CONTINUOUS RESERVOIR SIMULATION INCORPORATING UNCERTAINTY QUANTIFICATION AND REAL-TIME DATA

A Thesis

by

JAY CUTHBERT HOLMES

Submitted to the Office of Graduate Studies of
Texas A&M University
in partial fulfillment of the requirements for the degree of

MASTER OF SCIENCE

December 2006

Major Subject: Petroleum Engineering

CONTINUOUS RESERVOIR SIMULATION INCORPORATING

UNCERTAINTY QUANTIFICATION AND REAL-TIME DATA

A Thesis

by

JAY CUTHBERT HOLMES

Submitted to the Office of Graduate Studies of
Texas A&M University
in partial fulfillment of the requirements for the degree of

MASTER OF SCIENCE

Approved by:

Chair of Committee,	Duane A. McVay
Committee Members,	W. John Lee
	Akhil Datta-Gupta
	J. Eric Bickel
Head of Department,	Stephen A. Holditch

December 2006

Major Subject: Petroleum Engineering

ABSTRACT

Continuous Reservoir Simulation Incorporating

Uncertainty Quantification and Real-time Data. (December 2006)

Jay Cuthbert Holmes, B.A., Rice University

Chair of Advisory Committee: Dr. Duane A. McVay

A significant body of work has demonstrated both the promise and difficulty of quantifying uncertainty in reservoir simulation forecasts. It is generally accepted that accurate and complete quantification of uncertainty should lead to better decision making and greater profitability. Many of the techniques presented in past work attempt to quantify uncertainty without sampling the full parameter space, saving on the number of simulation runs, but inherently limiting and biasing the uncertainty quantification in the resulting forecasts. In addition, past work generally has looked at uncertainty in synthetic models and does not address the practical issues of quantifying uncertainty in an actual field. Both of these issues must be addressed in order to rigorously quantify uncertainty in practice.

In this study a new approach to reservoir simulation is taken whereby the traditional one-time simulation study is replaced with a new continuous process potentially spanning the life of the reservoir. In this process, reservoir models are generated and run 24 hours a day, seven days a week, allowing many more runs than previously possible and yielding a more thorough exploration of possible reservoir descriptions. In turn, more runs enabled better estimates of uncertainty in resulting forecasts. A new technology to allow this process to run continuously with little human interaction is real-time production and pressure data, which can be automatically integrated into runs.

Two tests of this continuous simulation process were conducted. The first test was conducted on the Production with Uncertainty Quantification (PUNQ) synthetic reservoir. Comparison of our results with previous studies shows that the continuous approach gives consistent and reasonable estimates of uncertainty. The second study was conducted in real time on a live field. This study demonstrates the continuous simulation process and shows that it is feasible and practical for real world applications.

ACKNOWLEDGEMENTS

I would like to thank my parents for their constant support of my pursuit of knowledge.

TABLE OF CONTENTS

Page

LIST OF TABLES

LIST OF FIGURES

FIGURE Page

FIGURE Page

INTRODUCTION

Reservoir management is generally considered a continuous process that should span the entire life of a reservoir.[1-2] Furthermore, reservoir simulation, with its unique predictive capabilities, is widely regarded as a critical tool in modern reservoir management practice.[3] Reservoir simulation yields an assessment of reservoir properties and, when a forecast run is made, an assessment of future production. These assessments feed directly into the decision-making process. In his rules for decision making, Howard[4] establishes that it is necessary to assign probabilities to all possible outcomes of uncertain events. Therefore, making a good decision requires taking into account all possible outcomes and so it is necessary to quantify the uncertainty in forecasts. Conversely, if the uncertainty quantification in a forecast is incomplete, or nonexistent, then the decision may be poor. For this reason it is necessary to rigorously quantify uncertainty in production forecasts.

Capen[5] demonstrated thirty years ago that people in the petroleum industry significantly underestimate uncertainty in their assessments. In keeping with this tendency, reservoir simulation engineers traditionally take only limited consideration of uncertainty and often times do not try to quantify it at all. Quantifying uncertainty in production forecasts, of course, is not a trivial undertaking. The reservoir parameter space, the set of all possible combinations of reservoir parameters, is literally infinite.

This thesis follows the style of *SPE Reservoir Evaluation & Engineering.*

Attempts at uncertainty quantification in more recent studies, specifically Floris *et al*.,[6] have shown that, even when we explicitly try to quantify uncertainty in simulation studies, we still tend to underestimate it. It is therefore worthwhile to explore reservoir simulation techniques aimed at better quantifying uncertainty in forecasts.

Typically, reservoir simulation is only utilized at discrete points in the life of a reservoir. Reservoir studies are expensive and time-consuming due to the time and manpower required to tune and history match a simulation model. As such, *traditional* simulation studies usually can only be justified when considering a major investment. Taken individually, smaller reservoir management decisions typically do not warrant the expense of a simulation study and thus must proceed without simulation results. Inaccurate forecasts or no forecasts at all can lead to sub-optimal operations and significant economic consequences. Clearly, reservoir management would benefit if a calibrated simulation model was available at any time.

One way to address these issues is to treat simulation as a continuous process, similar to how simulation is employed in weather forecasting. In continuous simulation history match runs will be made twenty-four hours a day, seven days a week over the course of the reservoir's life. When new data, both static and dynamic, are available they will be added to the history matching process. The wide-spread utilization of real time data acquisition systems makes it practical to build such a system. A continuous simulation system will provide ready access to an up-to-date model for use in day-to-day reservoir management. As such, the costs of the study can be amortized over the life of the field. Perhaps more importantly, continuous simulation offers years with which to conduct a more exhaustive search of the reservoir parameter space. In turn, more thoroughly exploring the parameter space should result in better uncertainty quantification.

This thesis explores this idea of continuous reservoir simulation through the use of a continuous reservoir simulation software system. This system is described in detail below. The system was tested on two reservoirs. The first test was conducted on a synthetic model from the Production forecasting with UNcertainty Quantification (PUNQ) study.[6-8] The other test was performed on a live producing field. Results of the synthetic test show reasonable agreement with previous PUNQ work. The field test demonstrates that it is practical to apply this process to a real field.

BACKGROUND

Uncertainty Quantification Techniques

In the past decade, there has been a significant amount of work towards developing more rigorous uncertainty quantification techniques. Of particular interest is the work coming out of the PUNQ study. A joint effort of several industrial and academic partners, this study used multiple synthetic reservoirs to test numerous history matching and uncertainty quantification techniques.

