpolates
!what the streamline defined regions should be at that location.

do it=1,10

do k = 1, nz
do j = 1, ny
do i = 1, nx
   todd= idum(i,j,k)
      if(todd>0.1) then

        idu(i,j,k)=idum(i,j,k)
        if (idu(i+it,j,k) < 0.5)  idu(i+it,j,k)=idum(i,j,k)
        if (idu(i-it,j,k) < 0.5)  idu(i-it,j,k)=idum(i,j,k)
        if (idu(i,j+it,k) < 0.5)  idu(i,j+it,k)=idum(i,j,k)
        if (idu(i,j-it,k) < 0.5)  idu(i,j-it,k)=idum(i,j,k)
        if (idu(i+it,j+it,k) < 0.5)  idu(i+it,j+it,k)=idum(i,j,k)
        if (idu(i-it,j+it,k) < 0.5)  idu(i-it,j+it,k)=idum(i,j,k)
        if (idu(i+it,j-it,k) < 0.5)  idu(i+it,j-it,k)=idum(i,j,k)
        if (idu(i-it,j-it,k) < 0.5)  idu(i-it,j-it,k)=idum(i,j,k)
        if (idu(i,j,k+it) < 0.5)  idu(i,j,k+it)=idum(i,j,k)
        if (idu(i,j,k-it) < 0.5)  idu(i,j,k-it)=idum(i,j,k)
        if (idu(i+it,j,k+it) < 0.5)  idu(i+it,j,k+it)=idum(i,j,k)
        if (idu(i-it,j,k+it) < 0.5)  idu(i-it,j,k+it)=idum(i,j,k)
        if (idu(i,j+it,k+it) < 0.5)  idu(i,j+it,k+it)=idum(i,j,k)
        if (idu(i,j-it,k+it) < 0.5)  idu(i,j-it,k+it)=idum(i,j,k)
        if (idu(i+it,j+it,k+it) < 0.5)  idu(i+it,j+it,k+it)=idum(i,j,k)
        if (idu(i-it,j+it,k+it) < 0.5)  idu(i-it,j+it,k+it)=idum(i,j,k)
        if (idu(i+it,j-it,k+it) < 0.5)  idu(i+it,j-it,k+it)=idum(i,j,k)
        if (idu(i-it,j-it,k+it) < 0.5)  idu(i-it,j-it,k+it)=idum(i,j,k)
        if (idu(i+it,j,k-it) < 0.5)  idu(i+it,j,k-it)=idum(i,j,k)
        if (idu(i-it,j,k-it) < 0.5)  idu(i-it,j,k-it)=idum(i,j,k)
        if (idu(i,j+it,k-it) < 0.5)  idu(i,j+it,k-it)=idum(i,j,k)
        if (idu(i,j-it,k-it) < 0.5)  idu(i,j-it,k-it)=idum(i,j,k)
        if (idu(i+it,j+it,k-it) < 0.5)  idu(i+it,j+it,k-it)=idum(i,j,k)
        if (idu(i-it,j+it,k-it) < 0.5)  idu(i-it,j+it,k-it)=idum(i,j,k)
        if (idu(i+it,j-it,k-it) < 0.5)  idu(i+it,j-it,k-it)=idum(i,j,k)
        if (idu(i-it,j-it,k-it) < 0.5)  idu(i-it,j-it,k-it)=idum(i,j,k)

      end if

end do
```

```
end do
end do
idum=idu
end do

!writes the output

      do k = 1, nz
      do j = 1, ny
      do i = 1, nx

                  write(20,*) idu(i,j,k)
      end do
      end do
      end do

      close(10)
      close(20)

      END
```

SNESIM parameter file (cosnesim.par)

```
                  Parameters for COSNESIM
                  ***********************

START OF PARAMETERS:
data.dat                        - file with original data
1  2  3  4                      - columns for x, y, z, variable
3                               - number of categories
0   1   2                       - category codes
0.45  0.35  0.2                 - (target) global pdf
0                               - use (target) vertical proportions (0=n
vertprop.dat                    - file with target vertical proportions
0.5                             - servosystem parameter (0=no correction
0                               - debugging level: 0,1,2,3
snesim.dbg                      - debugging file
cosnesim.out                      - file for simulation output
1                               - number of realizations to generate
100    0.5   1.0                - nx,xmn,xsiz
100    0.5   1.0                - ny,ymn,ysiz
10     0.5   1.0                - nz,zmn,zsiz
       21375                         -random number seed
20                              - max number of conditioning primary dat
5                               - min. replicates number
1     0.9                       - condtion to LP (0=no, 1=yes), weight f
softdata.dat                    - file for local proportions
0                               - condition to rotation and affinity (0=
rotangle.dat                    - file for rotation and affinity
```

```
1                               - number of affinity categories
1.0  1.0  1.0                   - affinity factors (X,Y,Z)
5                               - number of multiple grids
new.out                         - file for training image
200  200  10                    - training image dimensions: nxtr, nytr,
1                               - column for training variable
50.0   30.0   8.0               - maximum search radii (hmax,hmin,vert)
45.0   0.0   0.0                - angles for search ellipsoid
```

Flow simulation deck (test.data)

```
--**********************NEW RUNSPEC SECTION**********************
RUNSPEC
TITLE
PPM Example Simulation

DIMENS
  100 100  10  /

NONNC

OIL
WATER

FIELD

EQLDIMS
    1  100   10    1    1 /

TABDIMS
    1    1   20   20    1    1 /

WELLDIMS
    5   10    1    5 /

NUPCOL
    3 /

START
   1 'JAN' 2000  /

NSTACK
60 /

UNIFOUT
UNIFIN

--================================================================
GRID
```

```
--INIT
NOGGF

DX
100000*10.0 /

DY
100000*10.0 /

EQUALS

    'TOPS' 2000.0 1 100 1 100 1 1/
    'PORO'   0.25 /
    'DZ'     10.0/
/


-----------------------------------------
--  permeability field
--   for PERMX
-----------------------------------------
INCLUDE
'perm.out'
/
COPY
     'PERMX' 'PERMY' /
     'PERMX' 'PERMZ' /
/
MULTIPLY
      PERMZ 0.1 /
/

--=================================================================

PROPS

--***   INCLUDE YOUR RELATIVE PERMEABILITY TABLES HERE    ***
SWFN
0.15 0.000 0.0
0.20 0.029 0.0
0.30 0.096 0.0
0.40 0.175 0.0
0.52 0.280 0.0
0.62 0.390 0.0
0.70 0.500 0.0
1.00 0.500 0.0/

SOF2
0.00 0.000
0.30 0.000
```

```
0.40 0.086
0.48 0.161
0.60 0.326
0.66 0.436
0.74 0.644
0.85 1.000/

PVTW
-- ref pr  bwi     cw      uw
      14.7  1.035 3.E-6    0.30  0  /

PVCDO
1600.0  0.9999 3.E-5 1.0 0.0 /

ROCK
930.0  8.E-6  /

DENSITY
54.2  62.97  26.74  /

SOLUTION
================================================================

EQUIL
-- Depth RfP  WOC    4  GOC    6      7      8       9
   2000  1200 2260.  3*  0 0  0/

RPTSOL
-- P   So  Sw  Sg  Rs  Rv  restart Inplace eq-data
   1   0   1   1   0   0     5       0       0   /

SUMMARY
================================================================

RUNSUM
SEPARATE
RPTONLY

-- Variables to output are greatly flexible.
-- Summary variables influence little on simulation.
-- Summary output influences only disk capacity.

WWCT
'P1'
'P2'
'P3'
'P4'
/

WBHP
```

```
'P1'
'P2'
'P3'
'P4'
'I1'
/

WWIR
'I1'
/

FOPR


--EXCEL
SCHEDULE
================================================================

WELSPECS
'P1' 'G1'   5   7     -1 'OIL'    /
'P2' 'G1'   95  2     -1 'OIL'    /
'P3' 'G1'   3   96    -1 'OIL'    /
'P4' 'G1'   98  94    -1 'OIL'    /
'I1' 'G1'   50  51    -1 'WATER'  /
/

COMPDAT
-- 1  2   3    4    5    6   7  8     9
'P1'  5   7     1    10 'OPEN' 1* 1*  0.375 /
'P2'  95  2     1    10 'OPEN' 1* 1*  0.375 /
'P3'  3   96    1    10 'OPEN' 1* 1*  0.375 /
'P4'  98  94    1    10 'OPEN' 1* 1*  0.375 /
'I1'  50  51    1    10 'OPEN' 1* 1*  0.375 /
/

-- balanced production

WCONPROD
-- 1      2      3  4     5  6  7   8   9
'P1' 'OPEN' 'LRAT' 3* 100.0 1*  100 /
'P2' 'OPEN' 'LRAT' 3* 100.0 1*  100 /
'P3' 'OPEN' 'LRAT' 3* 100.0 1*  100 /
'P4' 'OPEN' 'LRAT' 3* 100.0 1*  100 /
/
WCONINJ
'I1' 'WATER' 'OPEN' 'RATE' 400 3* 1350.0    /
--'I1' 'WATER' 'OPEN' 'BHP' 3000 3* 1250.0    /
/
--RPTRST
--'BASIC=4' /
```