To date, the PUNQ work is probably the most thorough treatment of uncertainty quantification in production forecasts. This study attempted a comprehensive survey of history matching/optimization techniques and uncertainty quantification methods. It is important here to distinguish between history matching and uncertainty quantification. History matching techniques take reservoir models and tune them so they match production and pressure data. In general, history matching works by generating and adjusting models in order to minimize an objective function. An objective function compares simulation results to observed data in order to quantitatively describe how well a simulation model represents an actual reservoir. Uncertainty quantification methods utilize the results of these simulation runs as the basis for a probabilistic forecast.

One group of history matching techniques investigated in the PUNQ study is gradient techniques. Gradient-based methods for optimization work by calculating sensitivity of the objective function to certain parameters. Using these sensitivities the reservoir model can be adjusted until one with a minimal objective value is found.

Gradient methods are attractive as they can be computationally efficient, but have the downside of being easily trapped in local minima. This shortcoming of easily getting stuck in local minima prevents gradient methods from fully characterizing a complex parameter space and thus they do not provide a good basis for uncertainty quantification.

A more successful technique studied in the PUNQ work and elsewhere[9-10] is the Genetic Algorithm (GA). GAs are a broad class of optimization algorithms with a variety of applications.[11] GAs are based loosely around the rules that govern genetics in nature. In a GA, "generations" of unique reservoir models are created by mixing parameter values of previously run models in a process known as "breeding." For each model in the generation a simulation run is made and an objective function value calculated. These models then serve as the basis for creating a subsequent generation. In addition, parameter values of new models are randomly changed in a process known as "mutation." In some implementations, some members of the previous generation "survive" and are included in the new generation. In addition, a generation can include new members that are generated using the same process used to create the first generation. In time this process will sample a significant portion of the reservoir parameter space and result in the creation of some very good history matched models.

Genetic Algorithms have the desirable property of being a powerful global optimization tool capable of very accurate history matches while concurrently thoroughly exploring the parameter space for accurate uncertainty quantification. Unlike gradient methods, Genetic Algorithms can cope with multiple local minima. In order to accomplish this Genetic Algorithms must make a large number of runs, which is computationally intensive. For some optimization applications this is viewed as a drawback. For our application, however, a large number of runs more thoroughly explores the parameter space and should allow for better uncertainty quantification.

In addition to traditional Genetic Algorithms, the PUNQ study looked at the Markov Chain Monte Carlo (MCMC) which is statistically more rigorous. This technique can be considered to be a type of GA. Here a model is initially created by breeding members of a parent generation. After that the model is run, an objective function value calculated and then the model is mutated. If the mutation improves the objective function value the model is kept, otherwise it reverts to its previous state. The model is then mutated and run again and again until an acceptable objective function value is reached. At this point the model is considered “matched” and saved for later use in generating probabilistic forecasts. Like other GAs, the MCMC method results in an excellent exploration of the parameter space as well as good individual history matched models.

In addition to looking at history matching techniques, the PUNQ study investigated methods for quantifying the uncertainty in forecasts utilizing a set of matched reservoir models. While the forecast uncertainty quantification technique cannot be entirely separated from the history match method, the techniques fall into two broad classes. In the first, class forecasts are created by taking a set of runs around a single optimum or a handful of local optima and using these to construct a forecast. This type of uncertainty quantification can be used with any history matching technique, so long as you can identify minima. The second group of uncertainty methods attempts to fully sample the parameter space to more fully describe the uncertainty in the forecasts. This is accomplished by running models that attempt to sample the full parameter space rather than models focused around minima. This group is limited to techniques that sample a wide range of the parameter space, namely GAs and MCMC.

Real-time Data and Ensemble Kalman Filter

Another important issue in reservoir management is the management of data.[12] The last decade or so has seen a dramatic increase in the use of real-time data acquisition technology. This technology has been quite valuable for monitoring and short-term optimization.[13] However, despite the large investment companies have made in real-time data acquisition, it is not being used to its full potential in full-field reservoir simulation. Barden[14] has demonstrated a semi-analytical full-field modeling application employing real-time data. Real-time data appears to be an under-utilized resource that will facilitate a continuous simulation environment.

An automatic history matching and uncertainty quantification method that could make use of this real-time data, and which has recently gained a lot of attention, is the Ensemble Kalman Filter (EnKF).[15-16] The EnKF is distinctly different from the methods investigated in the PUNQ study in that it can *continuously* integrate data and update models. The PUNQ study methods attempt to quantify uncertainty at a fixed point in the reservoir's life. There is a fixed set of observed data which is history matched against and a fixed forecast period. In contrast to this, the EnKF works by updating static and dynamic model parameters at each time step for which observed data are available, as explained below.

In an EnKF an ensemble of unique initial reservoir models is created. This ensemble of models is created so that the ensemble as whole represents the variability in the underlying reservoir parameters. For each model a simulation run is made. During this simulation run observed data are incrementally integrated into the model via an assimilation step. In this assimilation step reservoir properties, including static properties such as porosity and permeability, are modified so that the model matches the observed data. By constructing the initial ensemble to represent the variability in reservoir properties, the resulting set of production forecasts should in theory represent the uncertainty in future production. This technique is attractive as it is computationally very efficient. Unfortunately, the physically unrealistic practice of changing static properties in the assimilation step causes these properties to head towards extreme values. In addition, the individual members of the ensemble tend to converge to similar solutions. The EnKF is still a topic of active research for history matching and more work is needed.

Justification for Continuous Approach

The reservoir parameter space is usually extremely large, even with a coarse parameterization of the reservoir. Obviously, we cannot make a simulation run for every possible combination of reservoir parameters. Despite the vastness of the parameter space, the techniques presented in the PUNQ study attempt to quantify uncertainty with relatively few runs. Techniques like the gradient methods attempt to quantify uncertainty using just a few hundred runs. The GAs and the MCMC make more runs, ranging from one thousand to several thousand runs. Even with the GA and MCMC techniques, however, there are practical limitations because they are being applied in the context of one-time studies, where there are time and budget constraints. Indeed, this is a limitation of all one-time simulation studies, where only so many runs can be made in limited period of time. However, if it were possible to make many more runs, one could better explore the parameter space.

Even though techniques that make thousands of runs were examined, the PUNQ study offers little insight into practical implementation using realistically sized models. The PUNQ study used small simulation models that could be run quickly even on desktop computers. It is not uncommon, however, for real world simulation models to take hours or even days to run on powerful servers. One way to approach uncertainty quantification with large simulation models, and make as many runs as possible, is to treat history matching and uncertainty quantification as a continuous process. This continuous process will entail making history match runs continuously over the life of the reservoir. Even with large simulation models this offers the potential to make tens of thousands of simulation runs over the life of the reservoir. These thousands of runs should yield a more thorough exploration of the parameter space and better probabilistic forecasts.

Currently, the tools exist with which to build a system that continuously history matches a petroleum reservoir and generates probabilistic forecasts. The real-time data acquisition technology needed to build a system has already been widely implemented. Uncertainty quantification techniques have advanced to the point where they can be adapted for continuous simulation. Such a system promises to give better uncertainty estimates of future production through a more exhaustive search of possible reservoir combinations. The objectives of this study are to implement a continuous simulation process and to evaluate its practicality and effectiveness in generating probabilistic forecasts in producing oil and gas fields. The process will be evaluated on two reservoirs, the PUNQ synthetic reservoir and a live field.

CONTINUOUS SIMULATION PROCESS

Overview

Conducting simulation in the continuous manner described above requires the combination of several components. First, the reservoir must be analyzed in order to determine uncertain parameters and their associated uncertainty. Because we will be making many more simulation runs than traditional studies, we can consider many more parameters for our model. Next, a method of sampling the parameter space and generating reservoir models is needed. In turn this requires code to automatically run these simulations and read the results. An objective function is used to evaluate the ability of an individual model to reproduce the observed data. As new data are acquired from the field, they are added to the objective function calculation. Finally, the results of individual runs are combined into probabilistic forecasts. Below, I will describe the implementation of each of these elements in this study.

Parameter Space Exploration

Before we can begin making simulation runs it is necessary to first evaluate the underlying reservoir and determine which uncertain parameters will be considered. In general this is a manual process and relies on the ability of the reservoir engineer to make assessments based on the available data. In the tests conducted here the parameters considered are porosity and permeability. After we identify the parameters of interest, we assign distributions (either discrete or continuous) to quantitatively represent the uncertainty in these parameters. By identifying uncertain parameters and assigning distributions to model their uncertainty we define the parameter space. This process of identify uncertain parameters and assessing distributions is fairly consistent with what is traditionally done when assessing input uncertainty in a simulation study. One key difference here, though, is that it is not necessary to severely reduce the number of parameters in order to expedite the study.

In this system, the search of the parameter space is controlled by a Genetic Algorithm. The GA was chosen for its ability to optimize while at the same time exploring the parameter space. The discussion here will be limited to the specific GA implemented for this system. GAs are a broad class of algorithms, so for a more general discussion see Goldberg.[11] As described in the background, the Genetic Algorithm works by building and running "generations" of reservoir models.

The first generation of models is created by randomly sampling the probability distributions of the uncertain parameters. The synthetic and live field tests conducted here both used generations of 250 models. So in both tests, 250 random reservoir models were created to form the initial generation. Each initial model is run in the reservoir simulator. The results of the simulation run are read and used to calculate an objective function as described below.

For this process to be practical, simulation runs must automatically run without human interaction. In this study, a commercial simulator, Eclipse, was used for which we did not have access to the source code. This required the creation of a "wrapper" around the simulator. This simply entails additional code to create a file for each run, submit it to the simulator, and read the results. This process, obviously, could be streamlined by working directly with the simulator source code.

After creating the initial generation, subsequent generations are created by both “breeding” new models from the previous generation and generating new random models. In this implementation the GA uses a “tournament” breeding selection technique. To create a new model two pairs of models from the previous generation are selected at random. The model with the lower objective function value is then selected from each pair. A new model is then created by “cross-breeding” the selected models. Cross-breeding is used here to randomly sampling parameter values with equal weight from the two selected models to create a new model. The idea here is that selecting parents based on their ability to reproduce observed data, as measured by the objective function, will in turn lead to “child” models that better match the field history. This mirrors the natural concept of “survival of the fittest.”

Another concept from nature incorporated in the GA is “mutation.” In the mutation process the value of an individual parameter is replaced with a value randomly sampled from the probability distribution of that parameter. The purpose of mutation is to force a more thorough exploration of the parameter space, and thus prevent the GA from getting stuck in local minima, by investigating additional parameter values. In this work a mutation rate of 10% was used. This means that in generating a new reservoir model a given parameter has a 10% chance of being mutated and replaced with a new value sampled from the underlying distribution of that parameter.

In addition to models created by breeding, each generation contains a set of new random models created in the same manner as those in the initial generation. This is a form of “migration” in which a generation includes new “immigrant” models alongside models generated from the previous generation. The purpose of these models, like mutation, is to ensure that the search is broad and does not get trapped in isolated regions of the parameter space. Because it consists of models generated by randomly sampling parameter distributions, this set, as well as the initial generation, can be thought of as a small Monte Carlo method. Because these random models are run alongside models generated by the GA, they provide an opportunity to evaluate the GA’s ability to optimize and generate better than random history matches. **Fig. 1** shows the average objective function value for random models compared to those generated by the GA for the live reservoir test. From this figure we see that the GA produces models that are, on average, better than random models.