```
TUNING
0.1 30./
/
30 1 40 /

DATES
1 APR 2000 /
1 JUL 2000 /
1 OCT 2000 /
1 JAN 2001 /
1 APR 2001 /
1 JUL 2001 /
1 OCT 2001 /
1 JAN 2002 /
1 APR 2002 /
1 JUL 2002 /
1 OCT 2002 /
1 JAN 2003 /
1 APR 2003 /
1 JUL 2003 /
1 OCT 2003 /
1 JAN 2004 /
1 APR 2004 /
1 JUL 2004 /
1 OCT 2004 /
1 JAN 2005 /
1 APR 2005 /
1 JUL 2005 /
1 OCT 2005 /
1 JAN 2006 /
1 APR 2006 /
1 JUL 2006 /
1 OCT 2006 /
1 JAN 2007 /
1 APR 2007 /
1 JUL 2007 /
1 OCT 2007 /
1 JAN 2008 /
1 APR 2008 /
1 JUL 2008 /
1 OCT 2008 /
1 JAN 2009 /
1 APR 2009 /
1 JUL 2009 /
1 OCT 2009 /
1 JAN 2010 /
/
END
================================================================
```

B. Oil, Gas and Water Plots for WCA Reservoir

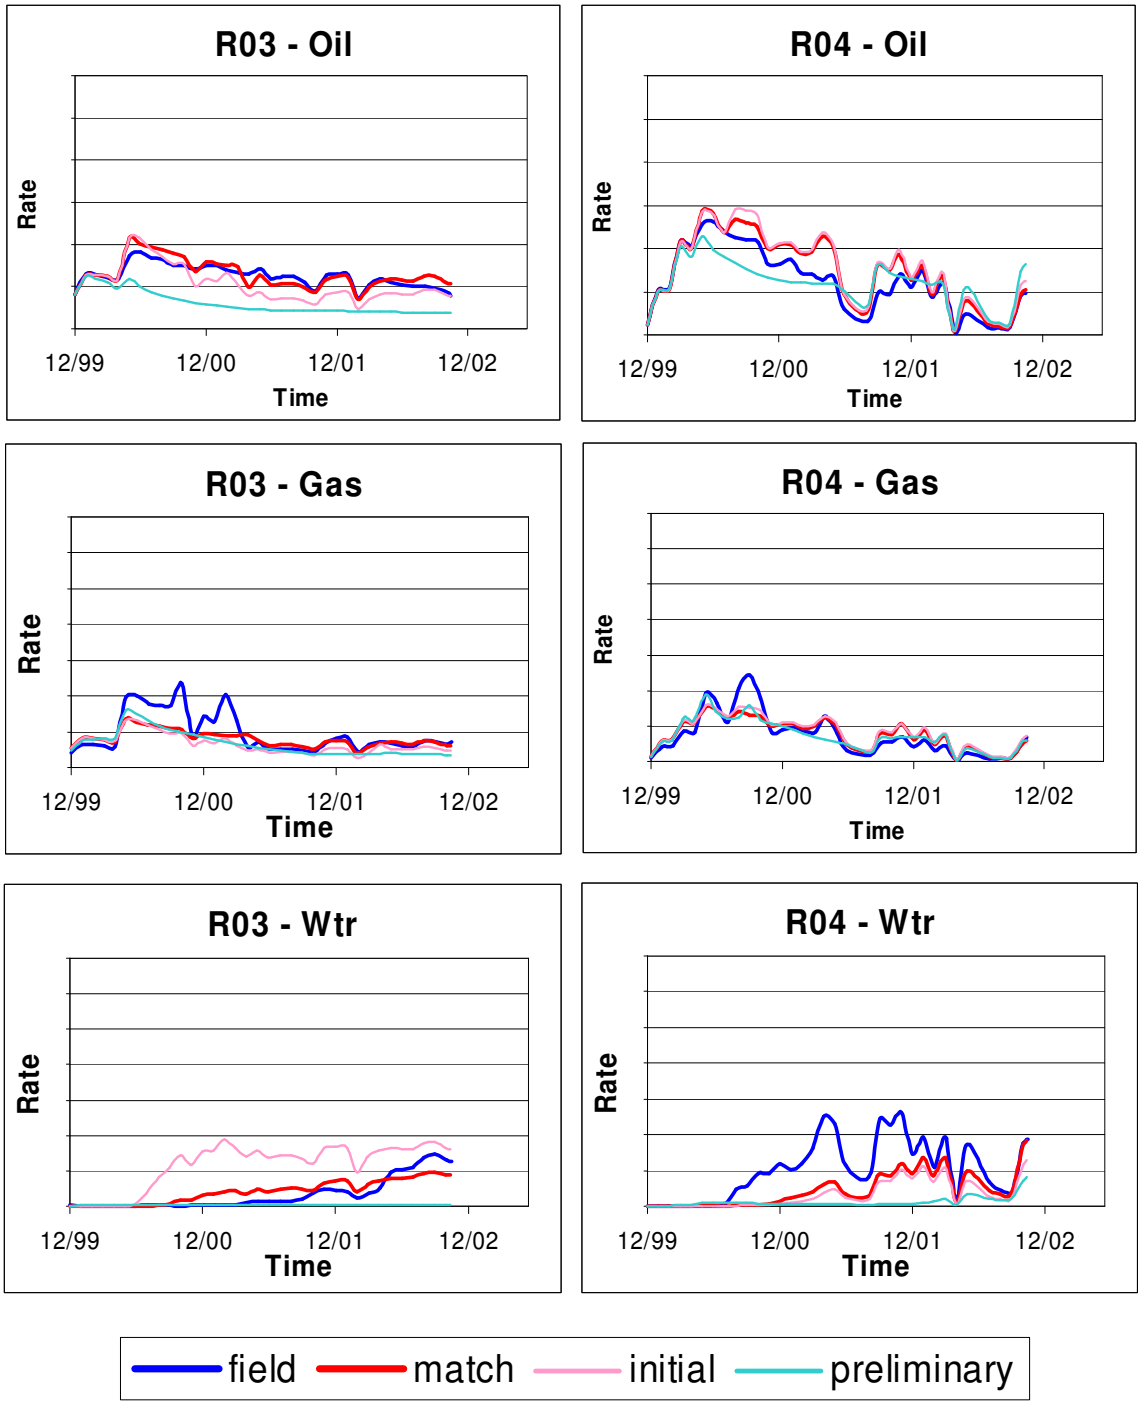

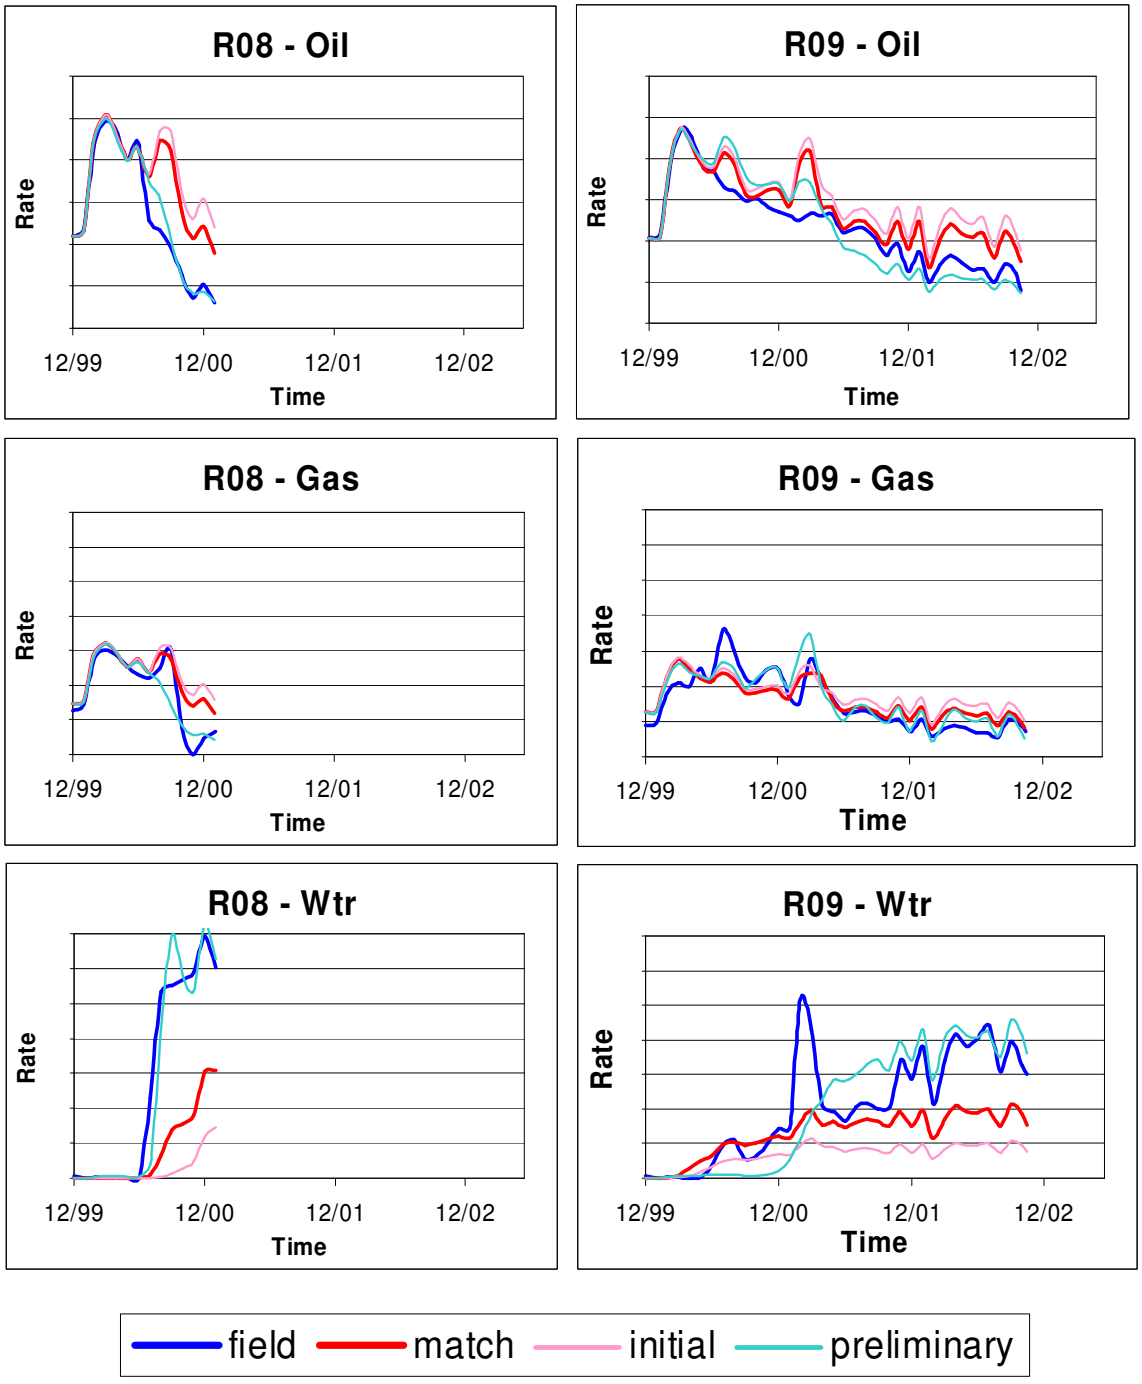
R08 - Oil
R09 - Oil
R08 - Gas
R09 - Gas
R08 - Wtr
R09 - Wtr
Rate
Time
12/99
12/00
12/01
12/02
field
match
initial
preliminary

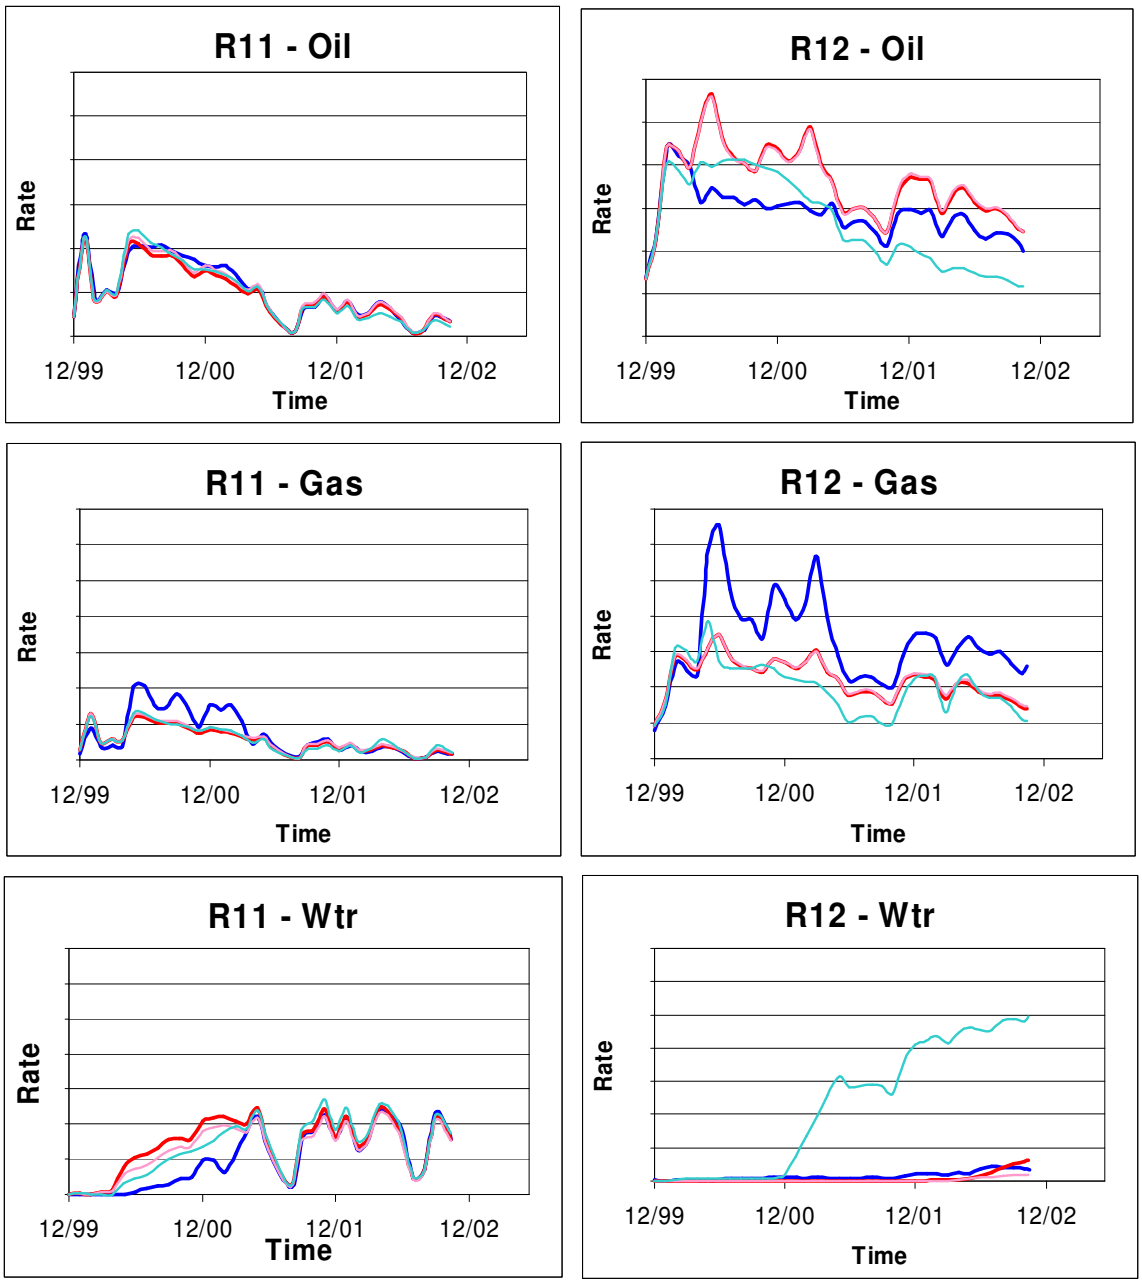
R11 - Oil
Rate
12/99
12/00
12/01
12/02
Time
R12 - Oil
Rate
12/99
12/00
12/01
12/02
Time
R11 - Gas
Rate
12/99
12/00
12/01
12/02
Time
R12 - Gas
Rate
12/99
12/00
12/01
12/02
Time
R11 - Wtr
Rate
12/99
12/00
12/01
12/02
Time
R12 - Wtr
Rate
12/99
12/00
12/01
12/02
Time
field
match
initial
preliminary

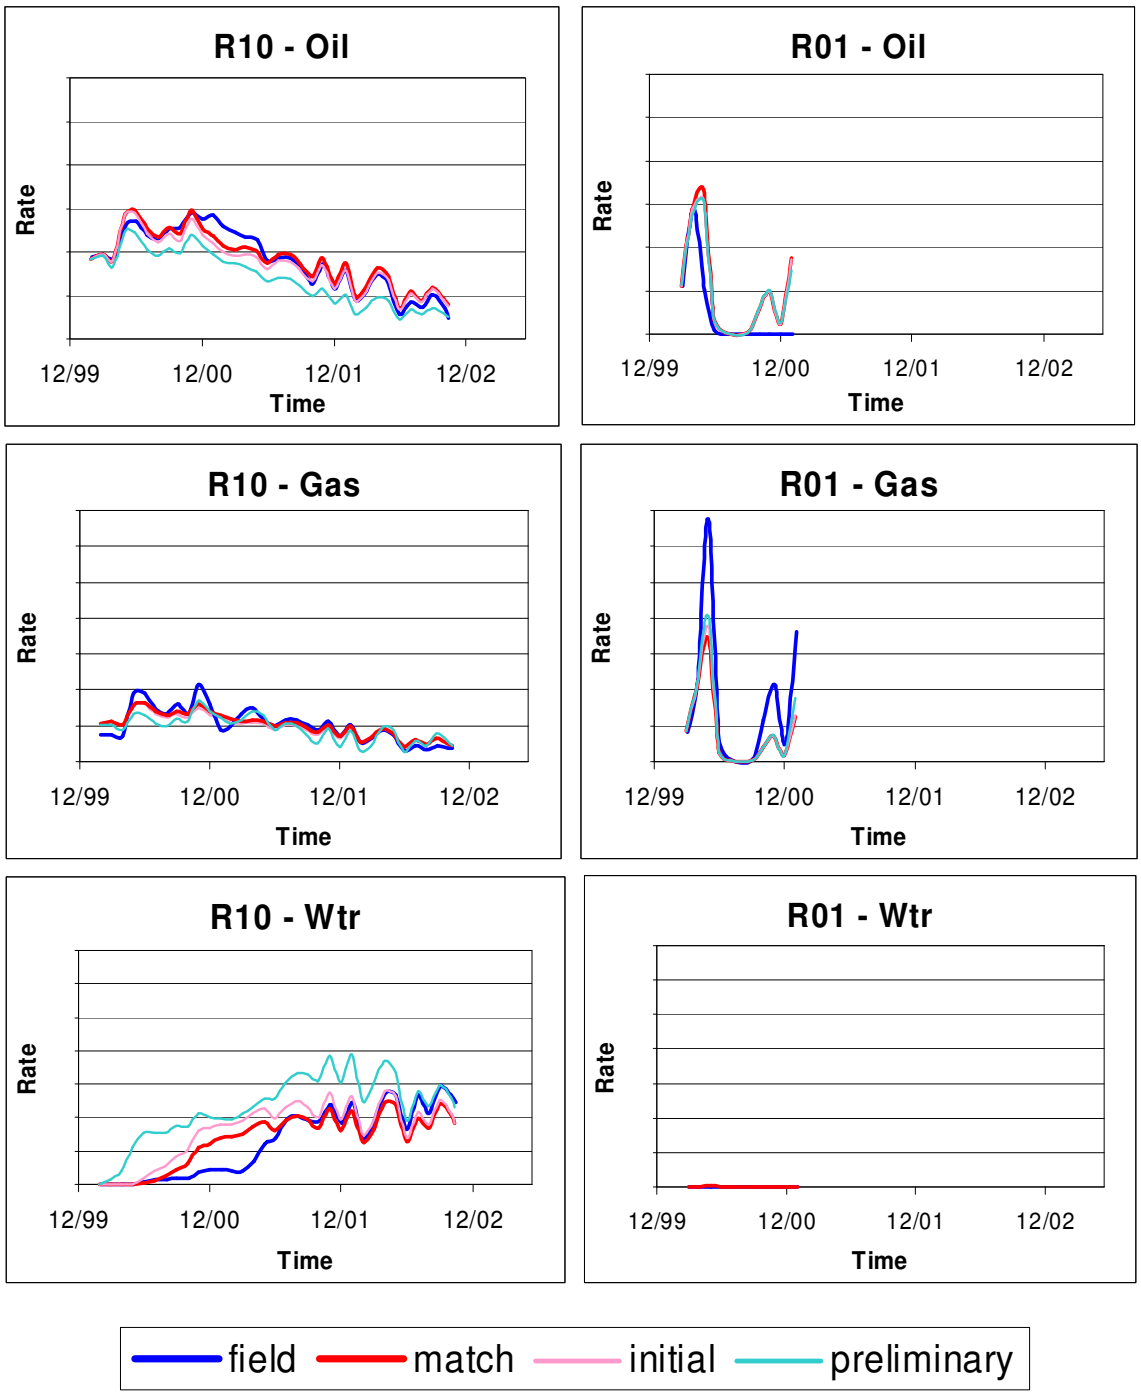
R10 - Oil
R01 - Oil
R10 - Gas
R01 - Gas
R10 - Wtr
R01 - Wtr
Rate
Time
12/99
12/00
12/01
12/02
field
match
initial
preliminary

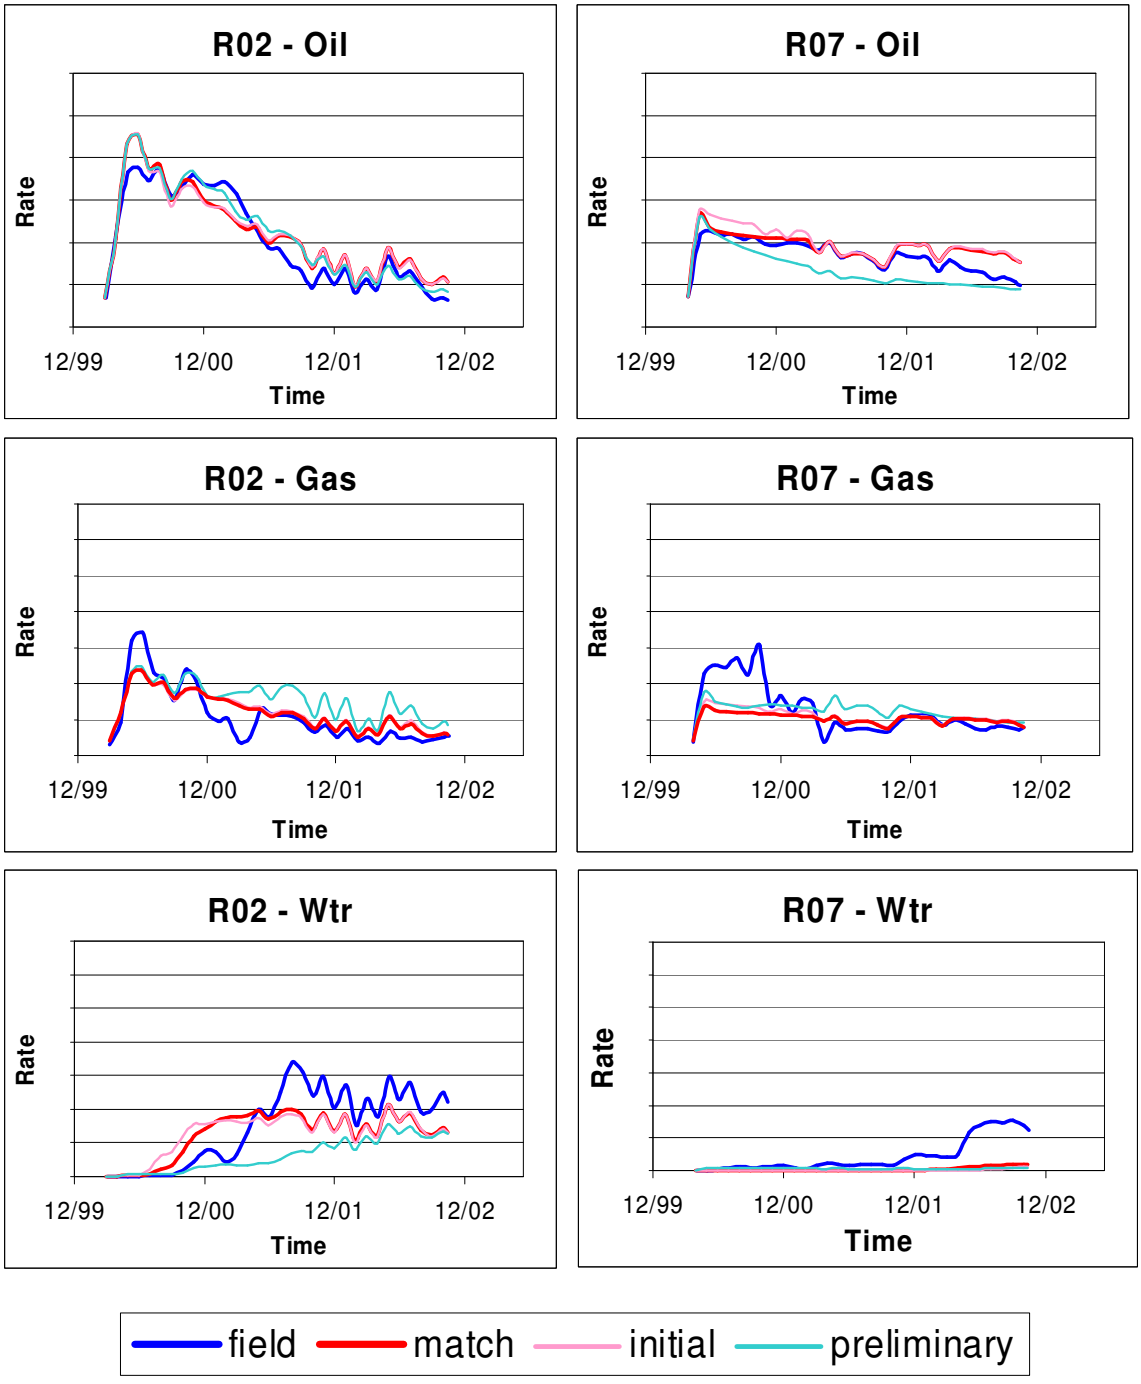
R02 - Oil
R07 - Oil
R02 - Gas
R07 - Gas
R02 - Wtr
R07 - Wtr
Rate
Time
12/99
12/00
12/01
12/02
field
match
initial
preliminary

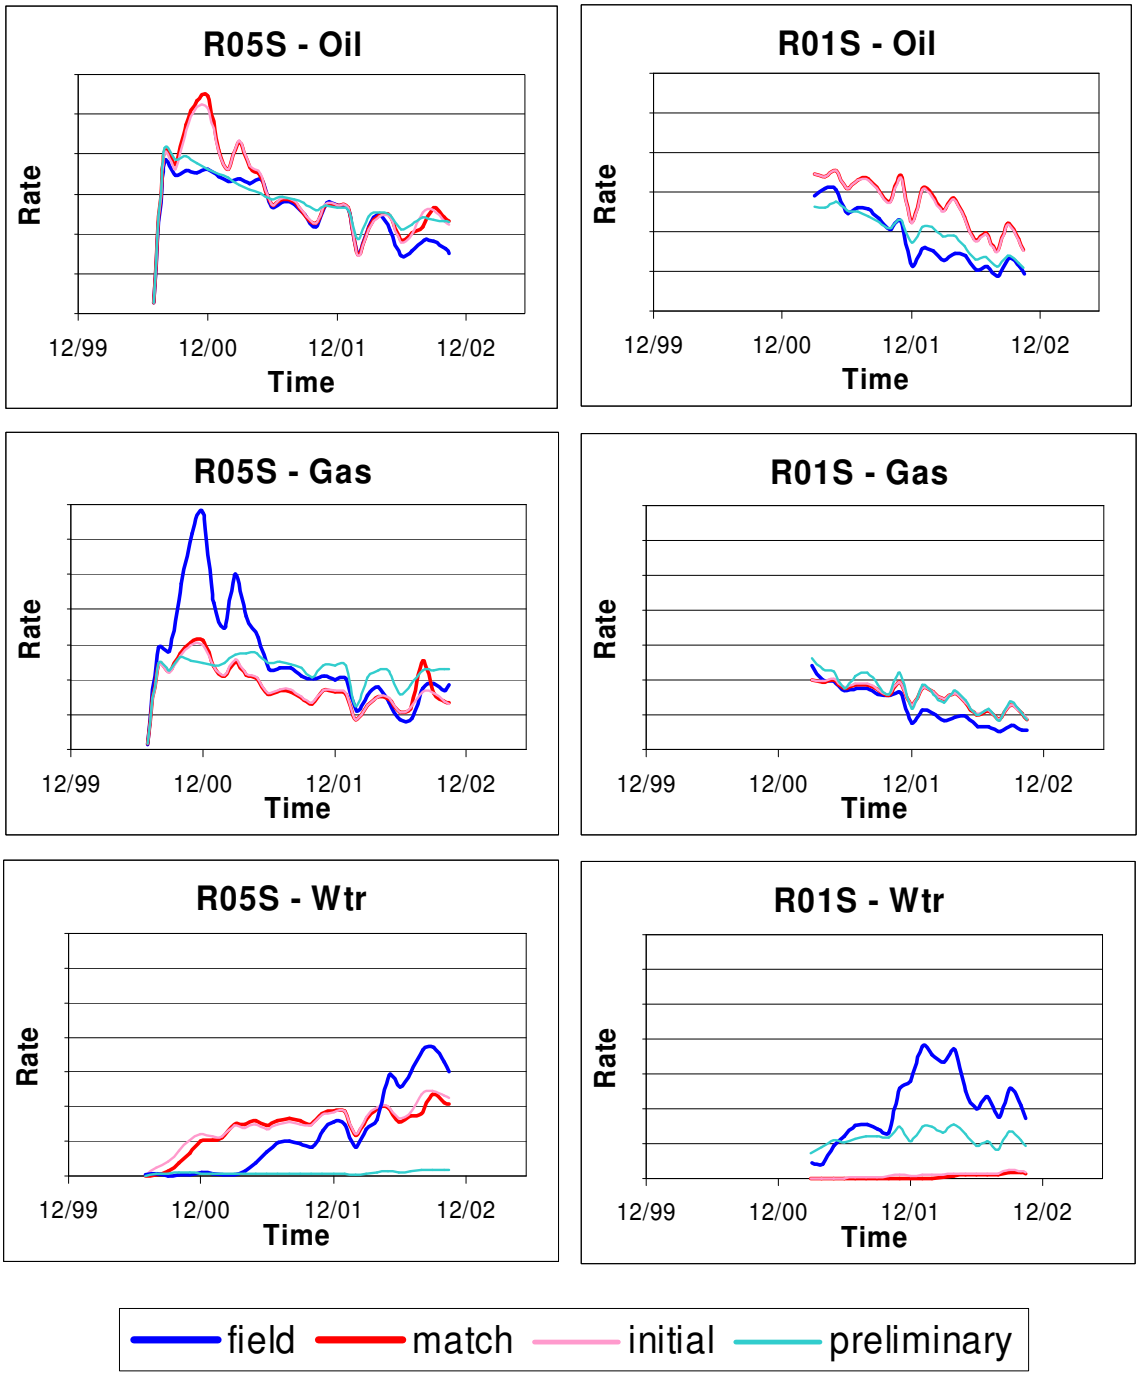
R05S - Oil
R01S - Oil
R05S - Gas
R01S - Gas
R05S - Wtr
R01S - Wtr
Rate
Time
12/99
12/00
12/01
12/02
field
match
initial
preliminary

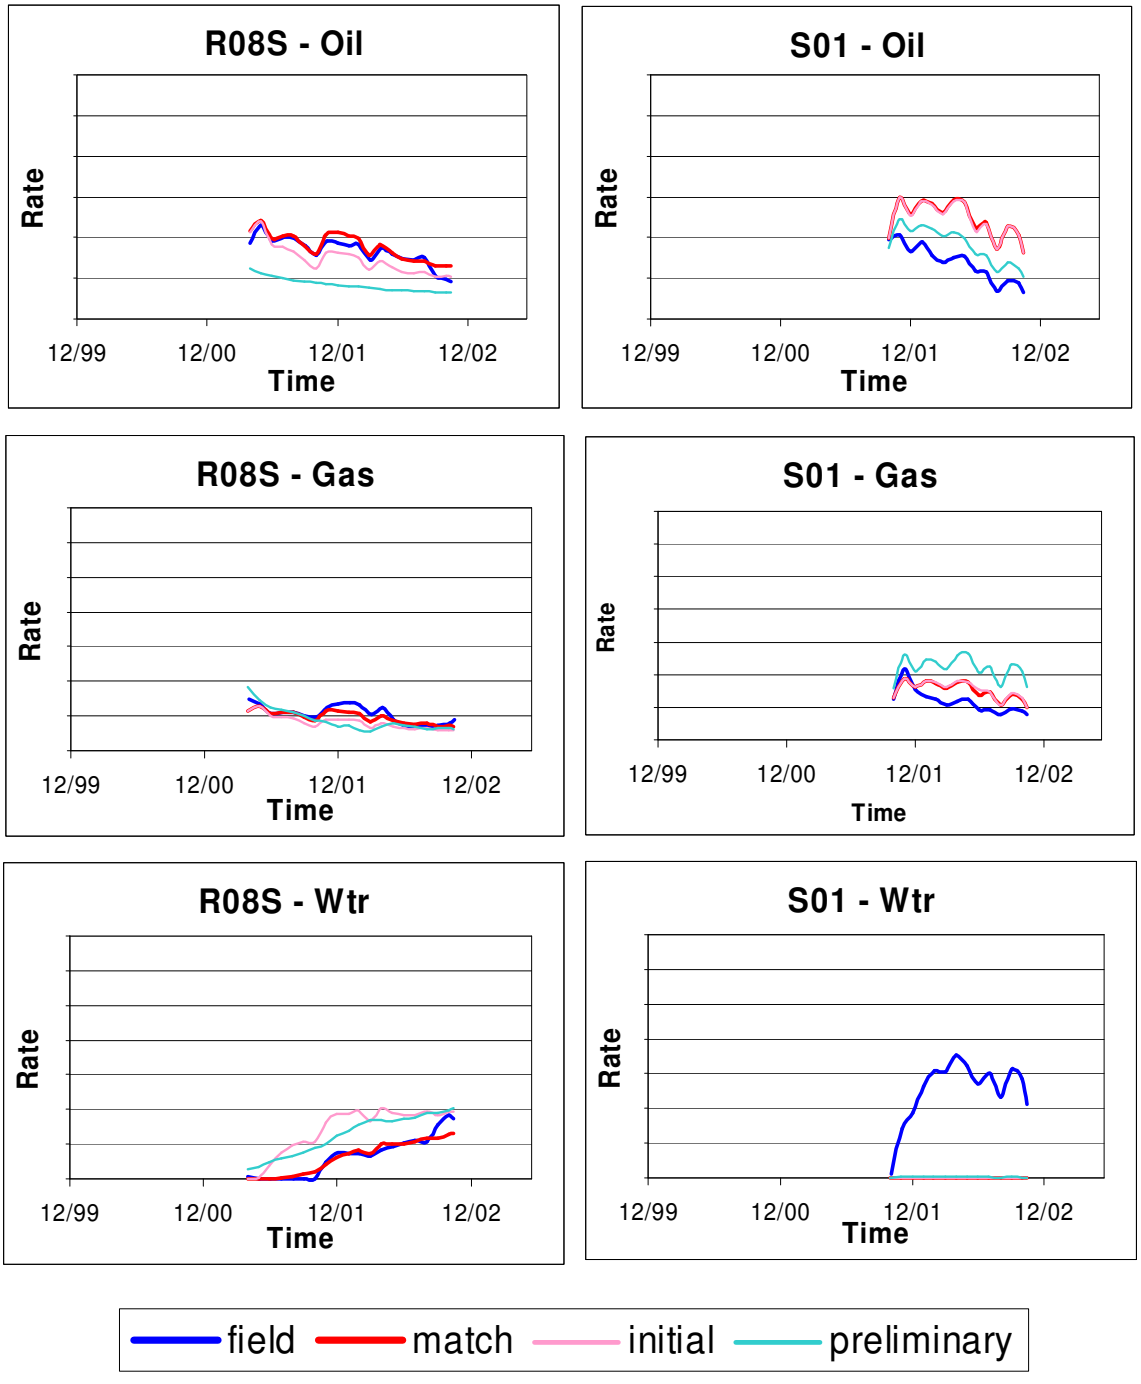
R08S - Oil
S01 - Oil
R08S - Gas
S01 - Gas
R08S - Wtr
S01 - Wtr
Rate
Time
12/99
12/00
12/01
12/02
field
match
initial
preliminary

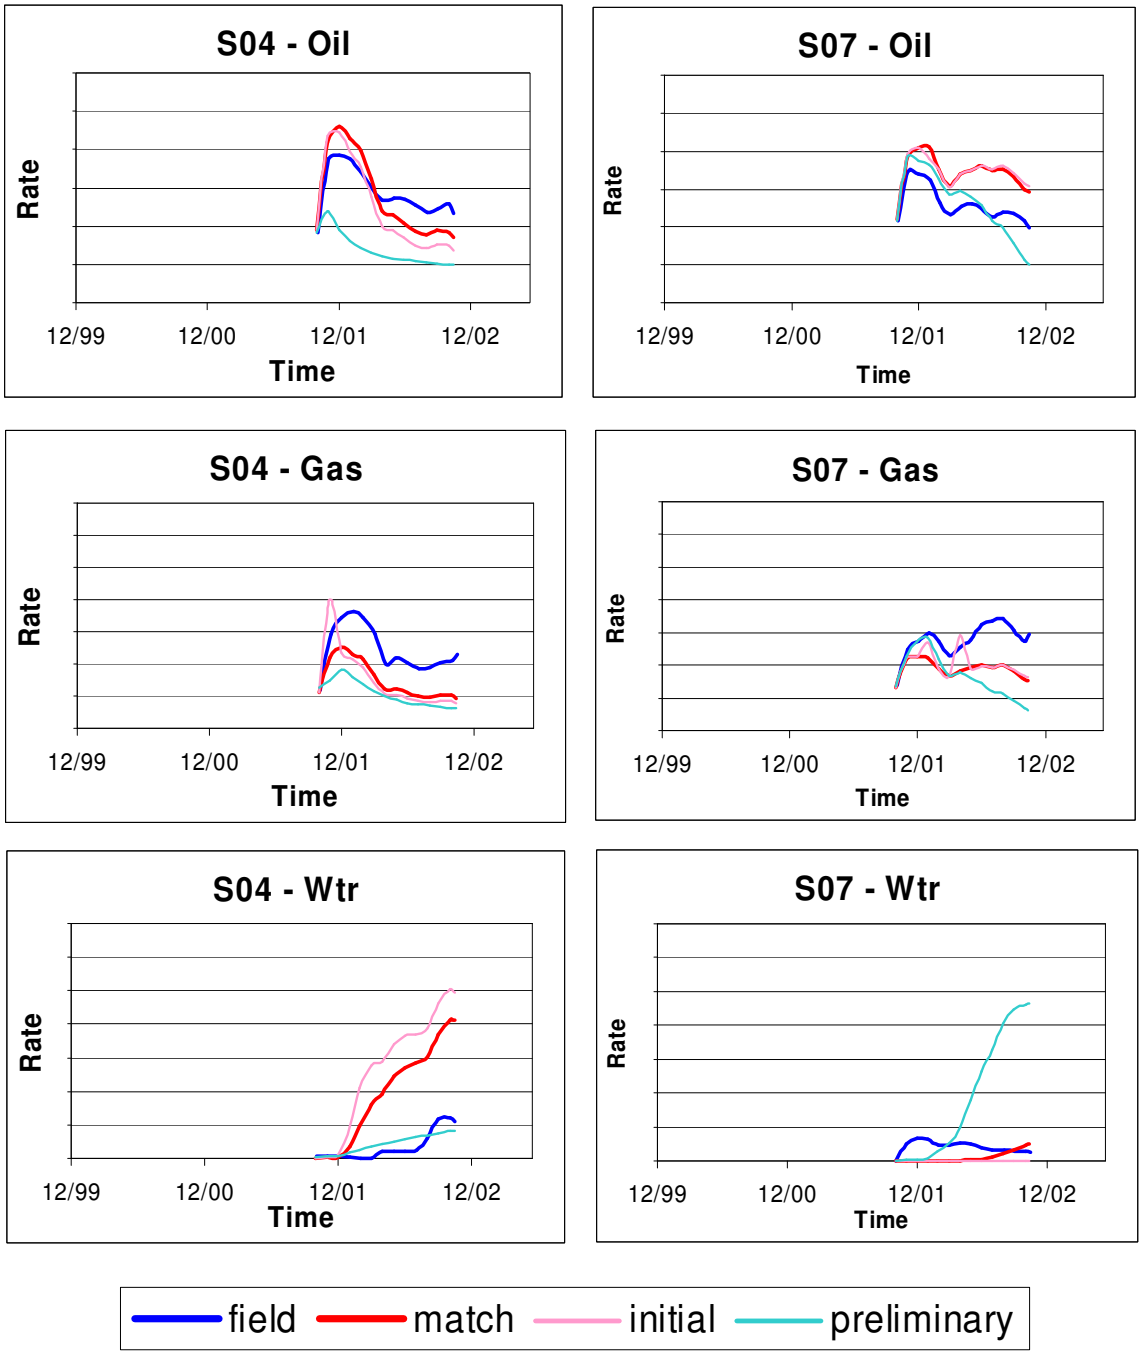
S04 - Oil
S07 - Oil
S04 - Gas
S07 - Gas
S04 - Wtr
S07 - Wtr
Rate
Time
12/99
12/00
12/01
12/02
field
match
initial
preliminary

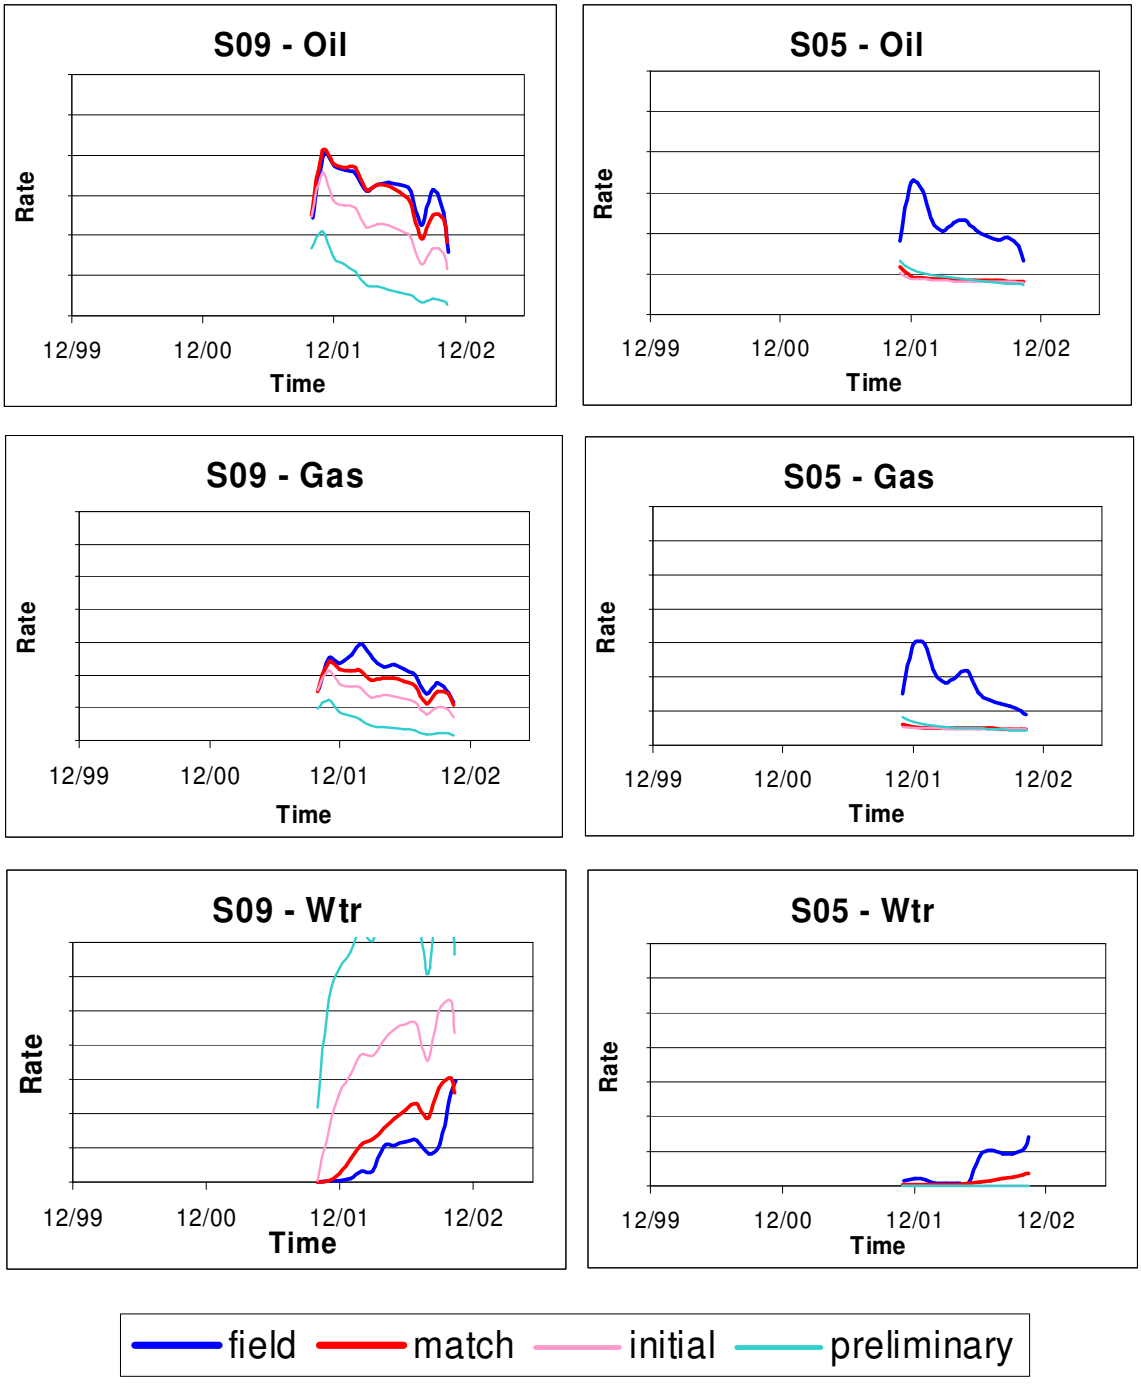
S09 - Oil
S05 - Oil
S09 - Gas
S05 - Gas
S09 - Wtr
S05 - Wtr
Rate
Time
12/99
12/00
12/01
12/02
field
match
initial
preliminary

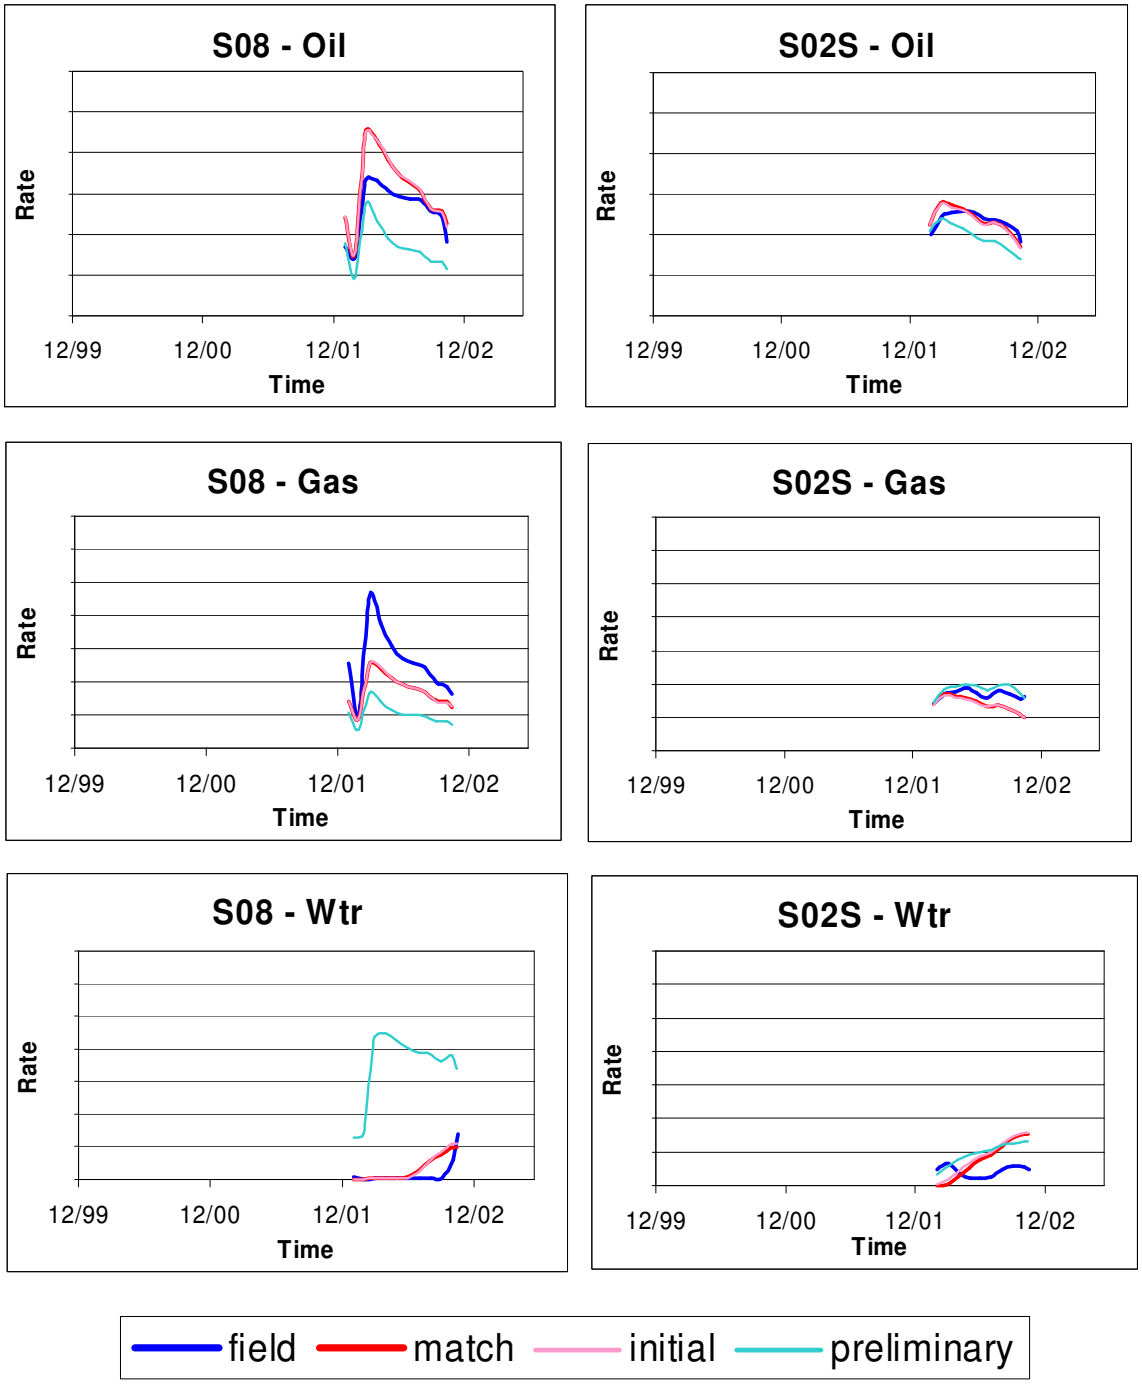
S08 - Oil
S02S - Oil
S08 - Gas
S02S - Gas
S08 - Wtr
S02S - Wtr
Rate
Time
12/99
12/00
12/01
12/02
field
match
initial
preliminary

C. Validation Runs for Six Models for WCA Reservoir

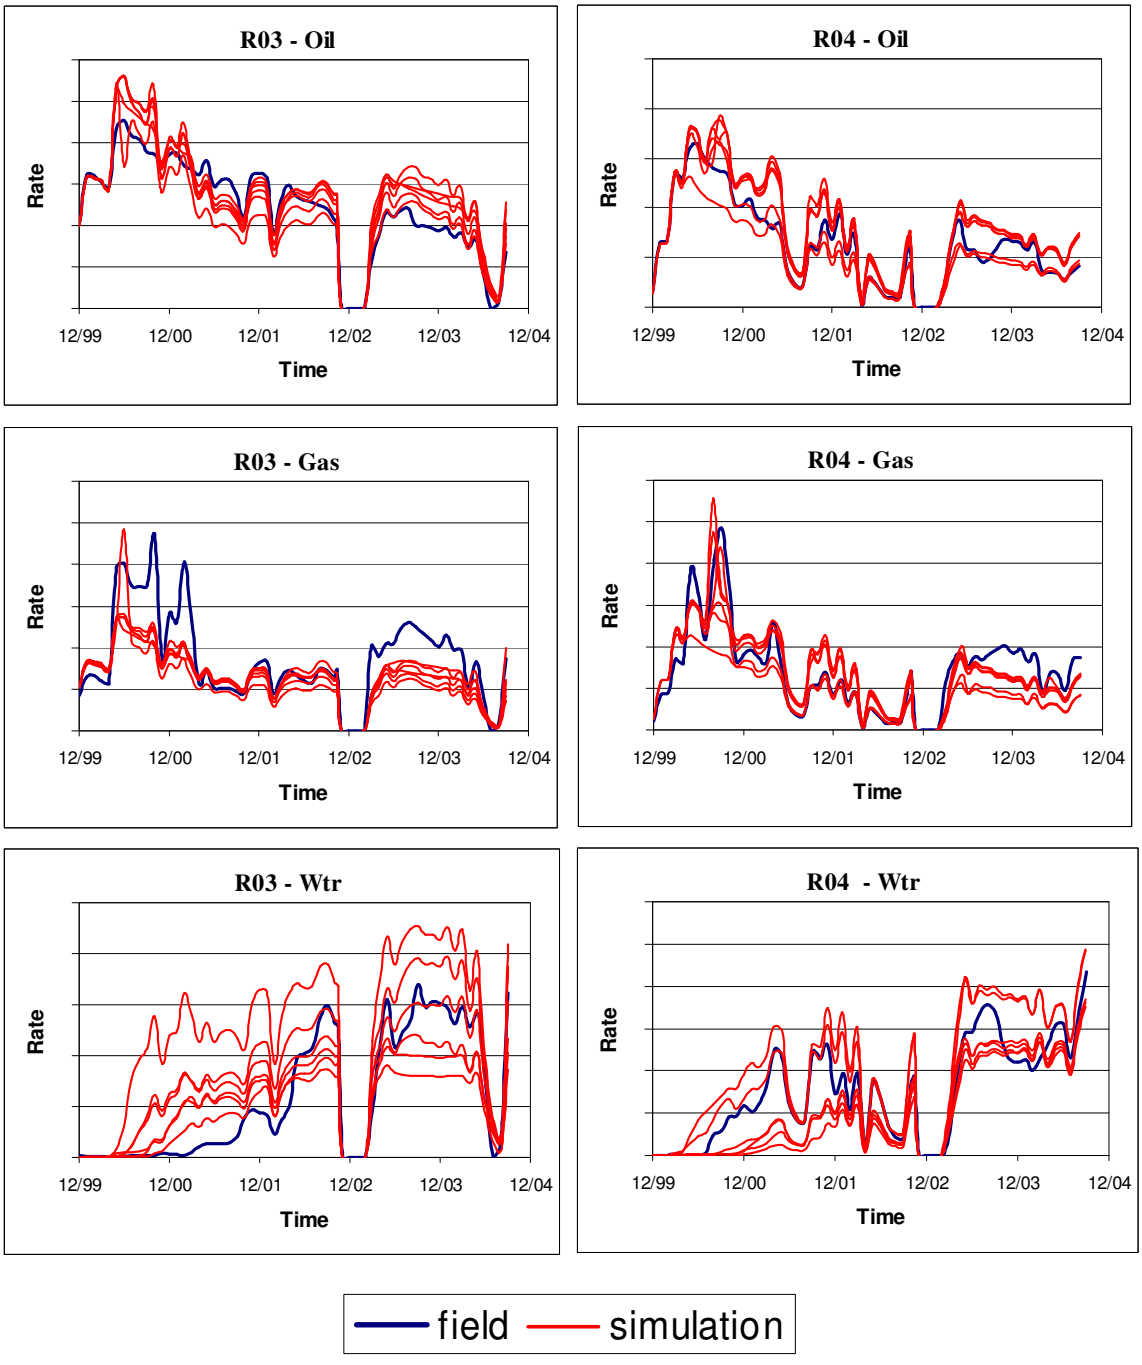

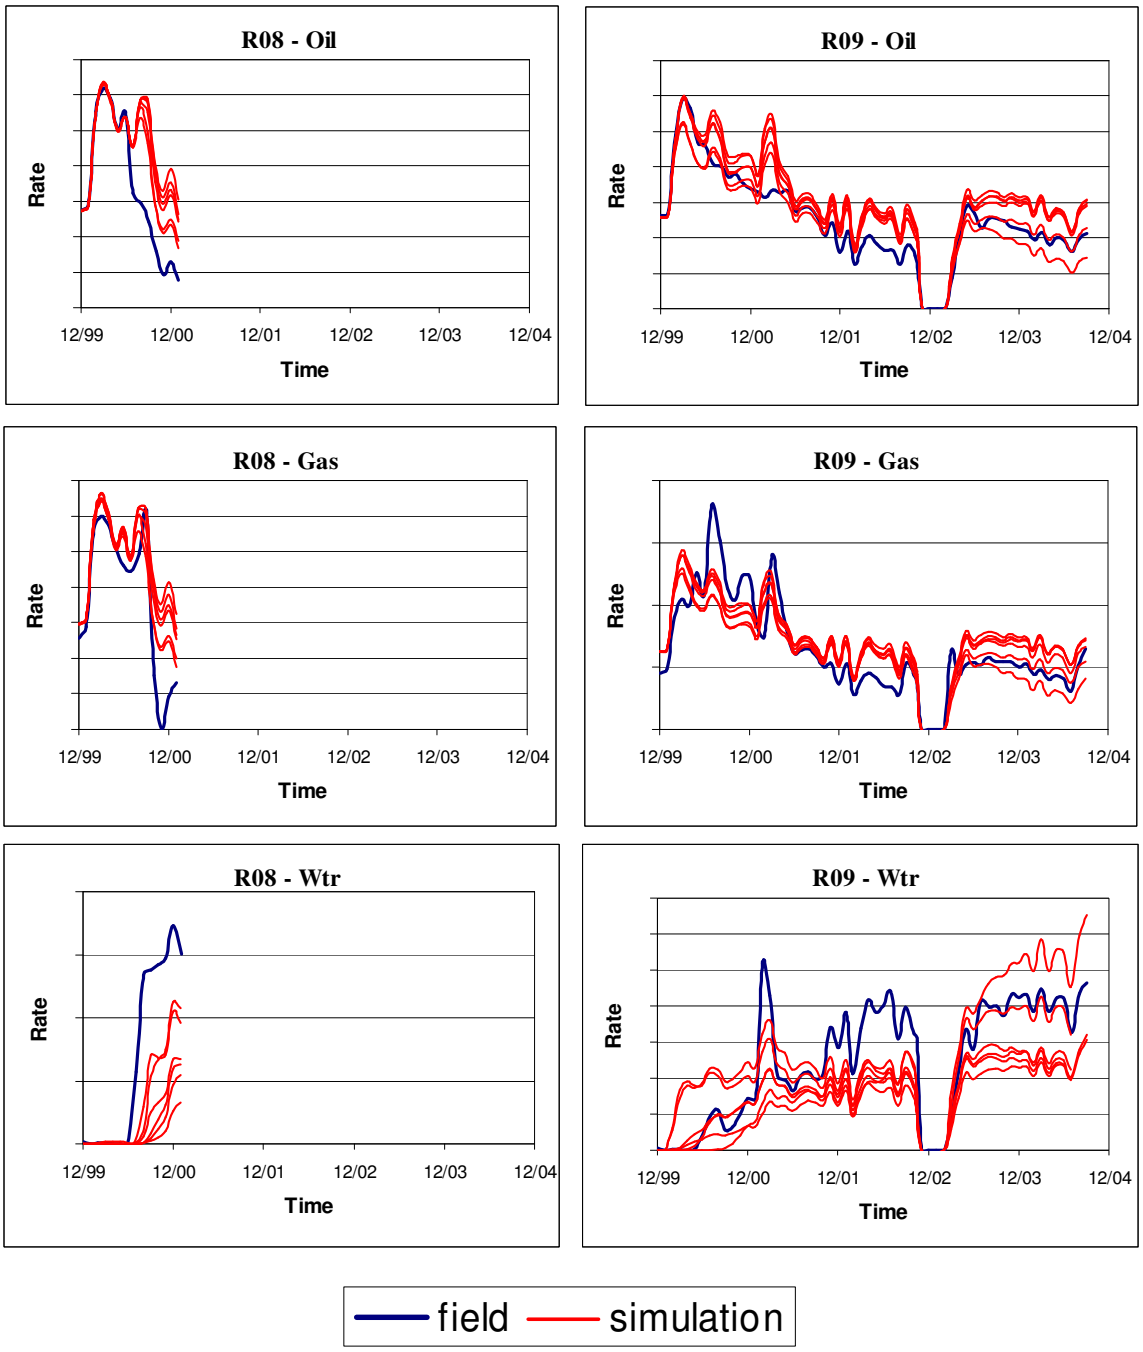
R08 - Oil
R09 - Oil
R08 - Gas
R09 - Gas
R08 - Wtr
R09 - Wtr
Rate
Time
12/99
12/00
12/01
12/02
12/03
12/04
field
simulation

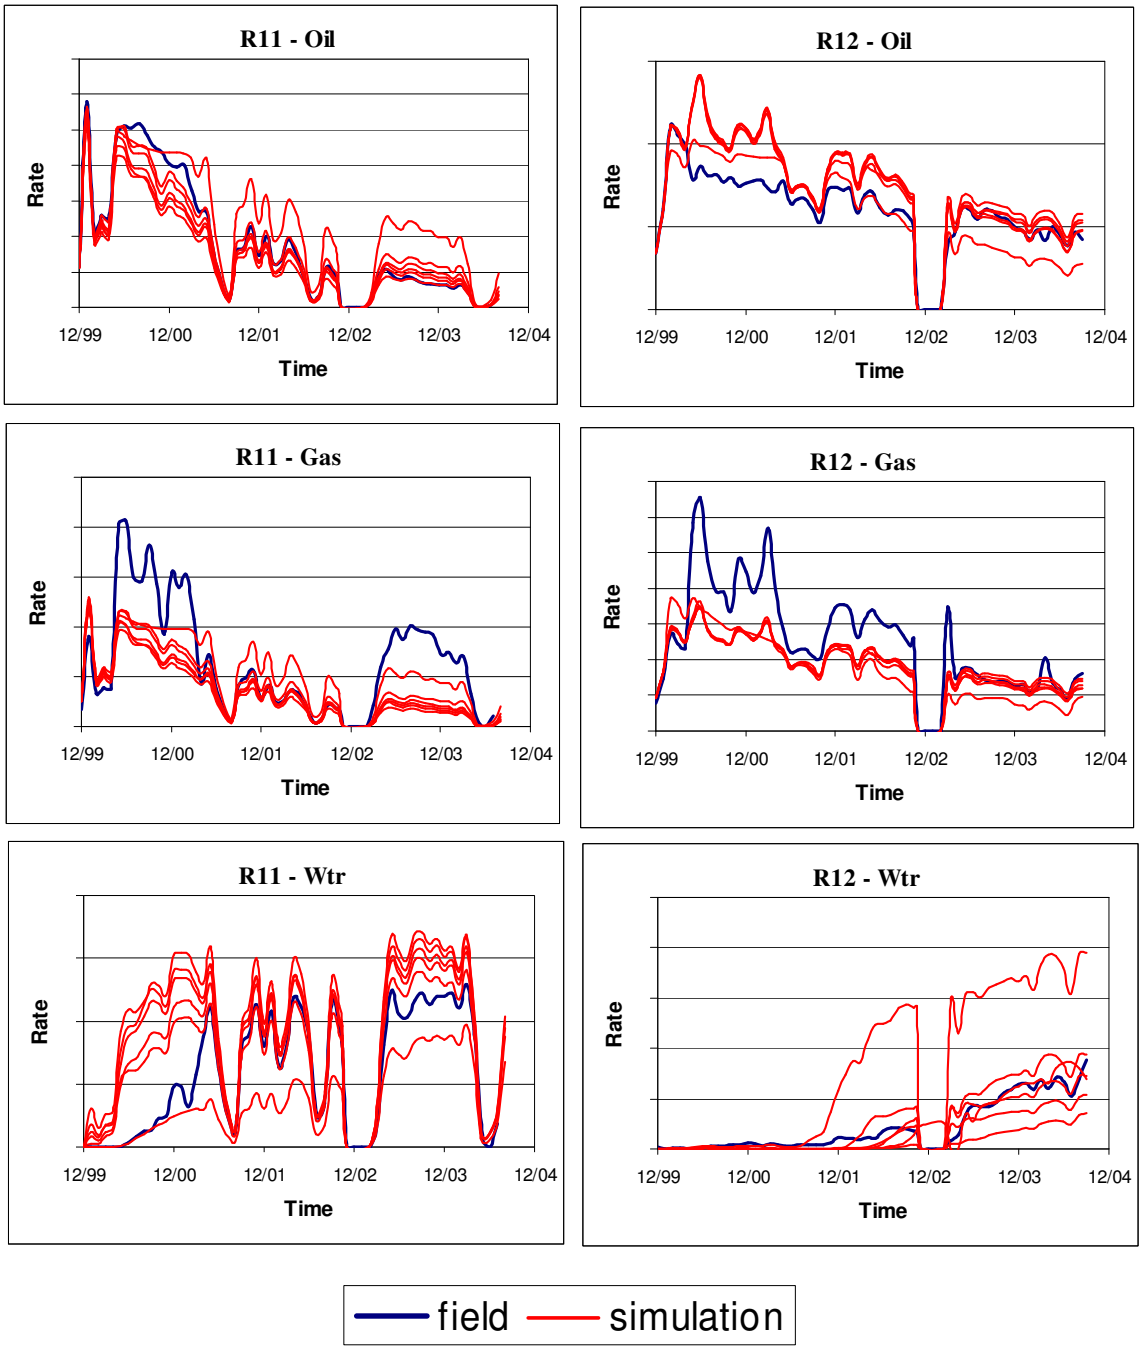
R11 - Oil
R12 - Oil
R11 - Gas
R12 - Gas
R11 - Wtr
R12 - Wtr
Rate
Time
12/99
12/00
12/01
12/02
12/03
12/04
field
simulation

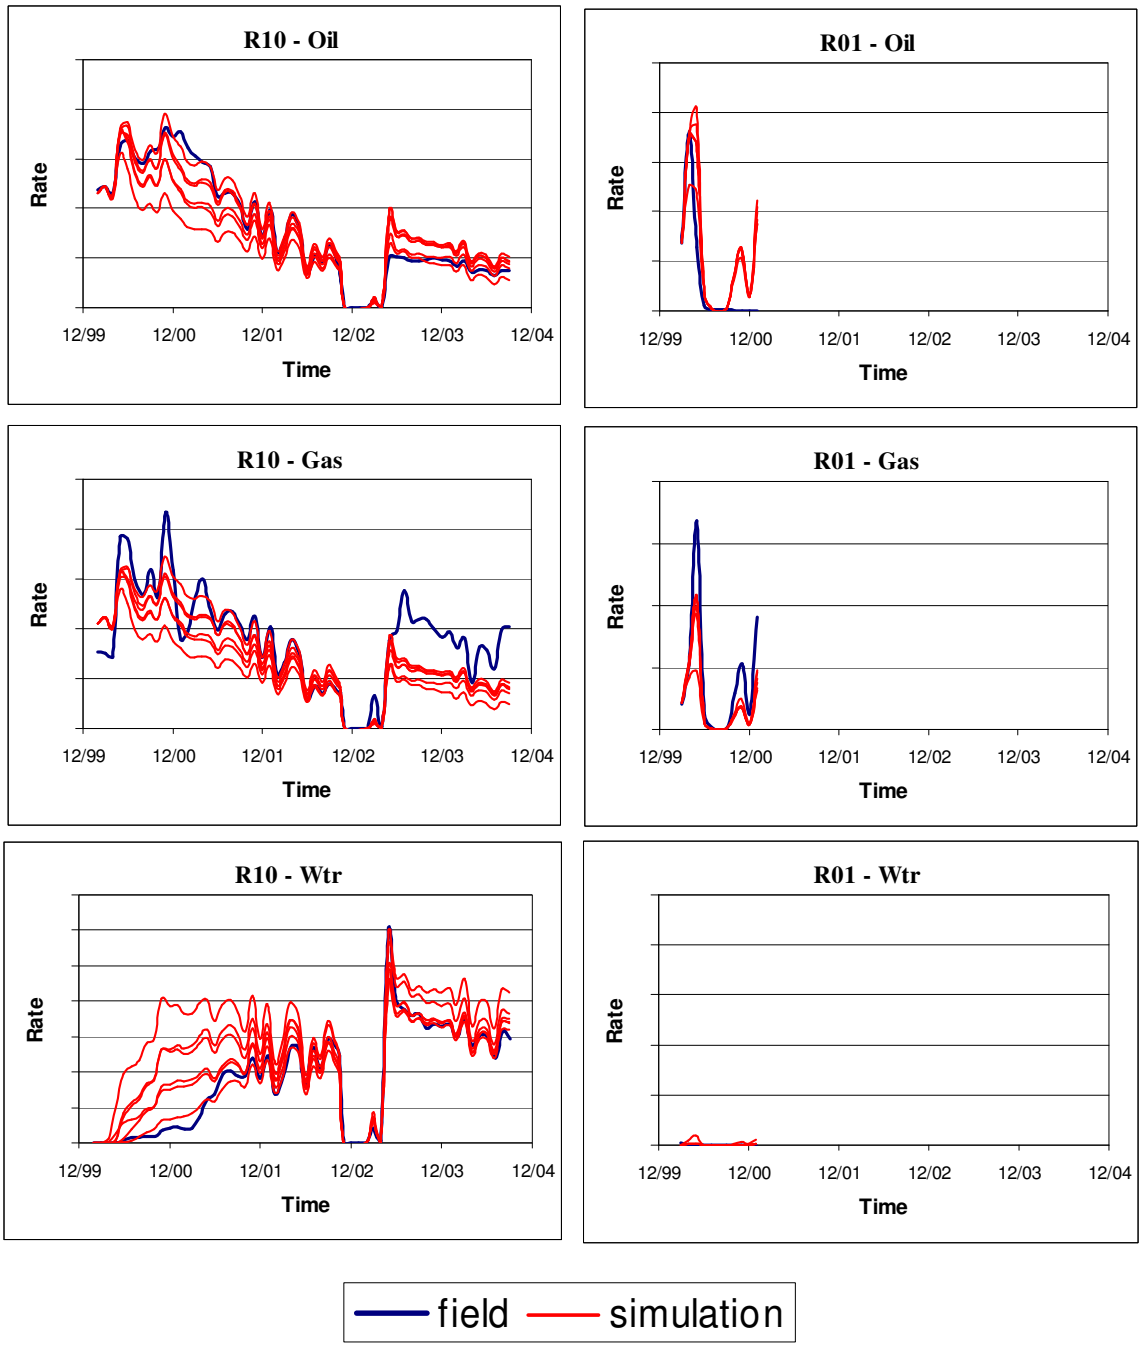
R10 - Oil
R01 - Oil
R10 - Gas
R01 - Gas
R10 - Wtr
R01 - Wtr
Rate
Time
12/99
12/00
12/01
12/02
12/03
12/04
field
simulation

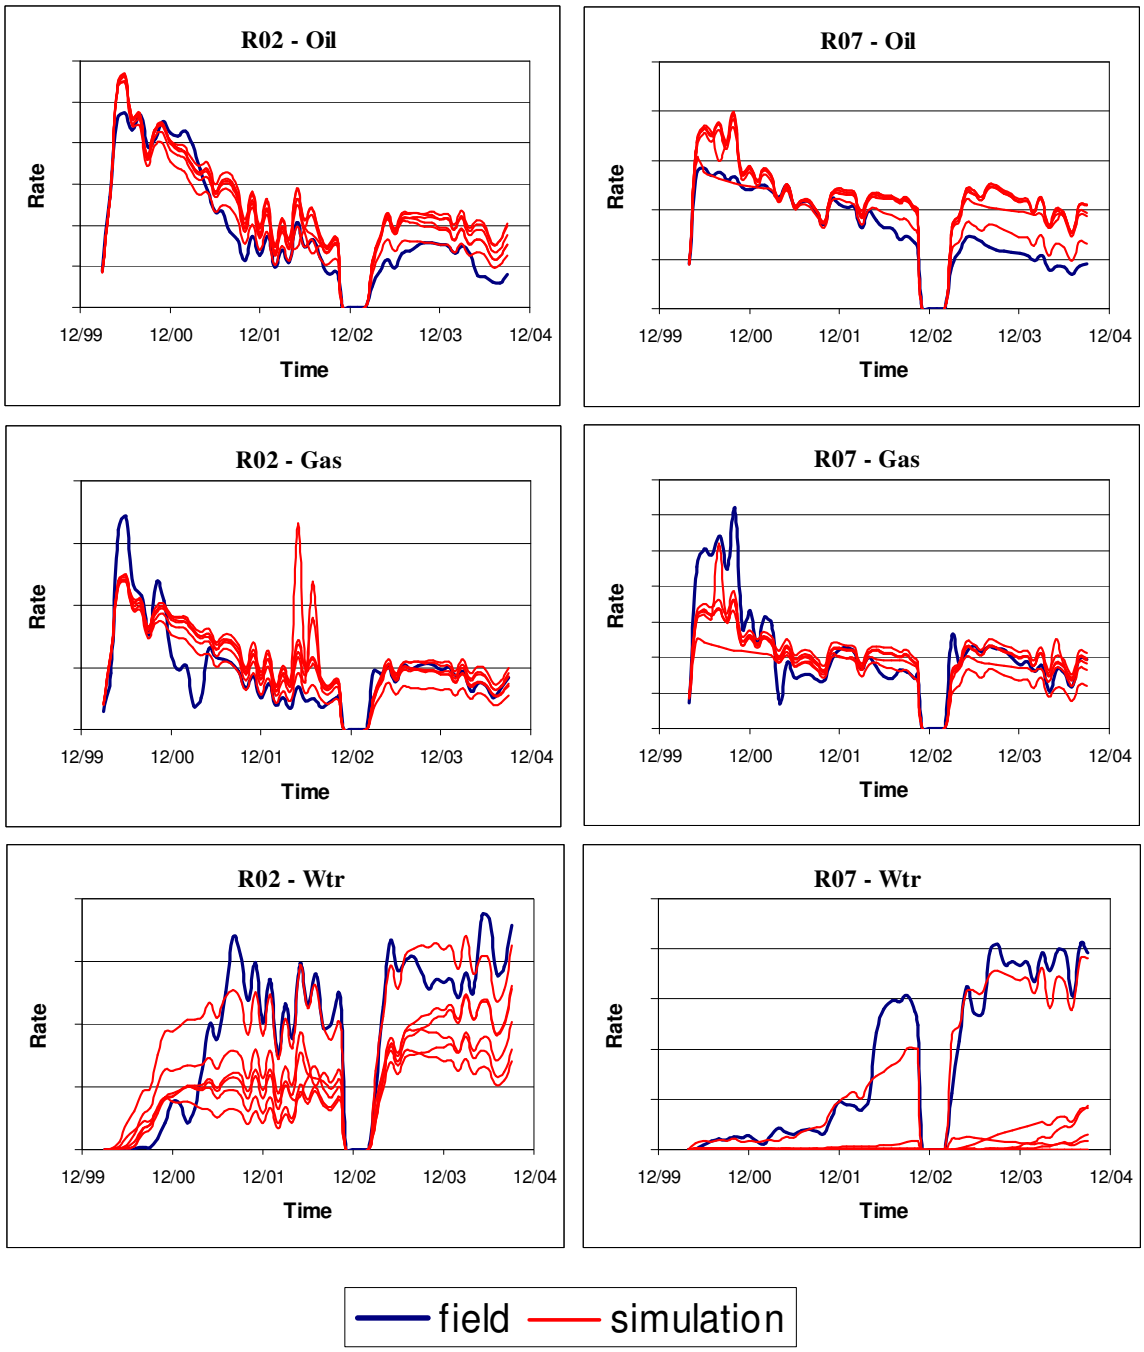
R02 - Oil
R07 - Oil
R02 - Gas
R07 - Gas
R02 - Wtr
R07 - Wtr
Rate
Time
12/99
12/00
12/01
12/02
12/03
12/04
field
simulation

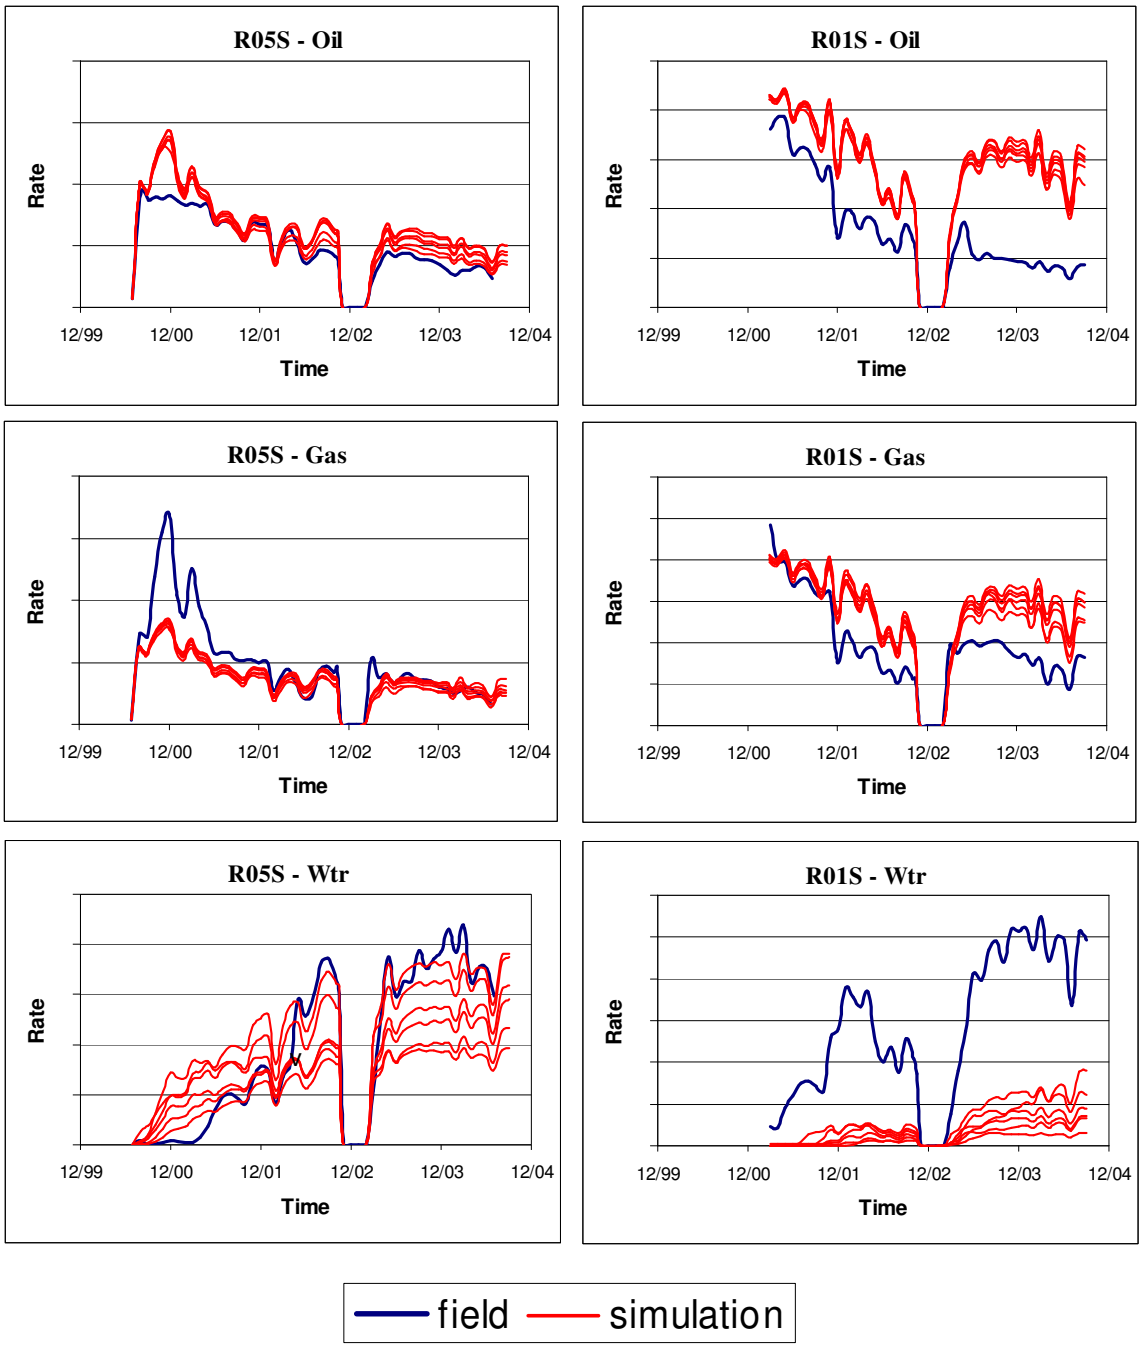
R05S - Oil
R01S - Oil
R05S - Gas
R01S - Gas
R05S - Wtr
R01S - Wtr
Rate
Time
12/99
12/00
12/01
12/02
12/03
12/04
field
simulation

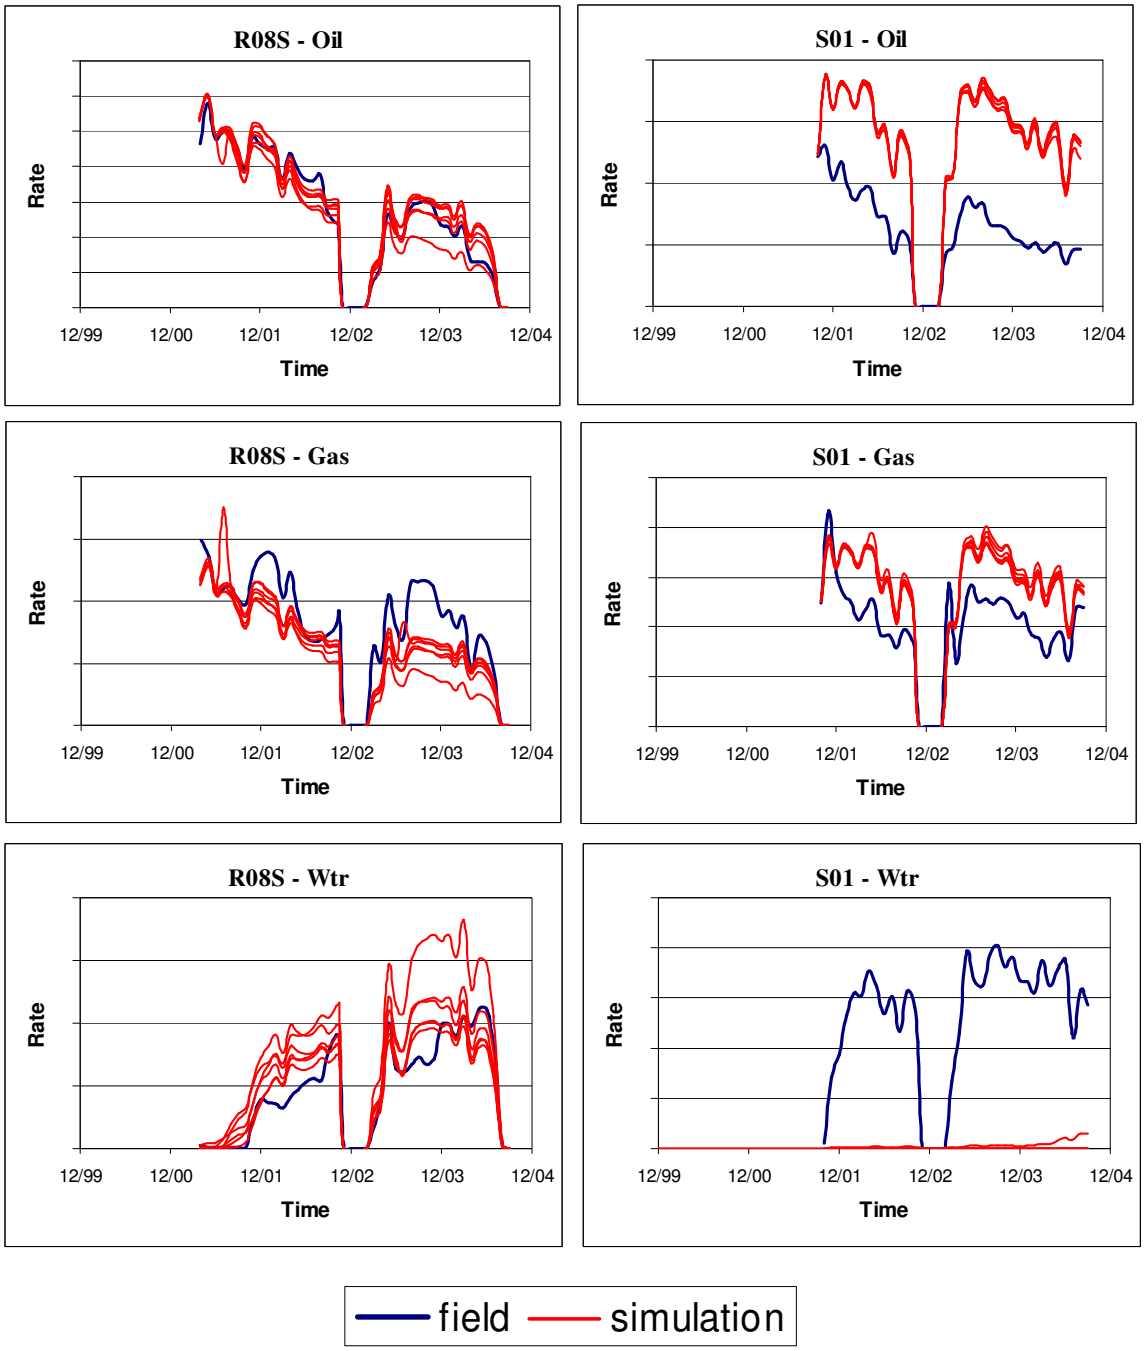
R08S - Oil
S01 - Oil
R08S - Gas
S01 - Gas
R08S - Wtr
S01 - Wtr
Rate
Time
12/99
12/00
12/01
12/02
12/03
12/04
field
simulation

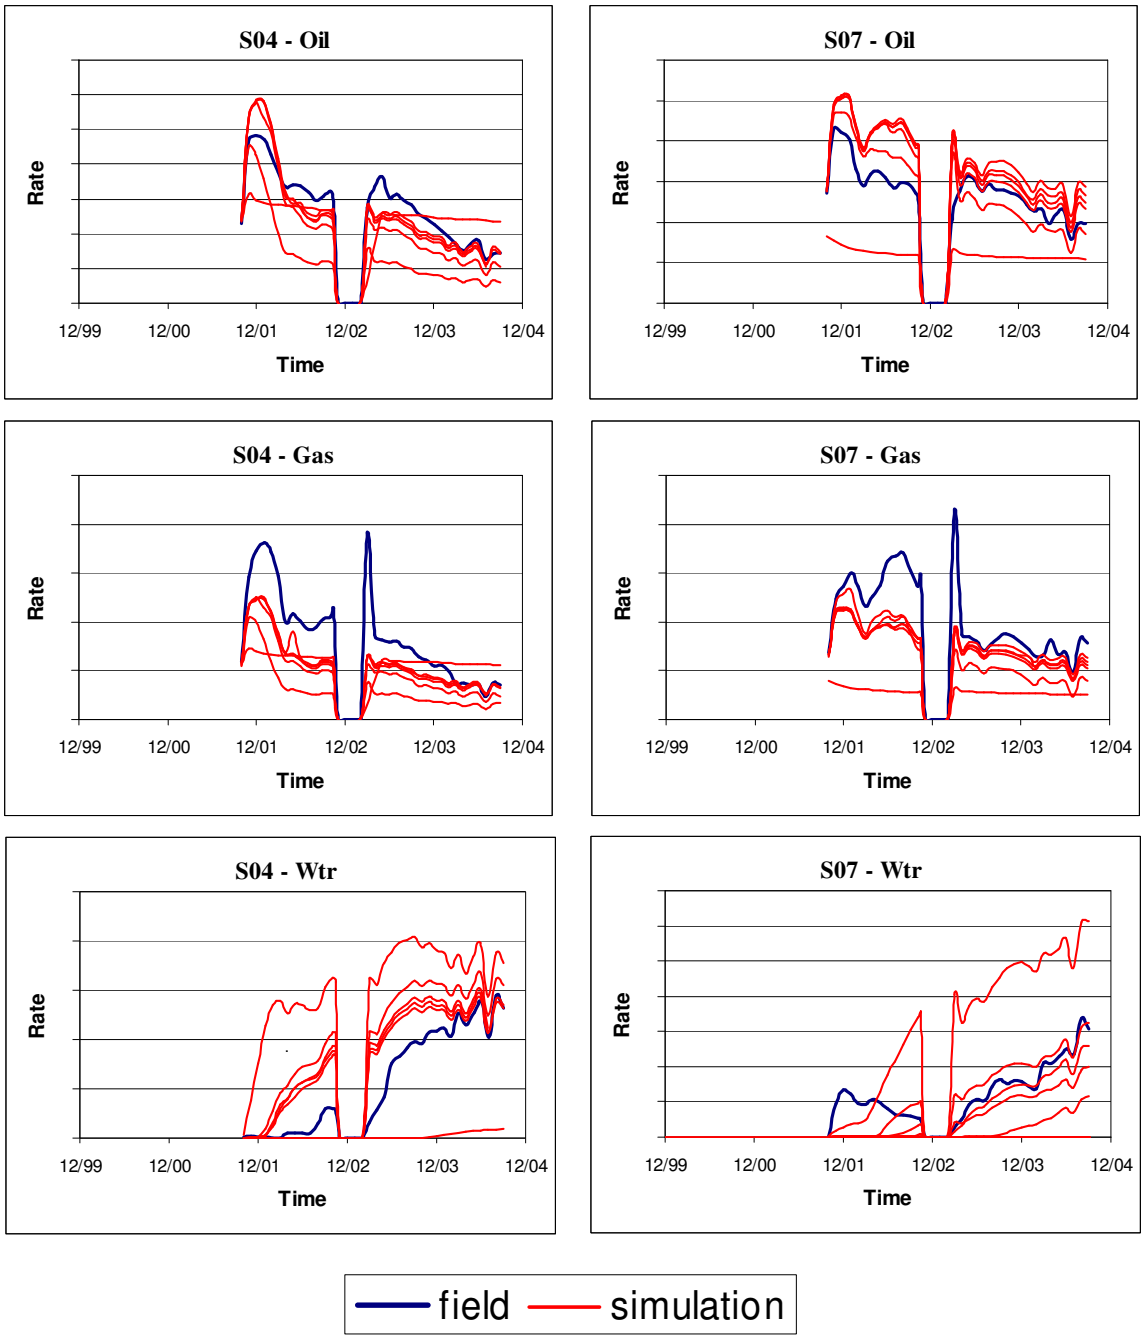
S04 - Oil
S07 - Oil
S04 - Gas
S07 - Gas
S04 - Wtr
S07 - Wtr
Rate
Time
12/99
12/00
12/01
12/02
12/03
12/04
field
simulation

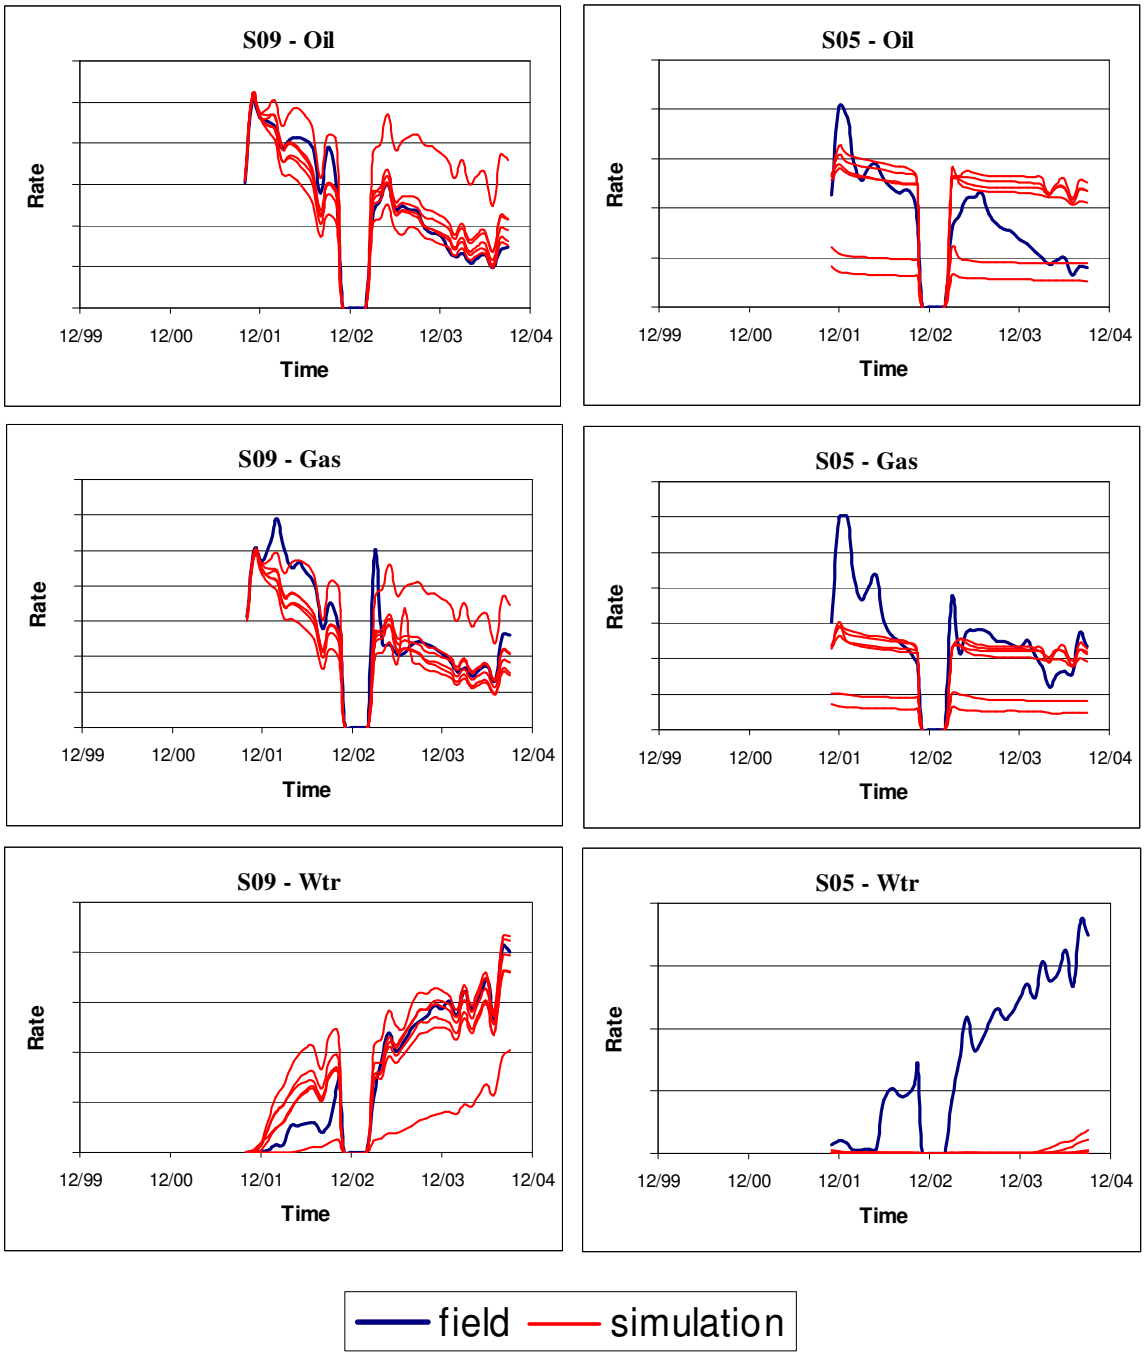
S09 - Oil
S05 - Oil
S09 - Gas
S05 - Gas
S09 - Wtr
S05 - Wtr
Rate
Time
12/99
12/00
12/01
12/02
12/03
12/04
field
simulation

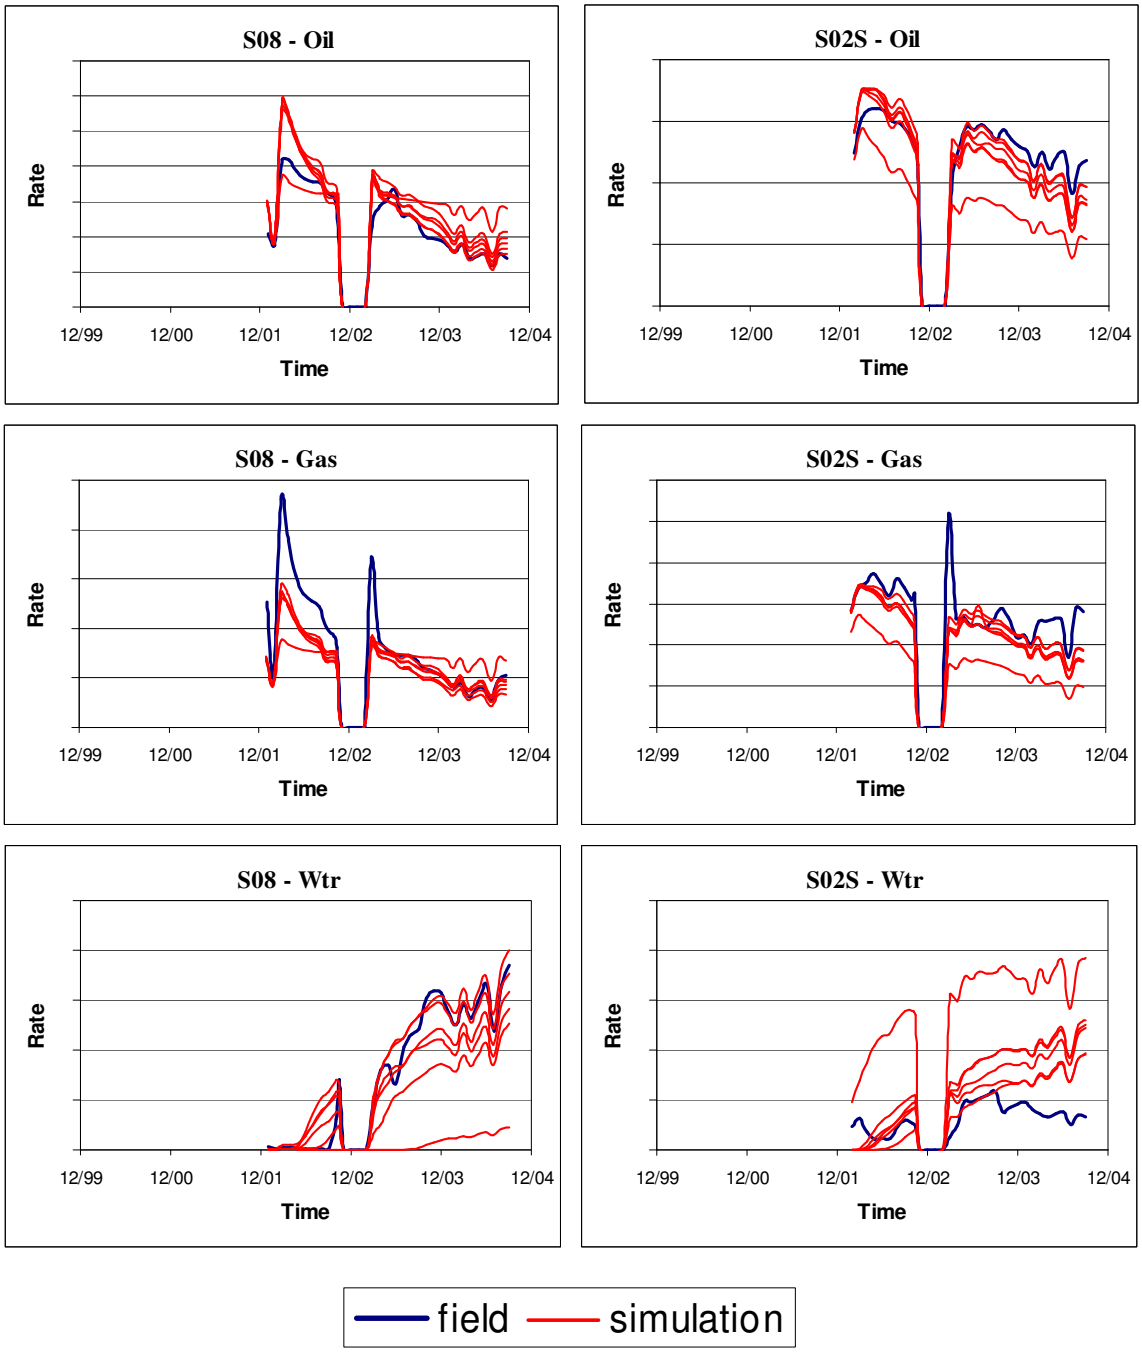
S08 - Oil
S02S - Oil
S08 - Gas
S02S - Gas
S08 - Wtr
S02S - Wtr
Rate
Time
12/99
12/00
12/01
12/02
12/03
12/04
field
simulation

D. Rates and Pressures for North Sea Reservoir

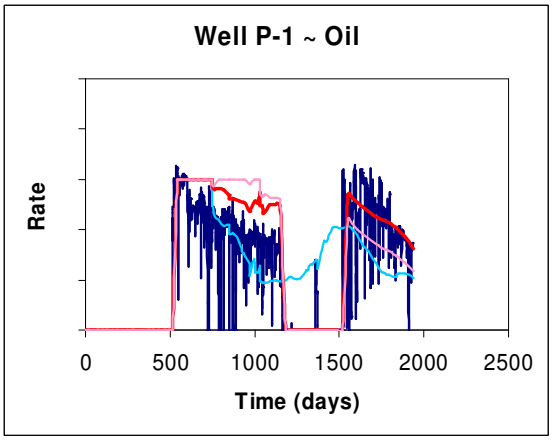

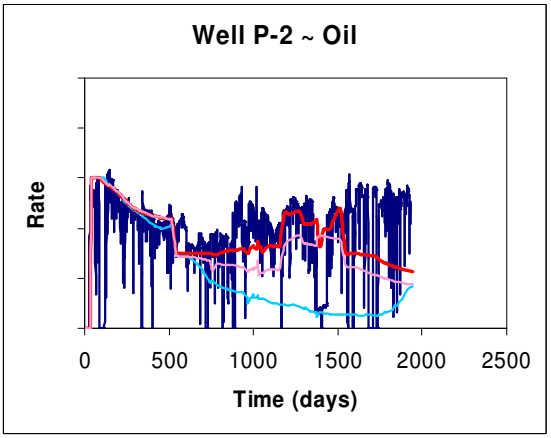

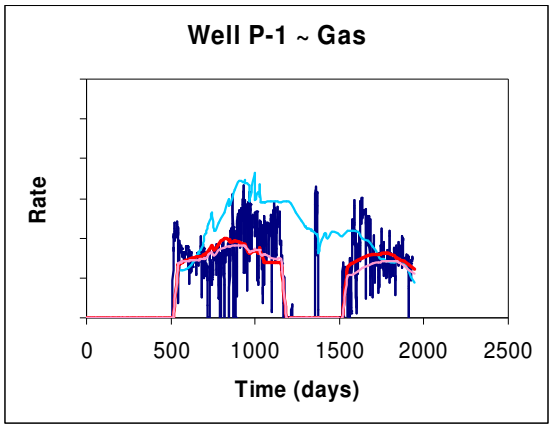

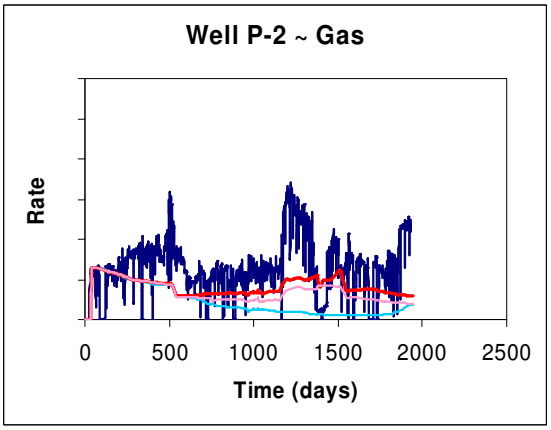

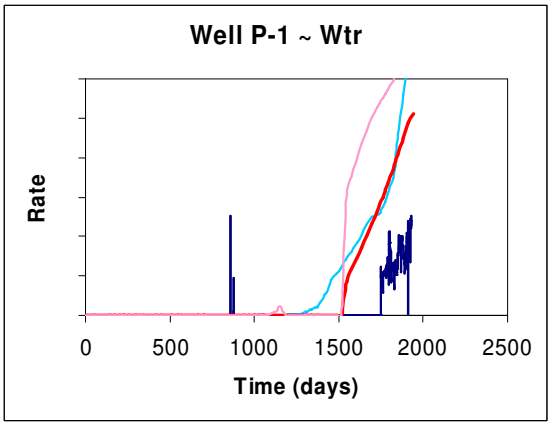

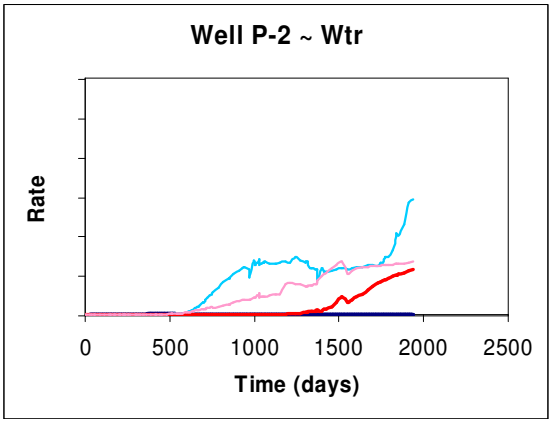

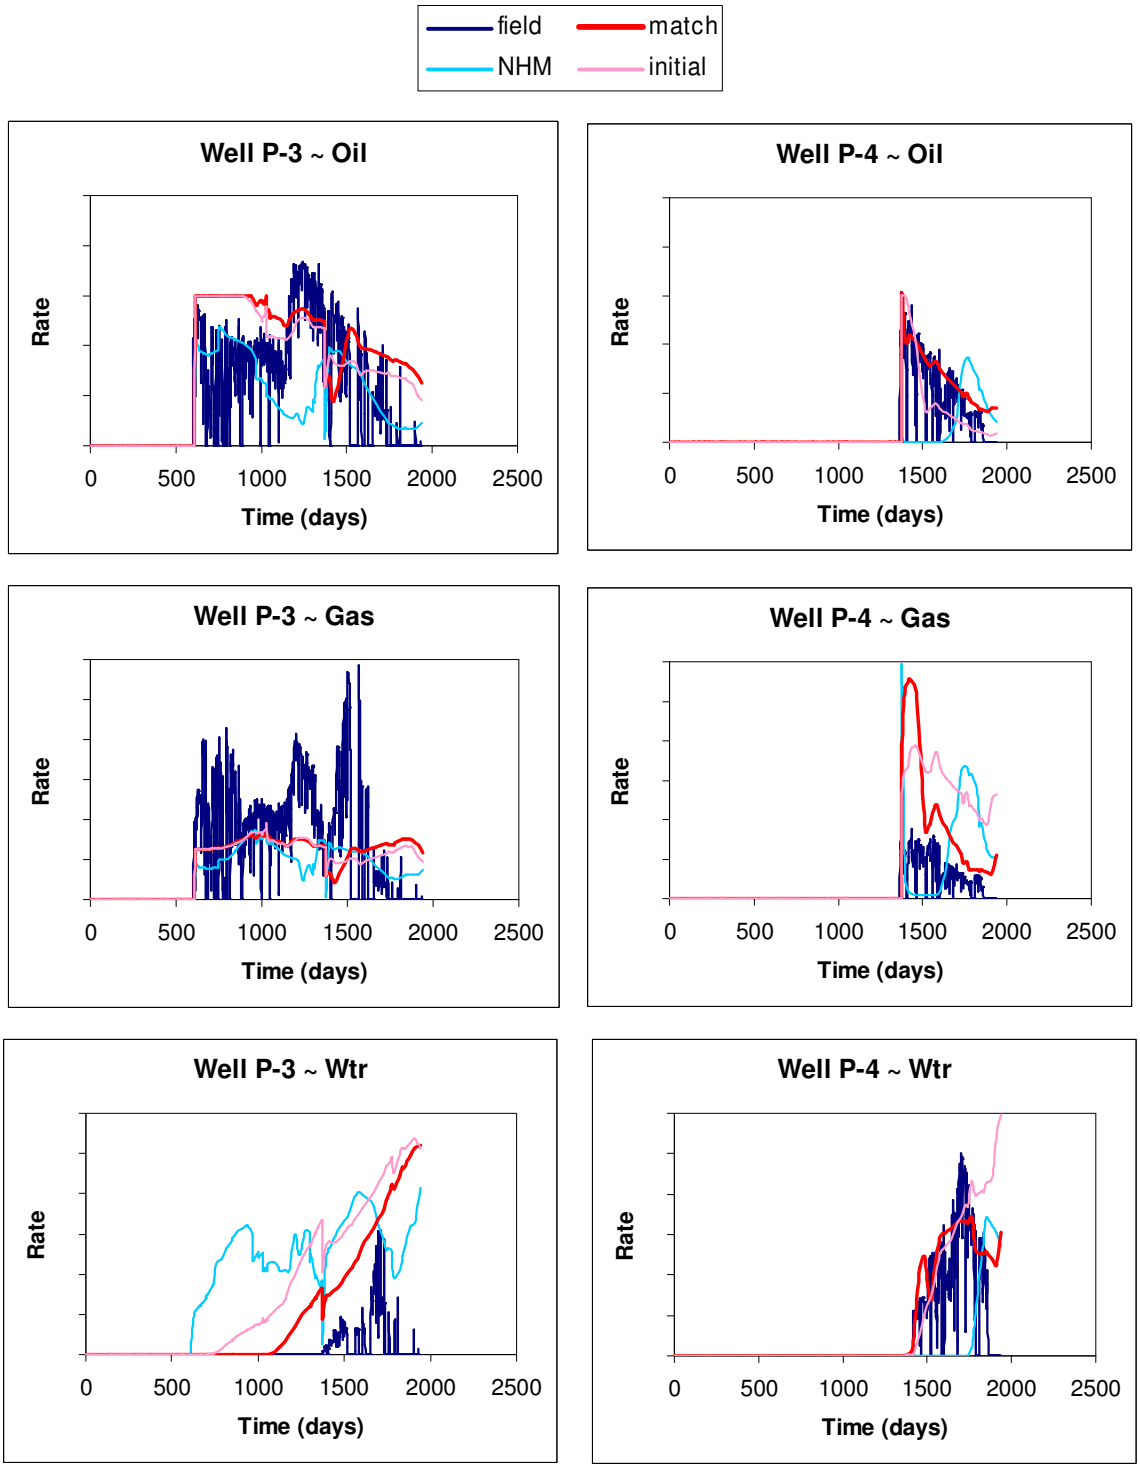
field
match
NHM
initial
Well P-3 ~ Oil
Well P-4 ~ Oil
Well P-3 ~ Gas
Well P-4 ~ Gas
Well P-3 ~ Wtr
Well P-4 ~ Wtr
Rate
Time (days)
0
500
1000
1500
2000
2500

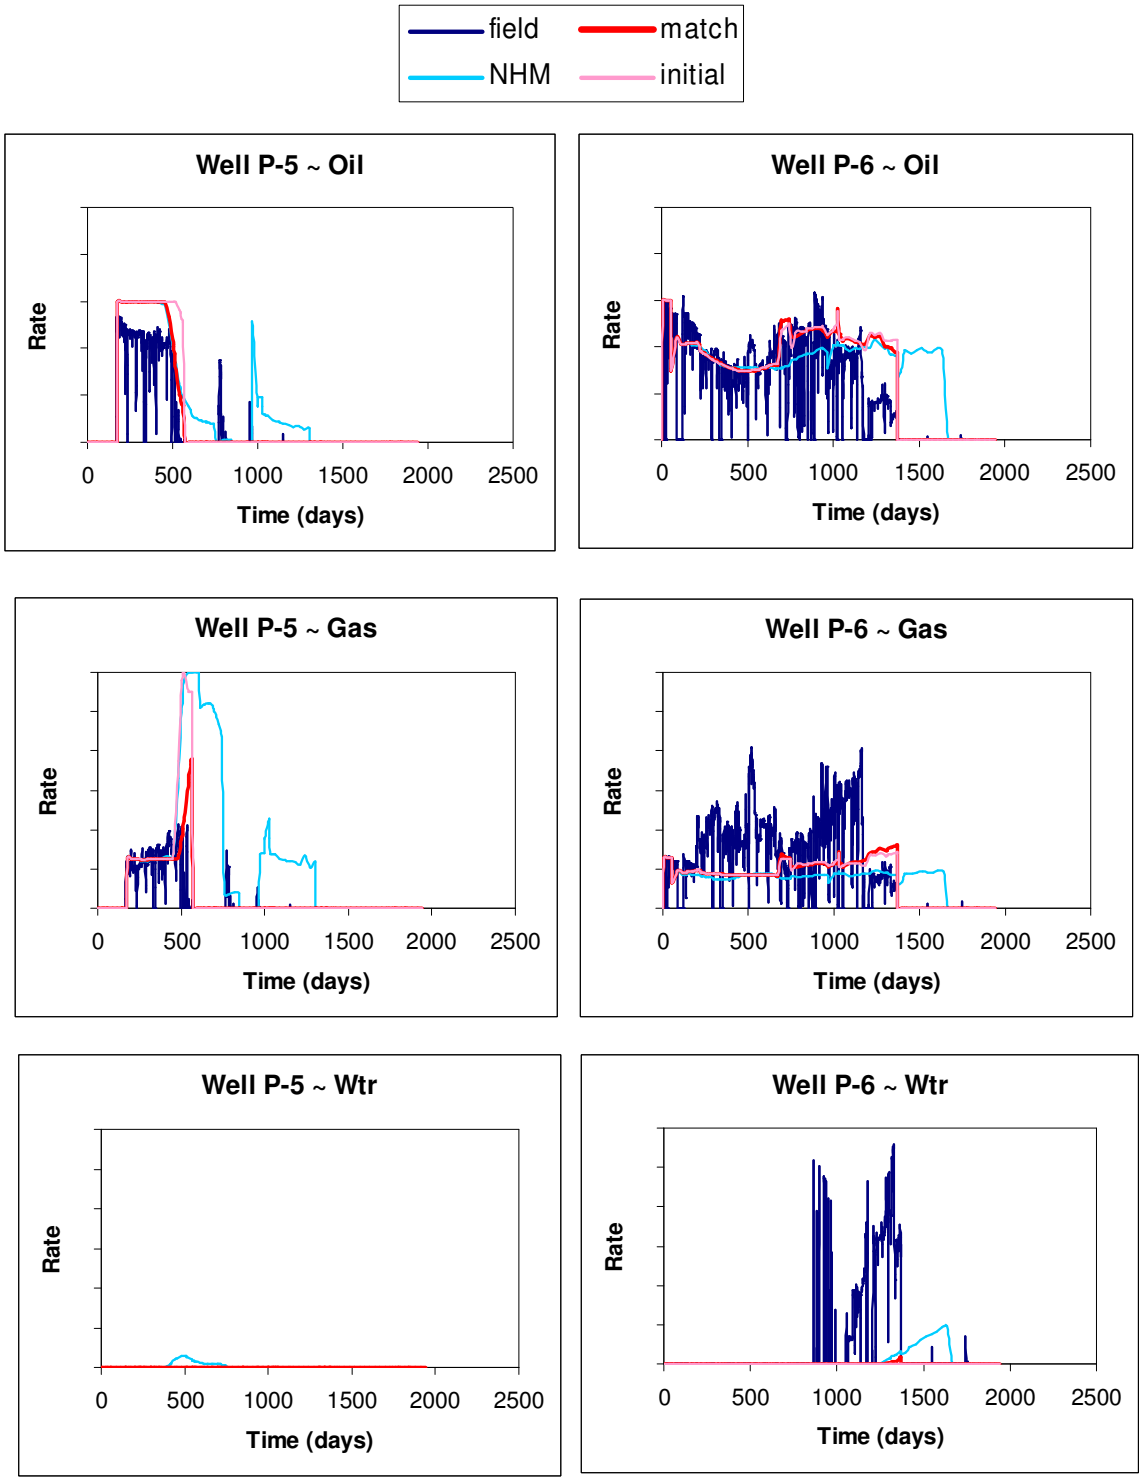
field
match
NHM
initial
Well P-5 ~ Oil
Well P-6 ~ Oil
Well P-5 ~ Gas
Well P-6 ~ Gas
Well P-5 ~ Wtr
Well P-6 ~ Wtr
Rate
Time (days)
0
500
1000
1500
2000
2500

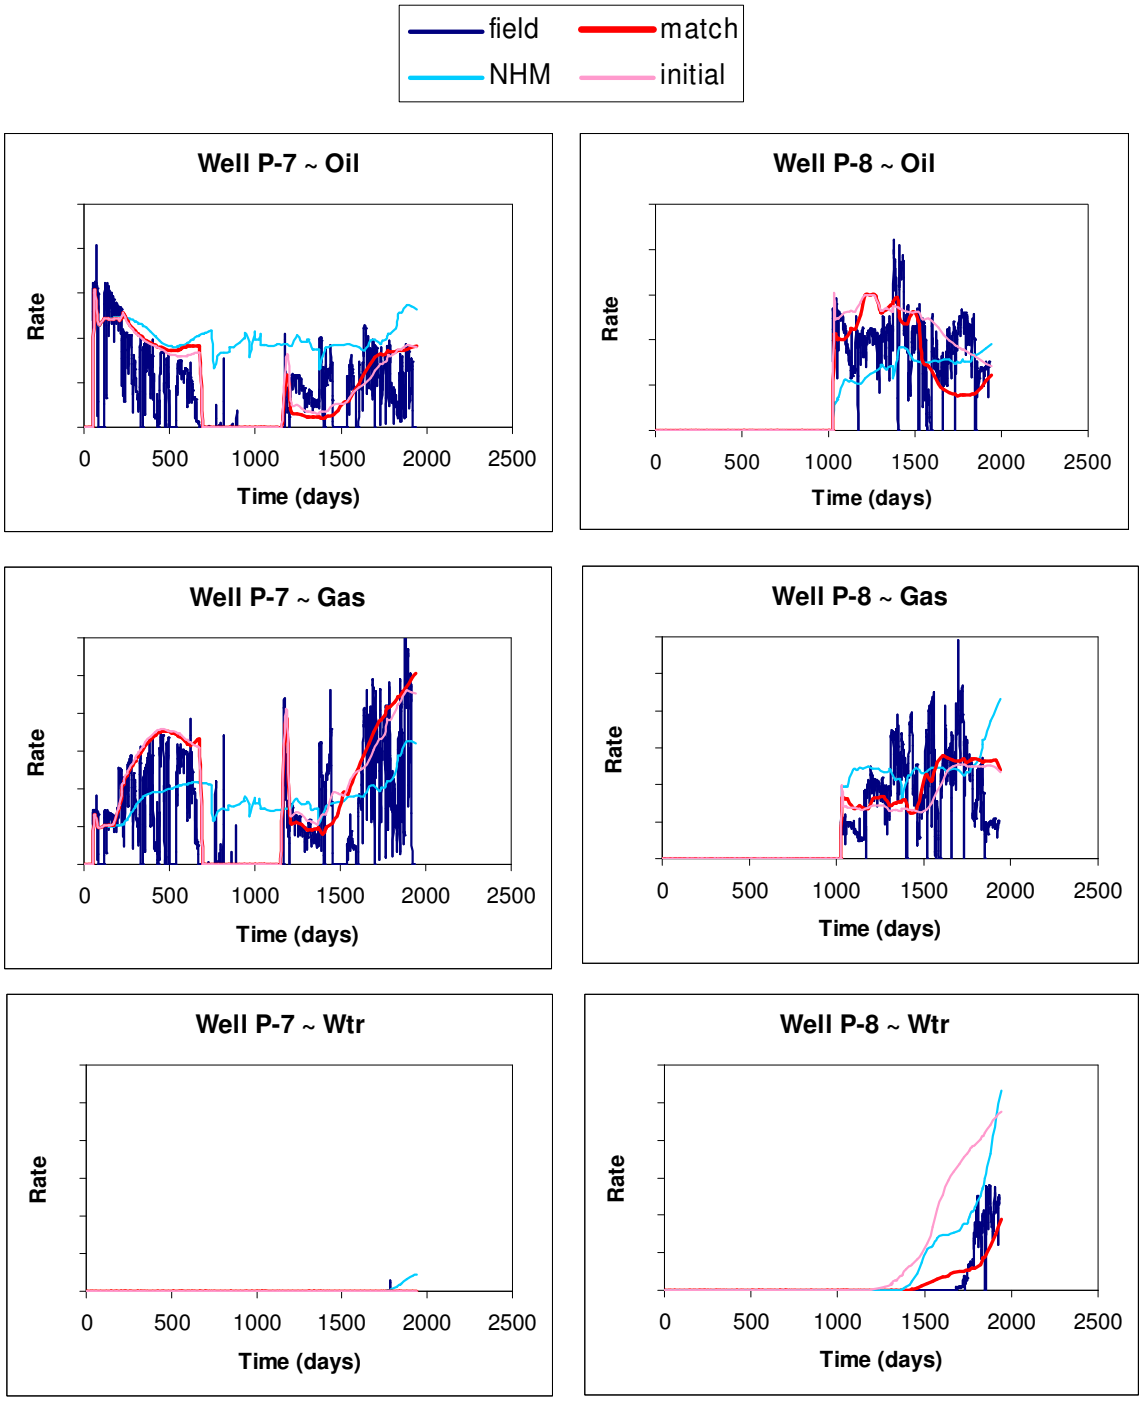
field
match
NHM
initial
Well P-7 ~ Oil
Well P-8 ~ Oil
Well P-7 ~ Gas
Well P-8 ~ Gas
Well P-7 ~ Wtr
Well P-8 ~ Wtr
Rate
Time (days)
0
500
1000
1500
2000
2500

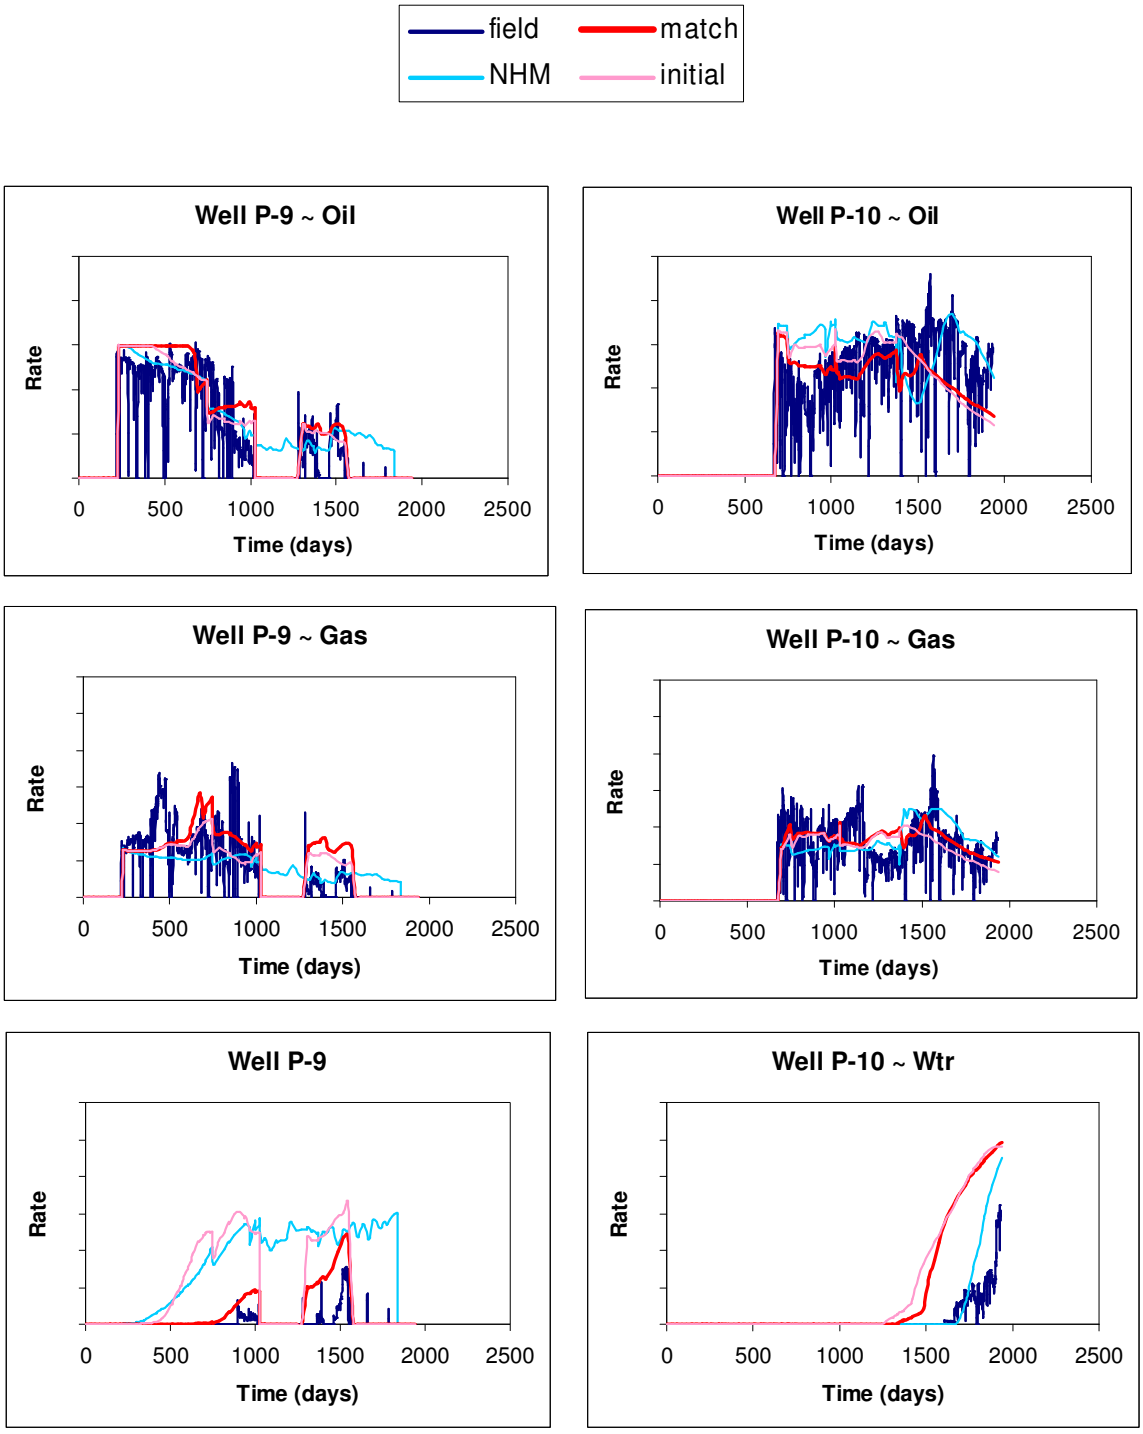
field
match
NHM
initial
Well P-9 ~ Oil
Well P-10 ~ Oil
Well P-9 ~ Gas
Well P-10 ~ Gas
Well P-9
Well P-10 ~ Wtr
Rate
Time (days)
0
500
1000
1500
2000
2500

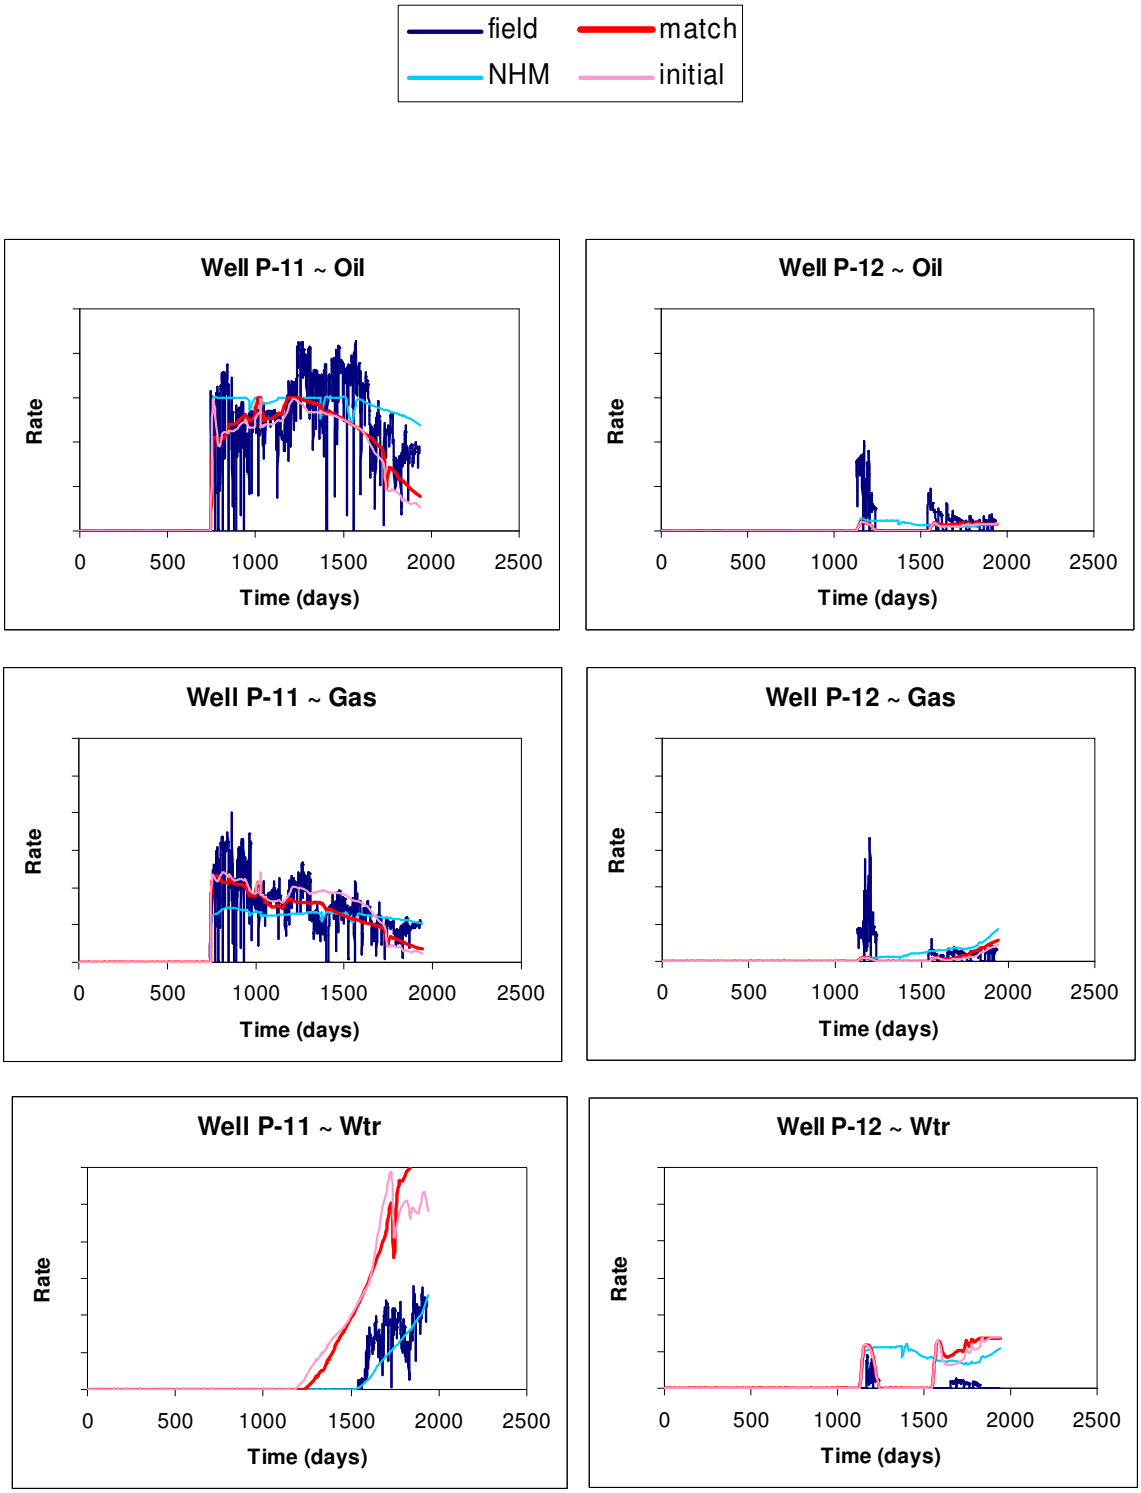
field
match
NHM
initial
Well P-11 ~ Oil
Well P-12 ~ Oil
Well P-11 ~ Gas
Well P-12 ~ Gas
Well P-11 ~ Wtr
Well P-12 ~ Wtr
Rate
Time (days)
0
500
1000
1500
2000
2500

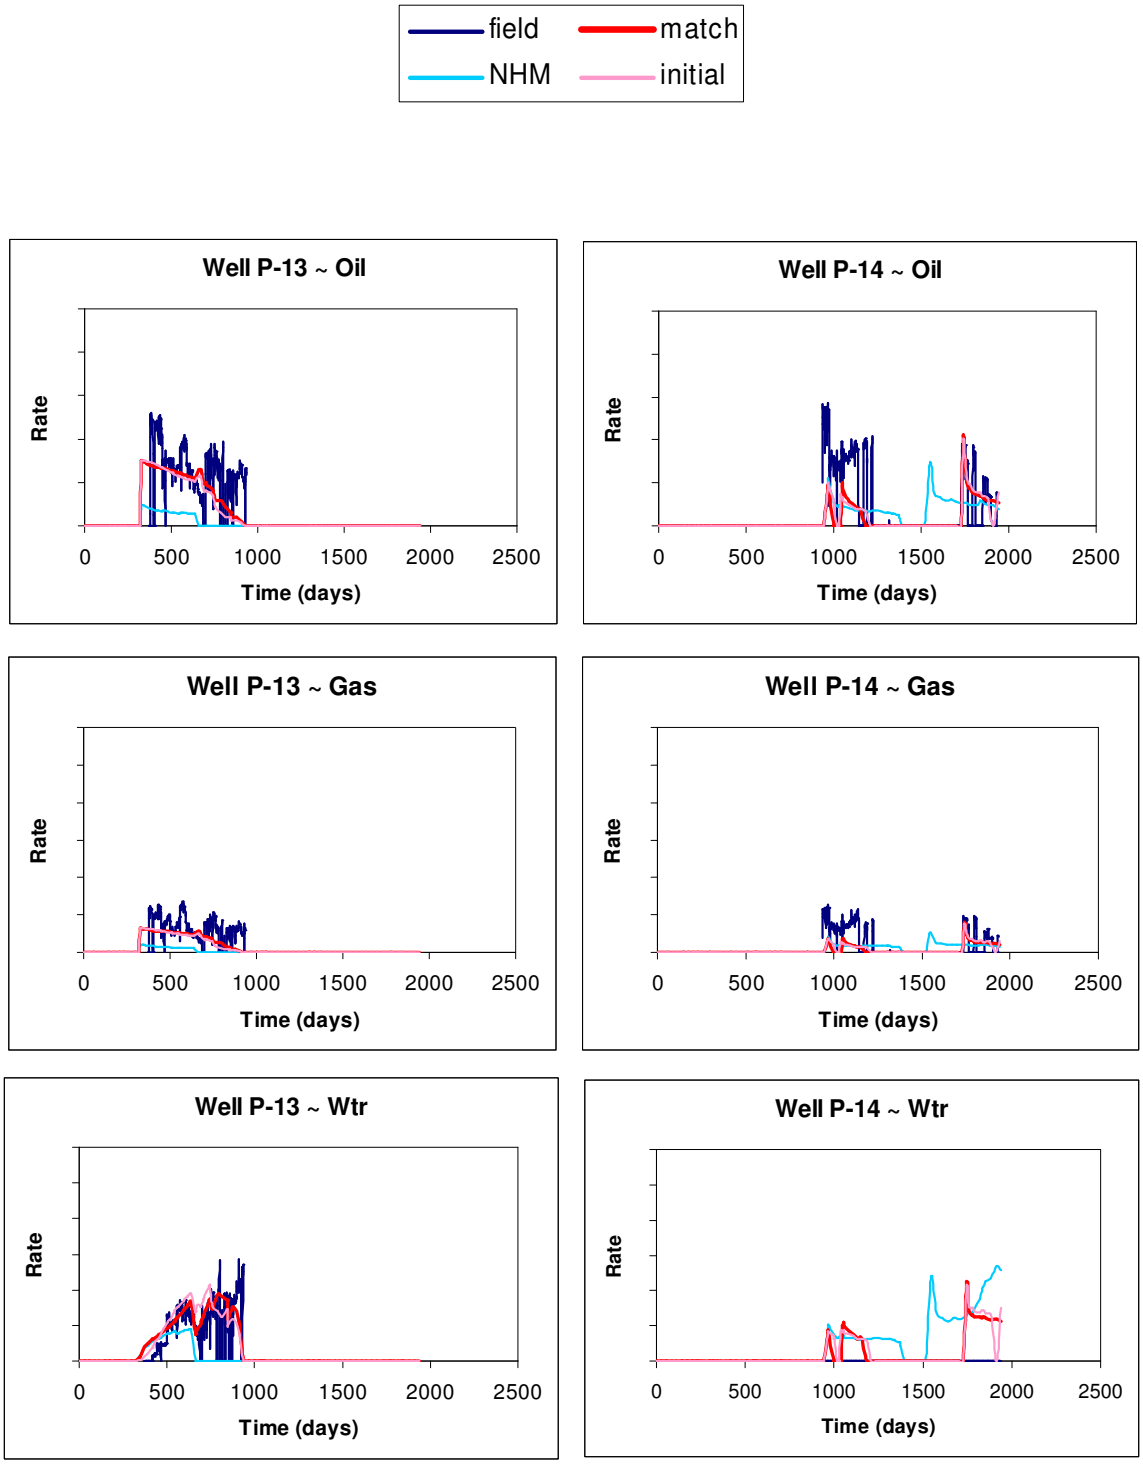
field
match
NHM
initial
Well P-13 ~ Oil
Well P-14 ~ Oil
Well P-13 ~ Gas
Well P-14 ~ Gas
Well P-13 ~ Wtr
Well P-14 ~ Wtr
Rate
Time (days)
0
500
1000
1500
2000
2500

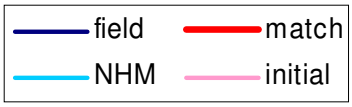
field
match
NHM
initial

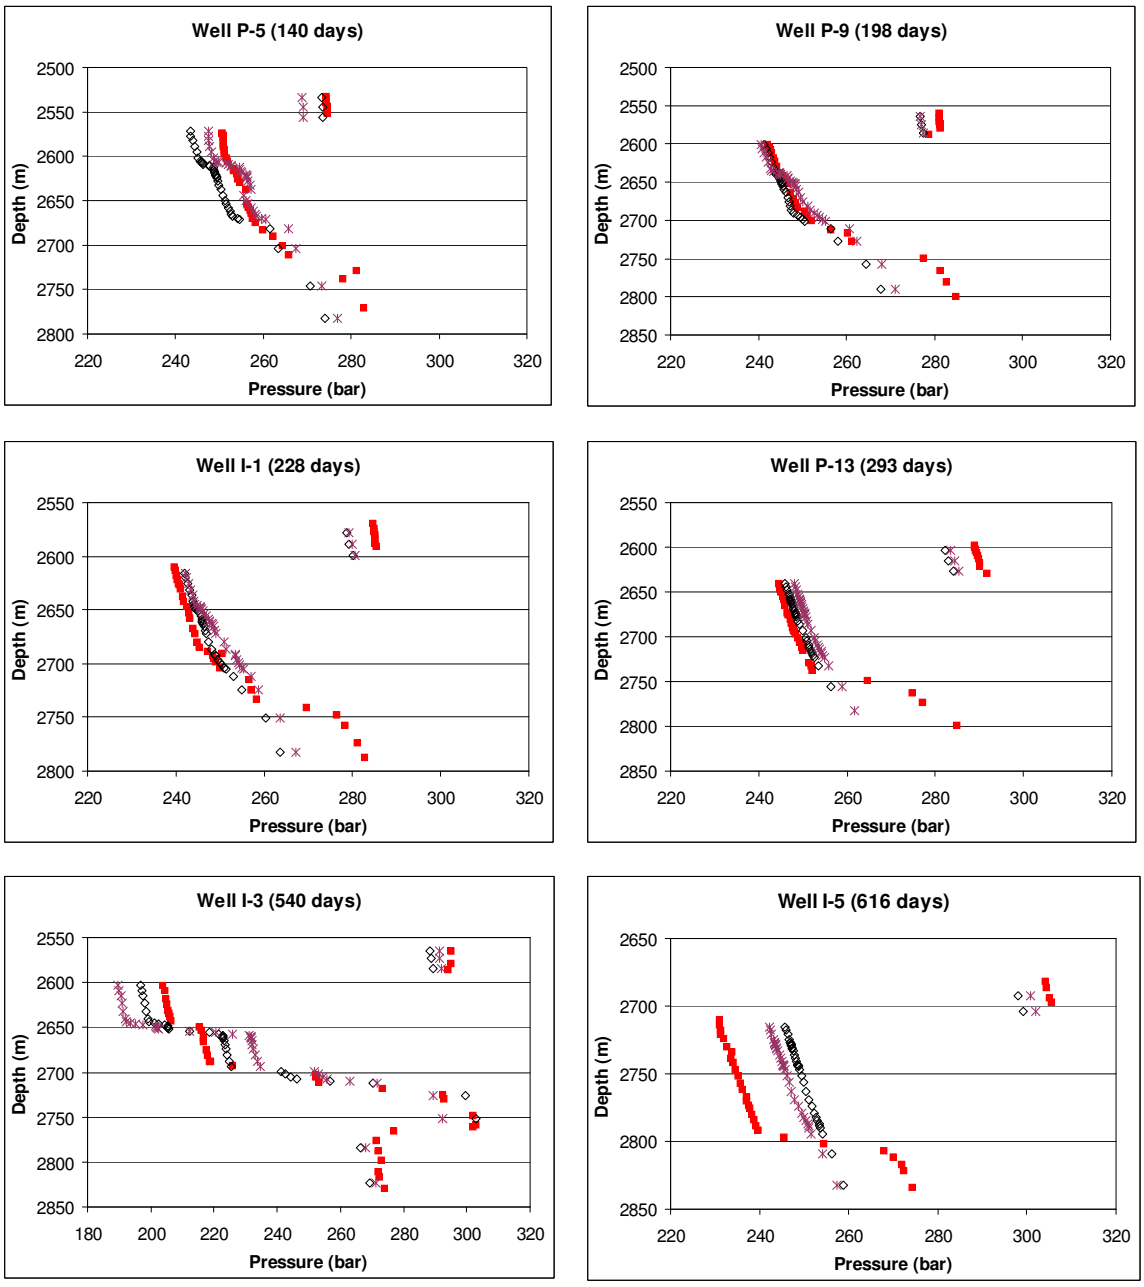
Well P-5 (140 days)
Depth (m)
Pressure (bar)
Well P-9 (198 days)
Depth (m)
Pressure (bar)
Well I-1 (228 days)
Depth (m)
Pressure (bar)
Well P-13 (293 days)
Depth (m)
Pressure (bar)
Well I-3 (540 days)
Depth (m)
Pressure (bar)
Well I-5 (616 days)
Depth (m)
Pressure (bar)
Field
NHM Model
Matched Model

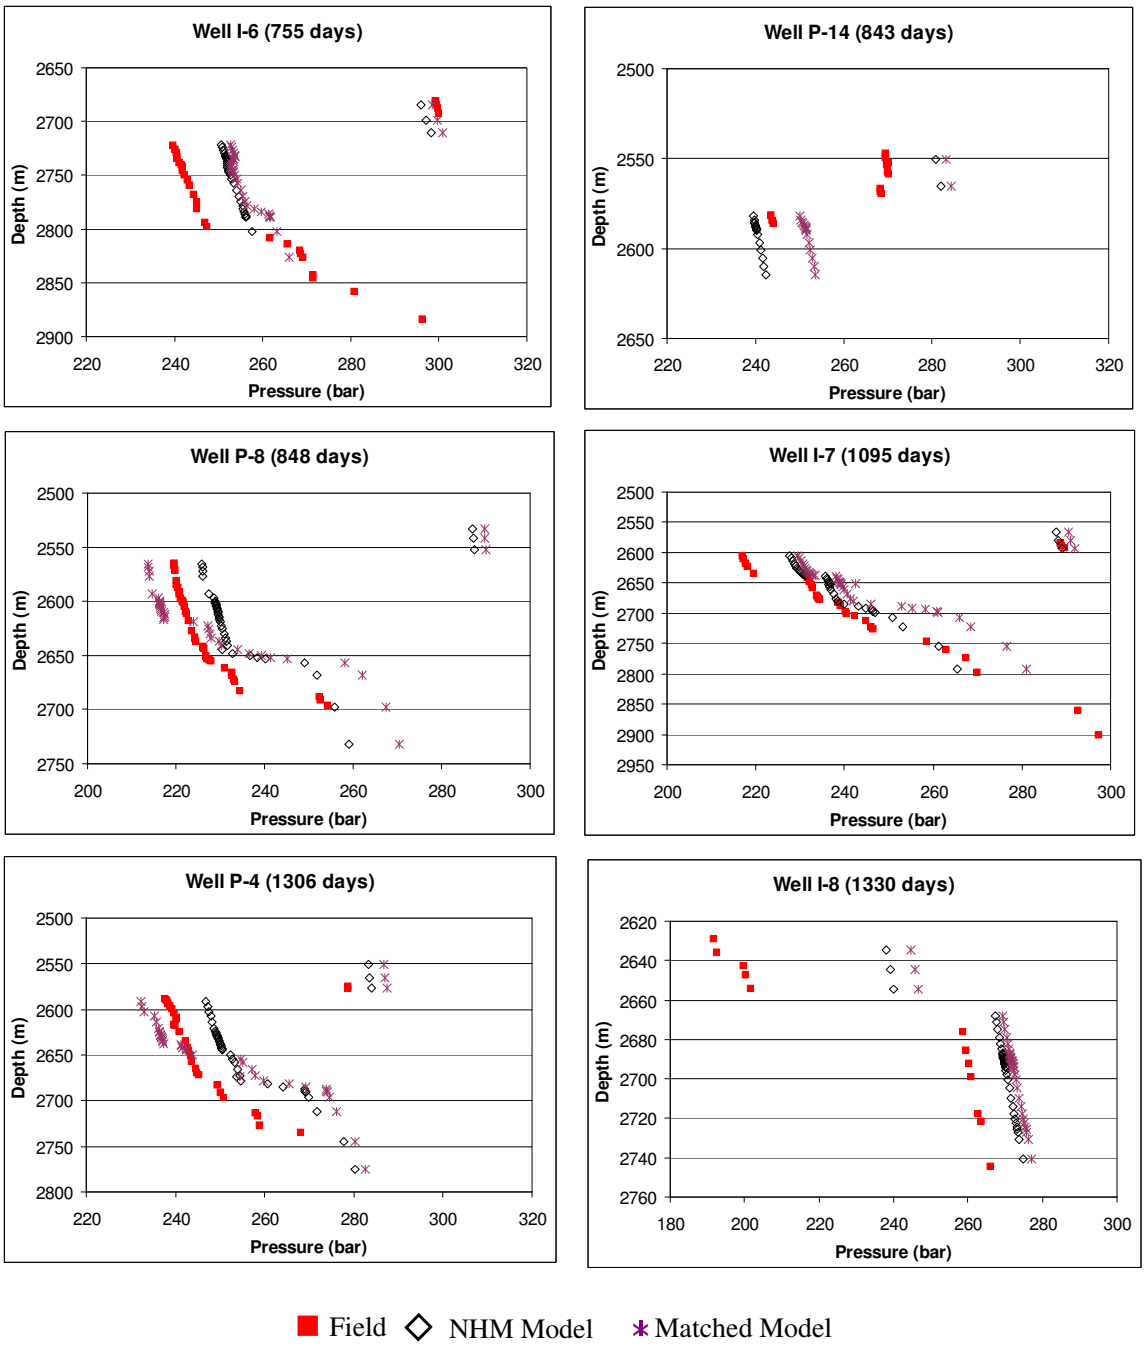
Well I-6 (755 days)
Depth (m)
Pressure (bar)
Well P-14 (843 days)
Depth (m)
Pressure (bar)
Well P-8 (848 days)
Depth (m)
Pressure (bar)
Well I-7 (1095 days)
Depth (m)
Pressure (bar)
Well P-4 (1306 days)
Depth (m)
Pressure (bar)
Well I-8 (1330 days)
Depth (m)
Pressure (bar)
Field
NHM Model
Matched Model

"Knowledge, in the end, is based on acknowledgement".

-Ludwig Wittgenstein

Bibliography

Aarts, E. H. L. and Korst, J.: Simulated annealing and Boltzmann machines, New York, Wiley, 1989.