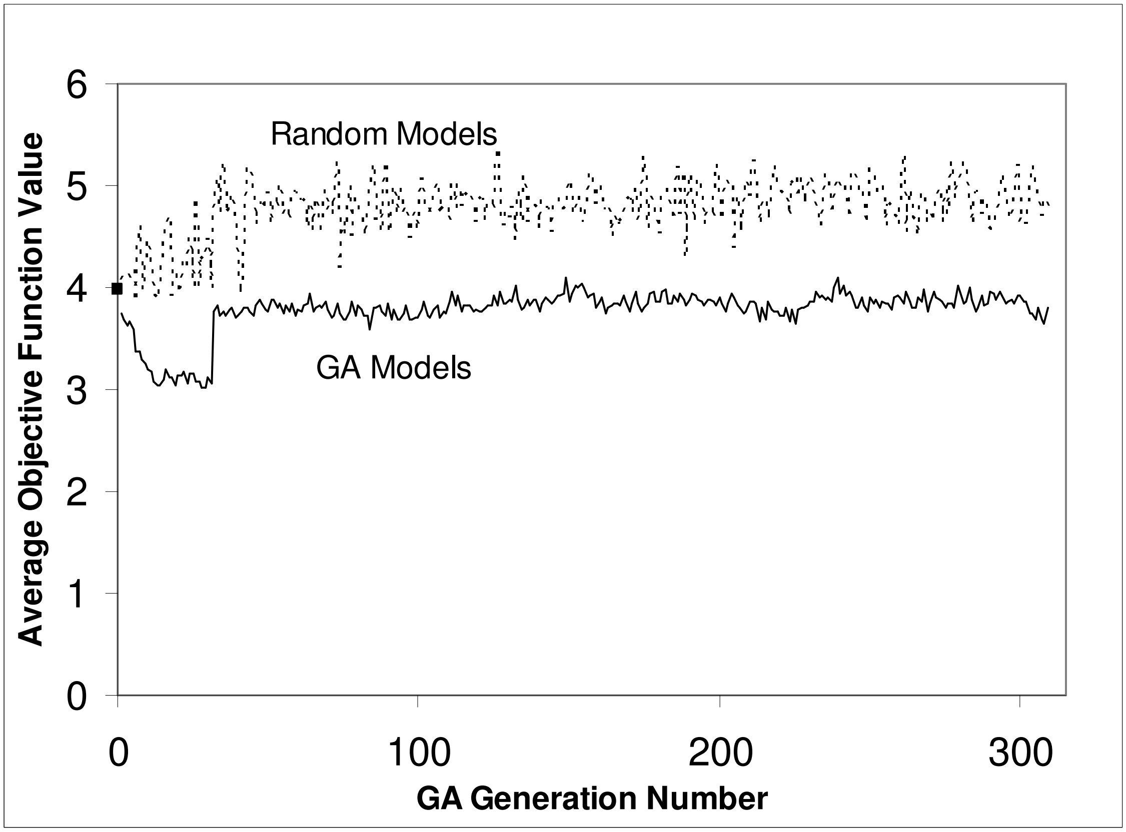

Fig. 1 – Value of optimization – Field A. On average, the GA generates models that better reproduce the field's history than do models randomly generated from the prior. Also, we see a major shift in the average objective function value in the 32nd generation. This shift corresponds to the introduction of tubing pressure data into the objective function calculation. In addition, after the first 10 generations or so the average objective function value of GA models appears to remain essentially constant with the exception of the shift due to the introduction of tubing pressure data. This stabilization is believed to be a function of the randomness introduced in each generation through both immigrant models and mutation.

Objective Function

An objective function is used to quantitatively evaluate how well an individual model reproduces the observed data from the field. This function is used in the construction of probabilistic forecasts and is utilized by the GA to guide the parameter space search. At a minimum a reservoir simulation objective function attempts to measure the ability of a model to reproduce dynamic field data, such as production and pressure data. Because the objective function is used to guide the parameter space search, the objective function also attempts to measure how well a model honors the prior static data. This is accomplished through the use of a prior term that is designed to measure the deviation of static parameter values in a given model from their expected value.

The objective function used in this work is similar to the objective function definition used in Eclipse's SimOpt package.[17] This function takes the form:

$$f = 0.75 * L + 0.25 * f_{prior} \qquad (1)$$

Where the likelihood term is the modified sum-of-squares term given by:

$$L = \left[\frac{1}{n} \sum_{i=1}^{n} \left[\frac{y_i^{calc} - y_i^{obs}}{\sigma_i} \right]^2 \right]^{0.5} \qquad (2)$$

The prior term is also modeled with a modified sum-of-squares. All parameters used in this study are modeled with a log normal distribution, so the prior term is given by:

$$f_{prior} = \left[\frac{1}{n} \sum_{i=1}^{n} \left[\frac{Log(\mu_p) - Log(p_i)}{\sigma_p} \right]^2 \right]^{0.5} \text{..............................(3)}$$

The prior term assists the search of the parameter space by preventing the Genetic Algorithm from tending toward extreme values. This term is not intended to be a rigorous statistical evaluation of the model's fit to prior data. All parameters used are log normally distributed multipliers with a mean value of 1 and Log(1) equals 0, so Eq. 3 can be reduced to:

$$f_{prior} = \left[\frac{1}{n} \sum_{i=1}^{n} \left[\frac{Log(p_i)}{\sigma_p} \right]^2 \right]^{0.5} \text{...(4)}$$

Continuous Data

At various points in time during this simulation process new data from the field will become available. It is advantageous to include new data in the process as quickly as possible because, at least in theory, more information about the field should lead to better forecasts and assessments of uncertainty. We can not simply add data as soon as it becomes available, however, as additional data will alter the objective function definition and could disrupt the selection process used by the GA to choose "parent" models. For this reason data are only added in between GA generations.

Forecasting

The final step in the continuous simulation process is combining the results of simulation runs into probabilistic forecasts. Forecasting is not done continuously, but rather at discrete points as needed. As in the Barker *et al.*[7] importance sampling and pilot-point approaches, forecasts are generated only using runs with an objective function below a certain acceptable threshold. This threshold value, 1.7 in Barker *et al.*[7] and the synthetic reservoir test, is the level at the match to observed data is deemed adequate. In addition, because the GA can over-sample a limited subsection of the parameter space, the acceptable models are compared to each other and if any two were deemed too close in terms of parameter values the model with the higher objective function value was removed. This pair by pair comparison is calculated as the sum of squares difference between the parameter values in one model and the corresponding parameters in the other. This results in an N-1 by N-1 comparison between all N acceptable models, which is computationally intensive. In practice this comparison resulted in the removal of just a handful of runs in the synthetic test and none in the live field test. After this filtering a forecast is generated in which the forecast values of acceptable models are equally weighted in the forecast distribution. This equal weighting is justified given the large sample of acceptable matches obtained by the continuous simulation process. Because this process gives a large sample of matched models, we can obtain smooth and complete forecast distributions with equally weighted forecasts. This is in contrast to other methods that attempt to define a forecast distribution based on at most a couple dozen matched models. In these methods weighting using the objective function is needed to infer a forecast distribution shape from a limited sample of the parameter space. Since we have a large and broad sample of the parameter space we can simply rely on relative frequency to construct our forecasts.

Barker *et al*.[7] provide an alternative approach to creating probabilistic forecasts which they claim is statistically rigorous. Rather than use equal weighting, Barker models uncertainty using the exponential likelihood function:

$$L = c * \exp\left[-\frac{1}{2}\sum_{i=1}^{n}\left[\frac{y_i^{calc} - y_i^{obs}}{\sigma_i}\right]^2 \right] \quad \ldots\ldots\ldots\ldots\ldots\ldots\ldots\ldots\ldots\ldots(5)$$

They state that Eq. 5 requires that the production data be independent measurements with normally distributed error. Unfortunately, the authors neither reference nor provide a derivation of this formula. However, Eq. 5 appears to be an adaptation of likelihood function for normal distributions, given by Vose[18] as:

$$L(\mu,\sigma) = \left(\frac{1}{\sqrt{2\pi\sigma^2}}\right)^n \exp\left[-\sum_{i=1}^{n} \frac{(x_i - \mu)^2}{2\sigma^2} \right] \quad \ldots\ldots\ldots\ldots\ldots\ldots\ldots..(6).$$

Here x_i is observation from an independent experiment. The major problem with adapting this formula for use in production forecasts is the assumption of *independent* measurements with normally distributed error. In production forecasts the same observation (such as the pressure in a given well) is made at multiple points in time. Obviously, the pressure in a well is not completely independent from the pressure at an earlier or later point in time. When dependant data points such as these are used in the likelihood function the assumption of independence is violated and the statistical validity of the approach is called in to question. Without any guidance from the authors in the form of a derivation or reference, this issue cannot be reconciled and, for this reason, I do not use this likelihood definition to weight forecasts. I have included this discussion simply to explain why my approach differs from previous work.