Anterion, F., Eymard, R., and Karcher, B.: "Use of Parameter Gradients for Reservoir History Matching," paper SPE 18433 presented at the SPE Symposium on Reservoir Simulation, Houston, TX, February, 6-8, 1989.

Agarwal, B. and Blunt, M. J.: "Full-Physics, Streamline-Based Method for History Matching Performance Data of a North Sea Field," paper SPE 66388 presented at the SPE Reservoir Simulation Syposium, Houston, TX, February, 11-14, 2001.

Bissell, R., Killough, J. E., and Sharma, Y.: "Reservoir History Matching Using the Method of Gradients on a Workstation," paper SPE 24265 presented at the SPE European Petroleum Computer Conference, Stavanger, Norway, May, 25-27, 1992.

Bissell, R., Sharma, Y., and Killough, J. E.: "History Matching Using the Method of Gradients: Two Case Studies," paper SPE 28590 presented at the SPE Annual Technical Conference and Exhibition, New Orleans, LA, September, 25-28, 1994.

Caers, J., Krishnan, S., Wang, Y. and Kovscek, A. R.: "A Geostatistical Approach to Streamline-Based HistoryMatching," *Soc. Pet. Eng. Journal* pp. 250-266, September 2002.

Caers, J., "Data conditioning with the probability perturbation method," In: Geostatistics 2004 Banff, Eds. Leuangthong and Deutsch, Canada, 2004.

Caers, J., and Hoffman, B.T.: "The probability perturbation method: A new look at Bayesian inverse modeling," *Mathematical Geology*, accepted for publication, 2005.

Carter, R. D., Pierce, A. C., Kemp, L., and Willians, D. L.: "Performance Matching With Constraints," *Soc. Pet. Eng. Journal* pp. 187-196, April 1974.

Carter J. N., and Ballester P. J.: "A Real Parameter Genetic Algorithm for Cluster Identification in History Matching," Proceeding of the European Conference on the Mathematics of Oil Recovery, Cannes, France, August 30 – September 2, 2004.

Caumon, G., Strebelle, S., Caers, J. K., and Journel, A. G.: "Assessment of Global Uncertainty for Early Appraisal of Hydrocarbon Fields," paper SPE 89943 presented at the SPE Annual Technical Conference and Exhibition, Houston, TX, September, 26-29, 2004.

Chavent, G., Dupuy, M., and Lemonnier, P.: "History Matching by Use of Optimal Theory," paper SPE 4627 presented at the SPE-AIME 48th Annual Fall Meeting, Las Vegas, NV, September 30-October 3, 1973.

Chen, W. H., Gavalas, G. R., Seinfeld, J. H., and Wasserman, M. L.: "A New Algorithm for Automatic History Matching," paper SPE 4545 presented at the SPE-AIME 48th Annual Fall Meeting, Las Vegas, NV, September 30-October 3, 1973.

Coats, K. H., Dempsey, J. R., and Henderson, J. H.: "A New Technique for Determining Reservoir Description from Field Performance Data," *Soc. Pet. Eng. Journal*, pp. 66-74 March, 1970.

Damsleth, E., Hage, A., and Volden, R., "Maximum Information at Minimum Cost: A North Sea Field Development Study with Experimental Design", *Journal of Petroleum Technology*, pp. 1350-1360, December, 1992.

Datta-Gupta, A., Vasco, D. W. and Long, J. C. S.: "Detailed Characterization of Fractured Limestone Formation Using Stochastic Inverse Approaches," paper SPE 27744 presented at the SPE/DOE Ninth Symposium on Improved Oil Recovery, Tulsa, OK, April 17-20, 1994.

Datta-Gupta, A., Lake, L. A., and Pope, G. A.: "Characterizing Heterogeneous Permeable Media with Spatial Statistics and Tracer Data using Sequential Simulated Annealing," *Mathematical Geology*, Vol. 27 Iss. 6, pp. 763-787, 1995.

Delhomme, J. P.: "Spatial variability and Uncertainty in Groundwater Flow parameters: A geostatistical Approach," *Water Resources Research*, Vol. 15 Iss. 5, pp. 269-280, 1979.

Deutsch, C. V., and Journel, A. G.: GSLIB: Geostatistical Software Library and User's Guide, Second Edition, New York, Oxford University Press, pp. 300, 1998.

Deutsch, C. V., and Wang, L.: "Hierarchical Object-Based Geostatistical Modeling of Fluvial Reservoirs," paper SPE 36514 presented at Annual Fall Meeting of SPE in Denver, CO, October 6-9, 1996.

Earth Decision Sciences: GOCAD 2.0: User's Manual, Houston, TX, pp. 1527, March, 2003.

Fasanino, G., Molinard, J., and de Marsily, G.: "Inverse Modeling in Gas Reservoirs," paper SPE 15592 presented at the SPE 61st Annual Technical Conference and Exhibition, New Orleans, LA, October, 5-8, 1986.

Feraille M., Roggero, F., Manceau, E., Hu, L.Y., Zabalza-Mezghani, I. and Costa Reis, L.: "Application of Advanced History Matching Techniques to an Integrated Field Case

Study," paper SPE 84463 presented 78th Annual Fall Meeting of SPE in Denver, CO, October 5-8, 2003.

Gavalas, G.R., Shah, P.C., Seinfeld, J.H.: "Reservoir History Matching by Bayesian Estimation," *Soc. Pet. Eng. Journal*, pp 337-350, Dec., 1976.

Geman S., and Geman, D.: "Stochastic relaxation, Gibb distribution and the Bayesian restoration of images," IEEE Transactions on Pattern Analysis and Machine Learning, pp. 721-741, 1984.

Gill, P. E., Murray, W., and Wright, M. H.: Practical Optimization, San Diego, CA, Academic Press, 1981.

Goldberg, D. E.: Genetic Algorithms in Search, Optimization, and Machine Learning, Reading, MA, Addison-Wesley Pub., 1989.

Gomez-Hernandez, J.J. and Journel, A.G. Conditional simulation of the Wilmington sand-shale sequence, Los Angeles basin, In Hydraulic Engineering. ed., M.A. Ports, publ. American Society of Civil Engineers, NY, pp. 801-806, 1989.

Gross, H., Thiele, M. R., Alexa, M. J., Caers, J. K., and Kovscek, A. R.: "Streamline-Based History Matching Using Geostatistical Constraints: Application to a Giant Mature Carbonate Reservoir," SPE 90069, presented at 79th SPE Annual Technical Conference and Exhibition, Houston, TX, September 26-29, 2004.

Haas, A. and Viallix, J. R.: "Krigeage applied to geophysics," *Geophysical Prospecting*, v. 24, pp. 49-69, 1976.

Hegstad, B.K. and Omre, H. "Uncertainty assessment in history matching and forecasting," In: Geostatistics 1996 Wollongong, Eds. Baafi and Schofield, Australia, Vol.1, p. 585-596, 1996.

Henriquez, A., Tyler, K. J., and Hurst, A.: "Characterization of Fluvial Sedimentology for Reservoir Simulation Modeling," *Soc. Pet. Eng. Journal FE*, September, 1990.

Hoffman, B. T. and Caers, J., "History Matching Under Geologic Control: Application to a North Sea Reservoir," In: Geostatistics 2004 Banff, Eds. Leuangthong and Deutsch, Canada, 2004.

Hoffman, B. T. and Caers, J., "Regional Probability Perturbation for History Matching," *Journal of Petroleum Science and Engineering*, Vol. 46, pp. 53-71, January, 2005.

Hirasaki G. J.: "Sensitivity Coefficients for History Matching Oil Displacement Processes," *Soc. Pet. Eng. Journal*, pp. 39-49, February, 1975.