Summary

Thus, we see that continuous simulation is a multi-step process. First, a suitable parameter space is defined. Next, the GA explores this parameter space by generating and running models. The GA is guided in its search by an objective function evaluation of each model, which is also used for generating forecasts. This objective function is updated with new data soon after it becomes available. Finally, the results of individual runs are combined into probabilistic forecasts.

SYNTHETIC RESERVOIR STUDY

Overview

The first test of the continuous simulation process was conducted on the PUNQ-S3 synthetic reservoir. This reservoir is a synthetic reservoir used in the PUNQ discussed above. In the PUNQ work simulation runs were matched against 8 years of observed data and forecasts were made out to 16.5 years of production. In this synthetic test the PUNQ-S3 field was continuously simulated starting in the 4th year of production and continuing through the end of the 8th, making forecasts out to 16.5 years. During this simulated 5 year period 45,000 simulation runs were made. The results of these runs were combined into probabilistic forecasts at several points during the test. Before examining these runs, however, we will look at the PUNQ-S3 model and the parameterization used in this test.

As mentioned, the PUNQ-S3 synthetic reservoir was used in the PUNQ study described above and has been extensively studied by others since. An Eclipse simulation model and other associated data for this reservoir are publicly available online.[19] By most standards the PUNQ-S3 reservoir is a small model with just 1761 active cells. On a modern desktop computer a single simulation run takes less than a minute, which is advantageous for making a large number of runs.

The PUNQ-S3 reservoir model is a five-layer, three-phase synthetic reservoir based on an actual field operated by Elf. The field contains six producing wells, which are shown on a structure map in **Fig. 2**. Layers one, three and five are of relatively high quality with maximum porosity of roughly 30% and maximum horizontal permeability of about 1 Darcy. Layers two and four are of substantially lower quality. The truth case porosity, horizontal and vertical permeability maps are shown in **Figs. 3-17**.

I parameterized the PUNQ-S3 reservoir using six homogenous regions per layer. Included in the online PUNQ-S3 dataset is a geological description. This geological description indicates that the reservoir is marked by wide southeasting high-quality streaks. For this reason I defined regions that approximate these streaks, rather than using rectangular regions. The regions are shown in **Fig. 18**. While a more rigorous parameterization based on geostatisical methods may be possible, the use of homogeneous regions was chosen in order to be consistent with the actual field case where a lack of data prevents a more complex parameterization.

Instead of using porosity and permeability values directly, the parameters used are porosity and permeability multipliers. These multipliers are applied to permeability and porosity base maps in running the simulation. The effect is the same as if porosity and permeability values were used directly, but this approach simplifies the implementation. The base maps used constant values of porosity, horizontal permeability and vertical permeability by layer and these are listed in **Table 1**.

Six regions per layer times five layers times three parameters (horizontal permeability, vertical permeability and porosity) gives a total of 90 parameters. We are able to use so many parameters because the continuous simulation process, when run over time, allows us to make many more runs than normally possible. To generate new models multipliers are randomly sampled from known distributions. Both vertical and horizontal permeability multipliers are modeled using a log-normal distribution with a mean of 1 and a standard deviation of 1. The log-normal distribution was chosen based on Craig *et al*.'s[20] use of this distribution with layered reservoirs. In order to prevent extreme and unrealistic values of permeability the distribution is capped on the upper end at a value of 4. If the log-normal distribution is not capped then it is theoretically possible to have multipliers approaching infinity, which of course is unreasonable. The porosity multiplier was modeled using a log-normal distribution with a mean of 1 and a standard deviation of one-half. Porosity was capped with a maximum value of 2.28. The distributions used for permeability and porosity multipliers are shown in **Figs. 19** and **20**.

Table 1 – Synthetic Test Base Porosity and Permeability.

Layer	Porosity	Horizontal Permeability (md)	Vertical Permeability (md)
1	22 %	500	200
2	10 %	50	15
3	22 %	500	200
4	16 %	250	40
5	22 %	500	200

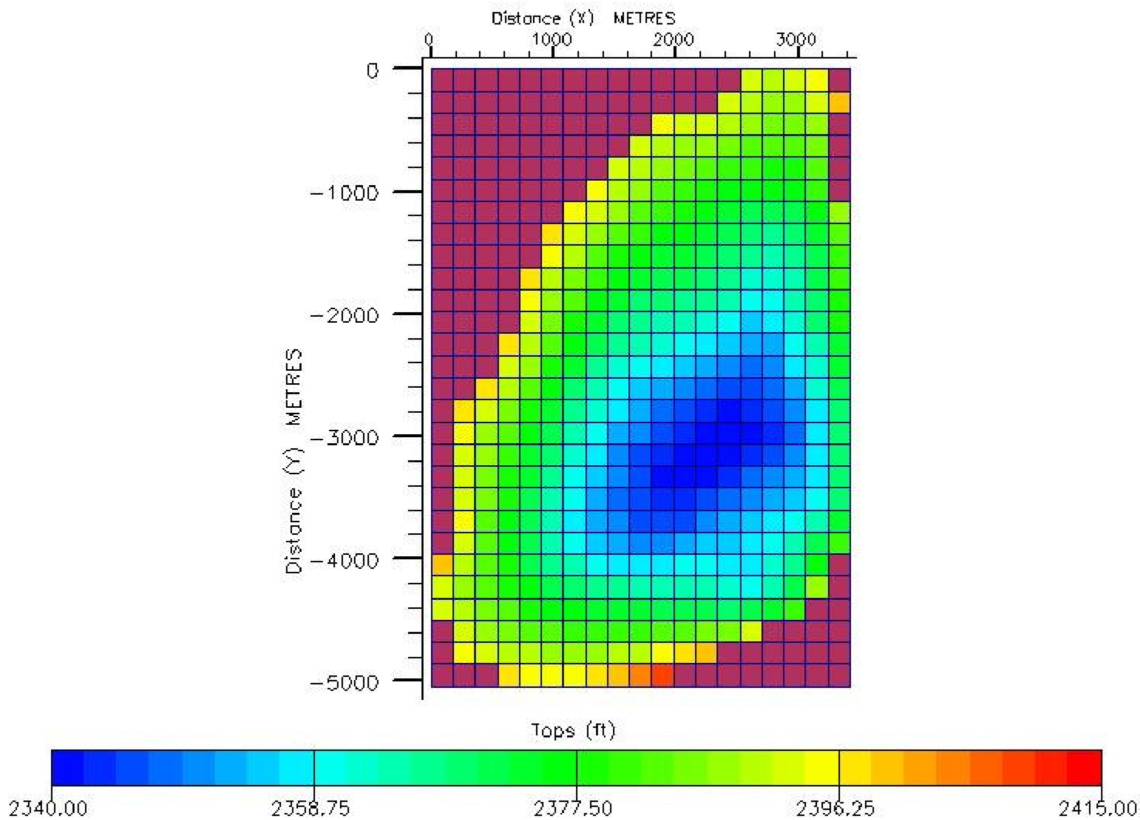

Fig. 2 – Structure of the PUNQ synthetic reservoir.

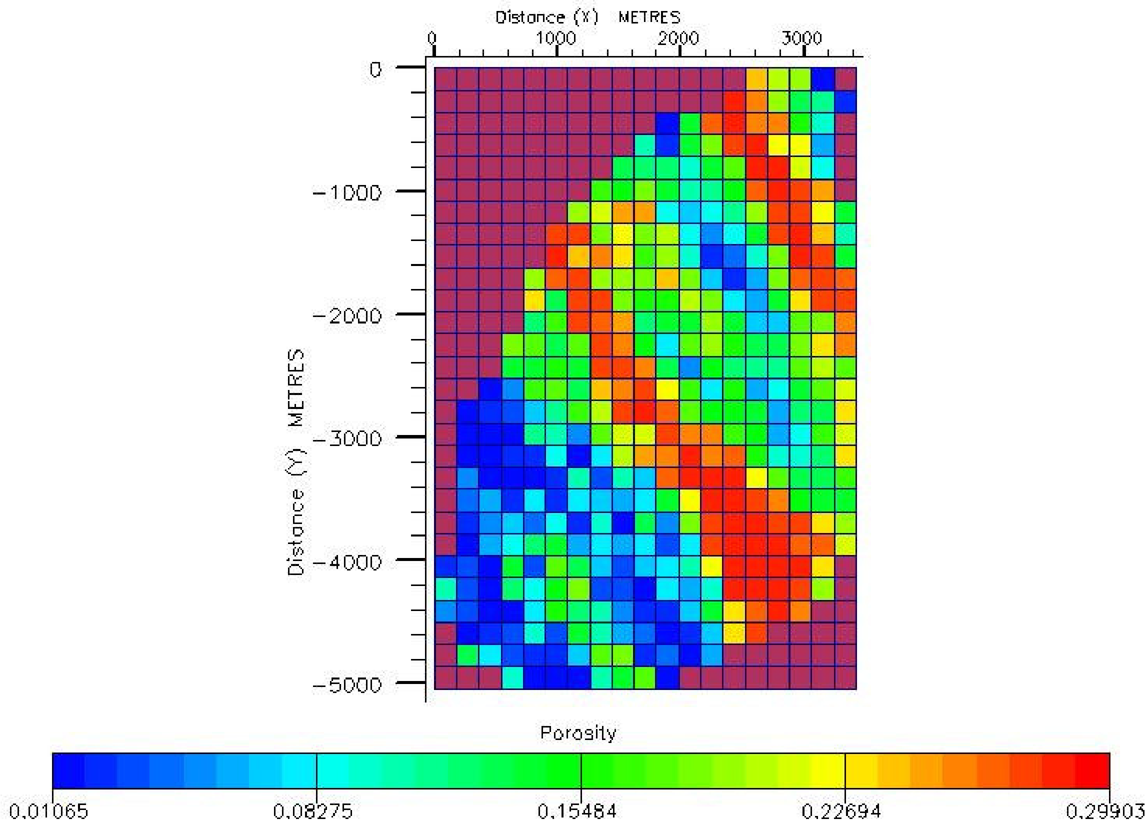

Fig. 3 – Truth case porosity for PUNQ reservoir - layer 1.

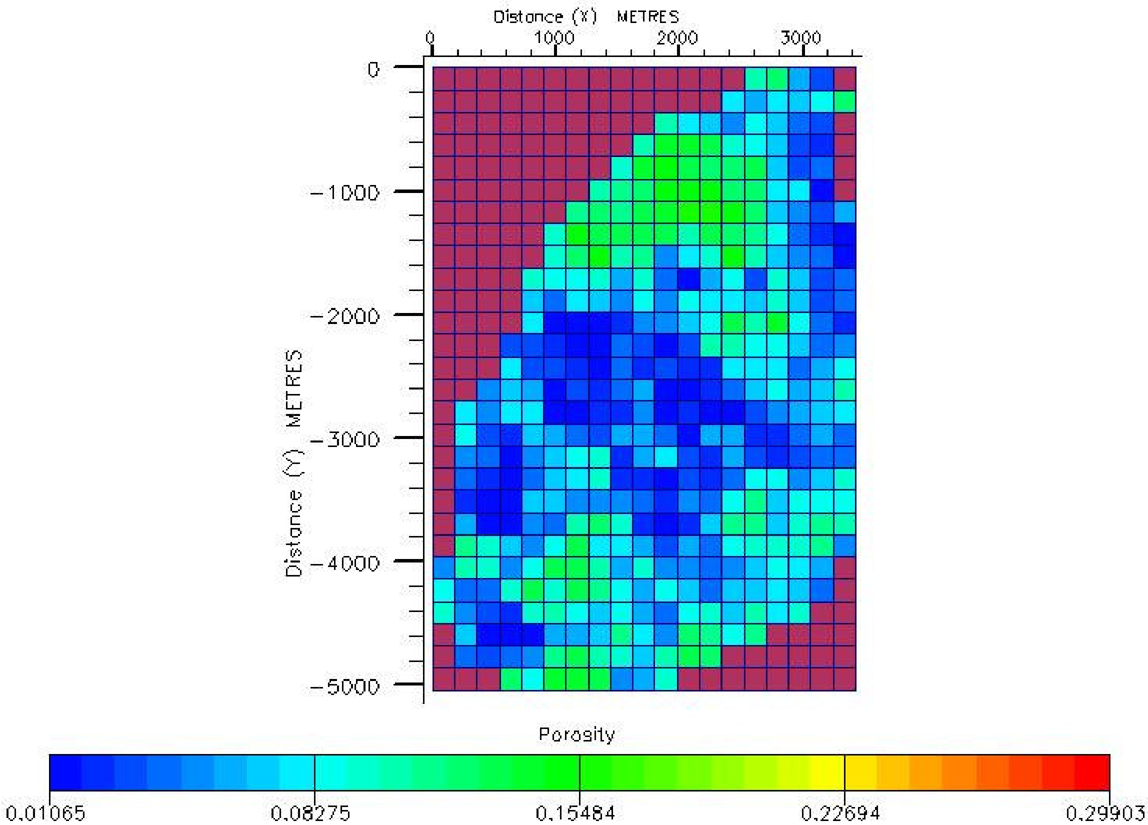

Fig. 4 – Truth case porosity for PUNQ reservoir - layer 2.