Jacquard, P. and Jain, C.: "Permeability Distribution from Field Pressure Data," *Soc. Pet. Eng. Journal*, pp. 281-294, December, 1965.

Journel, A. G.: "Combining Knowledge from Diverse Sources: an Alternative to Traditional Data Independence Hypothesis," *Mathematical Geology*, Vol. 34, Iss. 5, July 2002.

Journel A.G. and Alabert, F.A.: "Focusing on spatial connectivity of extreme-valued attributes: Stochastic indicator models of reservoir heterogeneities," paper SPE 18324 presented at the SPE Annual Technical Conference and Exhibition, Houston, TX, October, 2-5, 1988.

Journel, A.G. and Isaaks, E.: "Conditional indicator simulation: Application to a Saskatchewan uranium deposit," *Mathematical Geology*, Vol. 16, Iss.. 7, pp. 685-718.

Journel, A.J., and Huijbregts, C.: Mining Geostatistics. Academic Press, London, 600 p., 1978.

Kim, J. and Caers, J., "History Matching by Joint Perturbation of Facies Distribution and Net-to-Gross Ratio (N/G)," Stanford Center for Reservoir Forecasting (SCRF) Report No. 16, May 2003.

Krishnan, S.: The Tau Model to Integrate Prior Probabilities, Ph.D. Dissertation, Stanford University, December 2004.

Krige, D. G.: "A statistical Approach to Some Mine Valuations and Allied Problems at the Witwatersrand," Master's Thesis, University of Witwatersrand, 1951.

Landa, J. L.: "Technique to Integrate Production and Static Data in a Self-Consistent Way," paper SPE 71597 presented 76th Annual Fall Meeting of SPE in New Orleans, LA, September 30 - October 3, 2001.

Landa, J. L and Guyaguler B.: "A Methodology for History Matching and the Assessment of Uncertainties Associated with Flow Prediction," paper SPE 84465 presented 78th Annual Fall Meeting of SPE in Denver, CO, October 5-8, 2003.

Landa, J. L., Kamal, M. M., Jenkins, C. D., and Horne, R. N.: "Reservoir Characterization Constrained to Well Test Data: A Field Example," paper SPE 36511 presented at the SPE Annual Technical Conference and Exhibition, Denver, CO, October, 6-9, 1996.

Le Revalec-Dupin M., and Hu, L.Y.: "Gradual Deformation of Boolean Simulations," In: Geostatistics 2004 Banff, Eds. Leuangthong and Duetsch, Canada, 2004

Li R., Reynolds, A. C., and Oliver, D. S.: "History Matching of Three-Phase Flow Production Data," *Soc. Pet. Eng. Journal*, December, 2003.

Lodoen, O. P. and Omre, H.: "Improved Production Forecasts and History Matching Using Approximate Fluid-Flow Simulators," *Soc. Pet. Eng. Journal*, pp. 339-351, September, 2004.

Matheron, G.: Traité de géostatistique appliqué, Paris, France, Editions Technip, 1963.

Mattox, C. C., and Dalton, R. L.: Reservoir Simulation, Richardson, TX, SPE Monograph Volume 13, 1990

Mao, S.: "Generation of a reference petrophysical/seismic data set: the Stanford V reservoir," Stanford Center for Reservoir Forecasting (SCRF) Report No. 12, 1999.

Milliken, W. J., Emanuel, A. S., and Chakravarty, A.: "Applications of 3D Streamline Simulation To Assist History Matching," *SPE Journal RE&E*, pp. 502-508, December 2001.

Moosegard, K. and Tarantola, A., "Monte Carlo sampling of solutions to inverse problems," *J. Geophys. Res.*, B, v. 100, pp. 12431-12447, 1995.