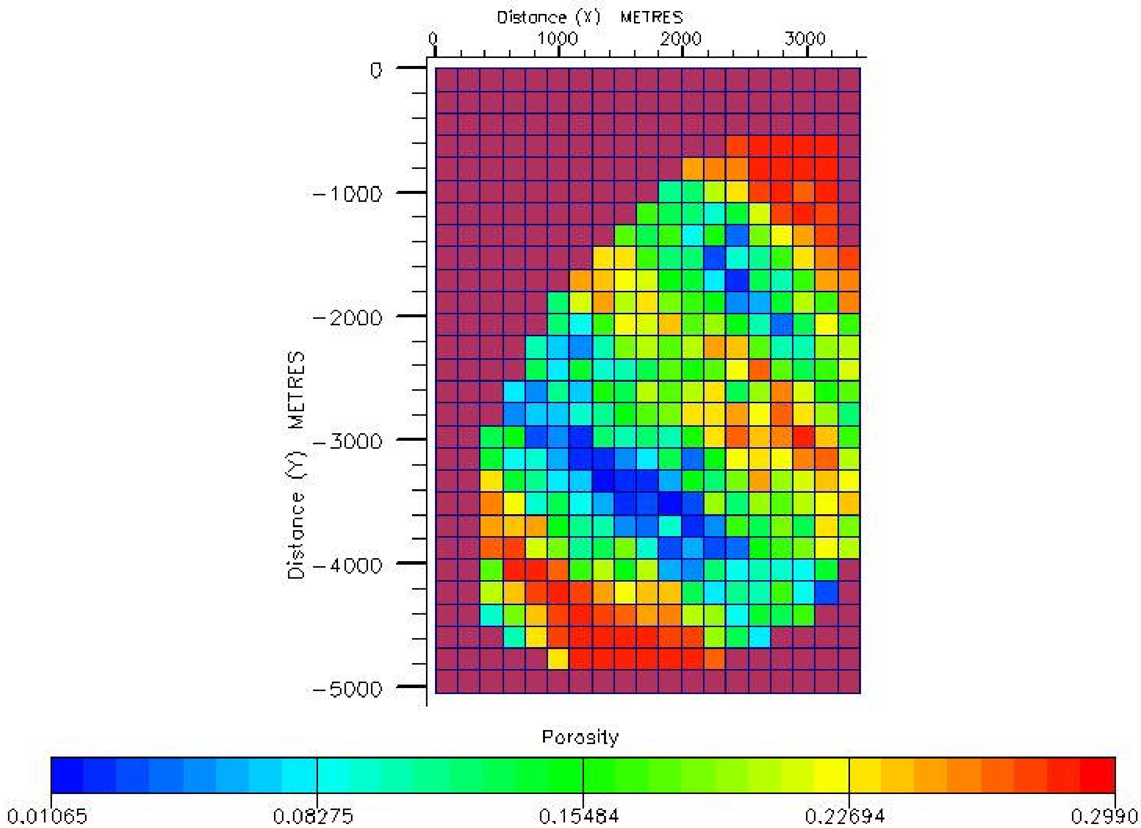

Fig. 5 – Truth case porosity for PUNQ reservoir - layer 3.

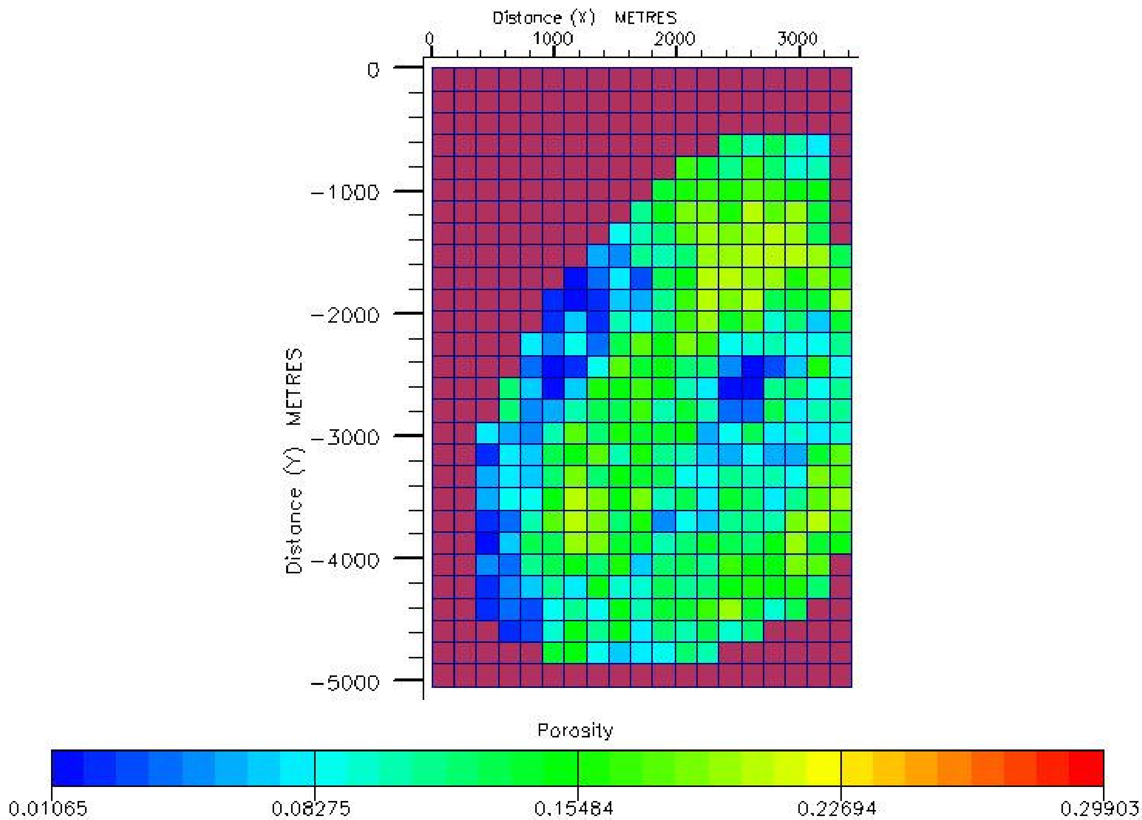

Fig. 6 – Truth case porosity for PUNQ reservoir - layer 4.