Oliver, D. S.: "Incorporation of Transient Pressure Data into Reservoir Characterization," *In Situ*, Vol. 18, No. 3, pp. 243-275, 1994.

Omre, H.: "Stochastic reservoir Models Conditioned to Non-Linear Production History Observation," In: Geostatistics 2000 Cape Town, Eds. Kleingeld and Krige, Cape Town, South Africa, Vol. 1, pp. 166–175, 2000.

Ouenes, A., and Sasd, N.: "A New, Fast Parallel Simulated Annealing Algorithm for Reservoir Characterization," paper SPE 26419 presented at the SPE Annual Technical Conference and Exhibition, Houston, TX, October, 3-6, 1993.

Phan, V., and Horne, R. L.: "Fluvial Channel Parameter Estimation Constrained to Static Production, and 4D Seismic Data," paper SPE 77518 presented at the SPE Annual Technical Conference and Exhibition, San Antonio, TX, September 29 - October 2, 2002.

Qassab, H., Khalifa, M., Pavlas, R., Afaleg, N., Ali, H., Kharghoria, A., He, Z., Lee, S. H. and Datta-Gupta, A.: "Streamline-based Production Data Integration under Realistic Field Conditions: Experience in a Giant Middle-Eastern Reservoir," paper SPE 84079 presented 78th Annual Fall Meeting of SPE in Denver, CO, October 5-8, 2003.

Roggero, F., and Hu, L. Y.: "Gradual Deformation of Continuous Geostatistical Models for History Matching," paper SPE 49004 presented at 73rd Annual Fall Meeting of SPE in New Orleans, LA, September 27 - 30, 1998.

Rojas, D.: "Determining likelihood of rock failure (fracturing) using an elastic simulator" Stanford Center for Reservoir Forecasting (SCRF) Report No. 18, May 5-6, 2005.

Saleri, N. G., Toronyi, R. M., and Shyder, D. E., "Data and Data Hierarchy," *Journal of Petroleum Technology,* SPE 21369, pp. 1286-1293, December, 1992.

Schlumberger: Simulation Software Manuals, Eclipse Technical Description, 2003.

Sen, M. K., Datta-Gupta, A., Stoffa, P. L., Lake, L. W., and Pope, G. A.: "Stochastic Reservoir Modeling Using Simulated Annealing and Genetic Algorithms," paper SPE 24754 presented at the SPE Annual Technical Conference and Exhibition, Washington, DC, October, 4-7, 1992.

StreamSim Technologies: 3DSL Streamline Simulator User's Manual, Version 1.8, San Francisco, CA, pp. 104, January, 2002.

Strebelle, S.: "Conditioning Simulation of Complex Structures Multiple-Point Statistics" *Mathematical Geology*, v. 34, Iss.. 1, January, 2002.

Sultan, A. J., Ouenes, A., and Weiss, W. W.: "Automatic History Matching for an Integrated Reservoir Description and Improving Oil Recovery," paper SPE 27712 presented at the SPE Permian Basin Oil and Gas Recovery Conference, Midland, TX, March, 16-18, 1994.

Suzuki, S., Daly, C., Caers, J., and Mueller, D.: "History Matching of Naturally Fractured Reservoir Using Elastic Stress Simulation, FracPerm Model and the Probability Perturbation Method," Stanford Center for Reservoir Forecasting (SCRF) Report No. 18, May 5-6, 2005.

Suzuki, S.: "Effect of Structure Geometry on History Matching of Faulted Reservoirs" Stanford Center for Reservoir Forecasting (SCRF) Report No. 17, May 7-8, 2004.

Suzuki, S.: "Determining Petrophysical Properties of Facies Using a Hierarchical History Matching Method" Stanford Center for Reservoir Forecasting (SCRF) Report No. 16, May 8-9, 2003.

Tan, T. B. and Kalogerakis, N.: "A Fully Implicit, Three-Dimensional, Three-Phase Simulator with Automatic History-Matching Capability," paper SPE 21205 presented at SPE 11th Symposium on Reservoir Simulation, Anaheim, CA, February, 17-20, 1991.

Tarantola, A.: Inverse Problem Theory: Methods for Data Fitting and Model Parameter Estimation, Elsevier, Amsterdam, The Netherlands, 1987.

Tjelmeland, H., and Omre, H.: "A Complex Sand-Shale Facies Model Conditioned on Observations from Wells, Seismic and Production," In: Geostatistics 1996 Wollongong, Eds. Baafi and Schofield, Austrailia, vol. 1, pp. 634-647, 1996.

Tran, T.: "The 'missing scale' and direct simulation of block effective properties," *Journal of Hydrology*, Vol. 183, pp. 37-56, August, 1996.

Tureyen, O. I.: A Parallel Modeling Approach to Reservoir Characterization, Ph.D. Dissertation, Stanford University, March 2005.

Veatch Jr., R. W., and Thomas, G. W.: "A Direct Approach for History Matching," paper SPE 3515 presented Annual Fall Meeting of SPE in New Orleans, LA, Oct. 3 – Oct. 6, 1971.

Voelker, J., "A Geostatistical Method for Characterizing Superpermeability From Flow-Meter Data: Application to Ghawar Field," paper SPE 84279 presented 78th Annual Fall Meeting of SPE in Denver, CO, Oct. 5 – Oct. 8, 2003.

Wahl, W. L., Mullins, L., Barham, R., and Bartlet, W.: "Matching the Performance of Saudi Arabian Oil Fields With an Electrical Model," paper SPE 414 presented at the SPE 37th Annual Fall Meeting, Los Angeles, CA, October, 7-10, 1962.

Wang, Y. and Kovscek A. R.: "Streamline Approach for History-Matching Production Data," *Soc. Pet. Eng. Journal*, pp. 353-362, December 2000.

Wang, Y. and Kovscek A. R., "Integrating Production History into Reservoir Models Using Streamlin-Based Time-of-Flight Ranking," *Petroleum Geoscience*, Vol. 9 Iss. 2 pp. 163-174, January, 2003.

Wen, X. H., Deutsch, C.V., and Cullick, A. S.: "High-Resolution Reservoir Models Integrating Multiple-Well Production Data," *Soc. Pet. Eng. Journal*, pp. 344-355, December, 1998.

White, C. D., Novakovic, D., Dutton, S. P., and Willis, B. J.: "A Geostatistical Model for Calcite Concretions in Sandstone," *Mathematical Geology*, Vol. 35, Iss. 5, pp. 549-575, July, 2003.

White, C. D., Willis, B. J., Narayanan, K., and Dutton, S. P.: "Identifying and Estimating Significant Geologic Parameters with Experimental Design," *Soc. Pet. Eng. Journal*, pp. 311-324, September, 2001.

Williams M. A., and Keating J. F.: "The Stratigraphic Method: A Structured Approach to History Matching Complex Simulation Models," paper SPE 38014 presented at SPE Reservoir Simulation Symposium, Dallas, TX, June, 8-11, 1997.

Yamada, T.: "Non-Uniqueness of History Matching," paper SPE 38014 presented at SPE Asia Pacific Conference, Yokohama, Japan, April, 25-26, 2000.

Success

To laugh often and much;
to win the respect of intelligent people
and the affection of children;
to earn the appreciation of honest critics
and endure the betrayal of false friends;
to appreciate beauty; to find the best in others;
to leave the world a bit better,
whether by a healthy child, a garden patch
or a redeemed social condition;
to know even one life has breathed easier
because you have lived.
This is to have succeeded.

-Ralph Waldo Emerson