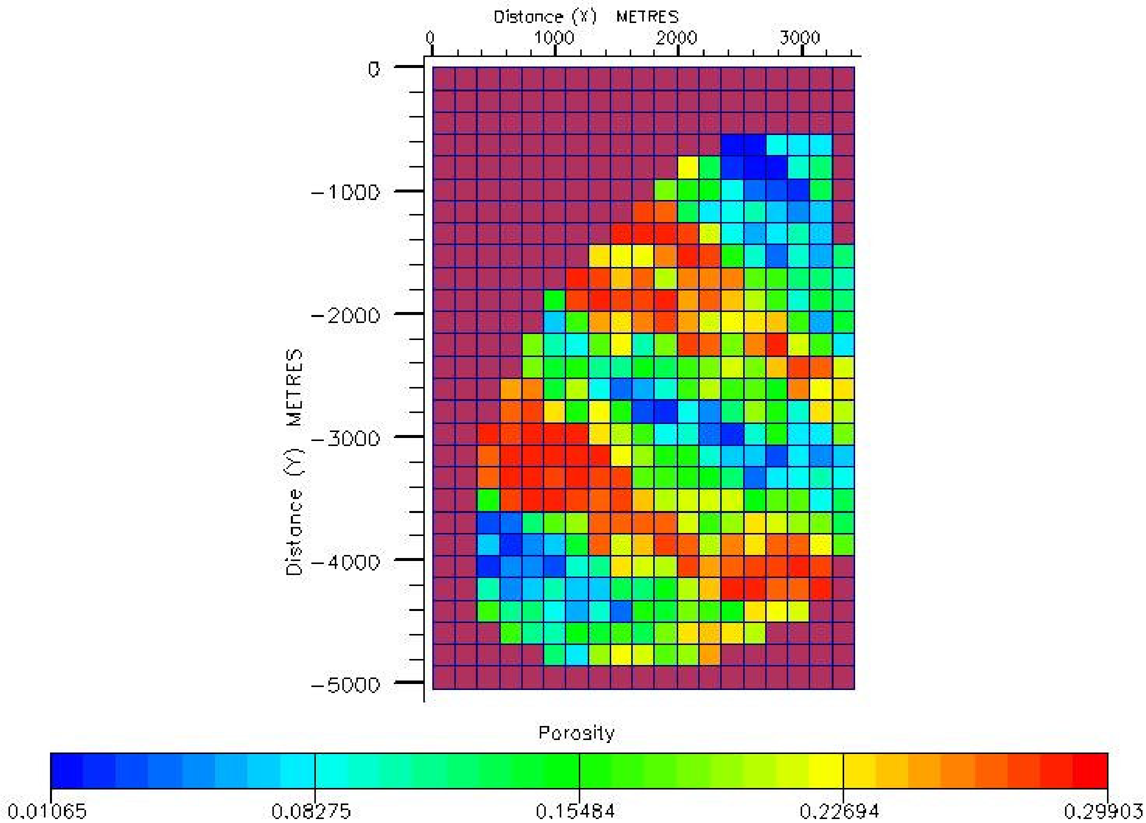

Fig. 7 – Truth case porosity for PUNQ reservoir - layer 5.

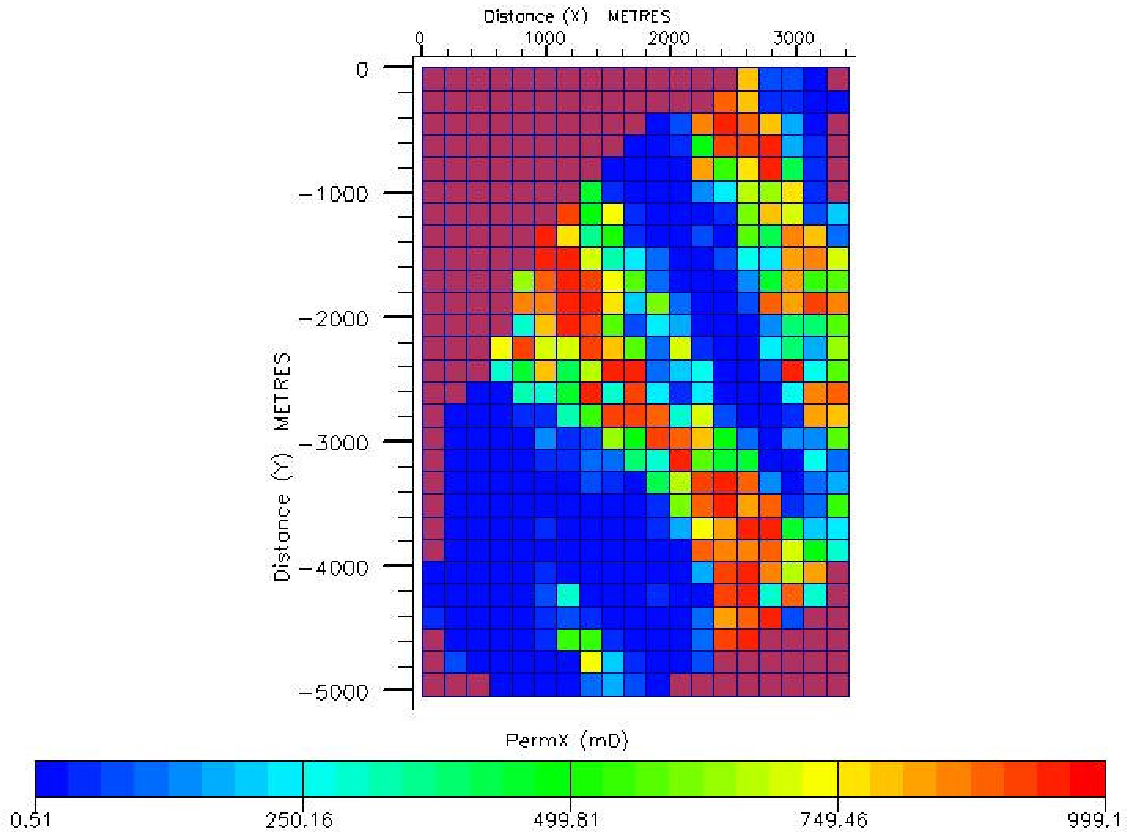

Fig. 8 – Truth case horizontal permeability for PUNQ reservoir - layer 1.

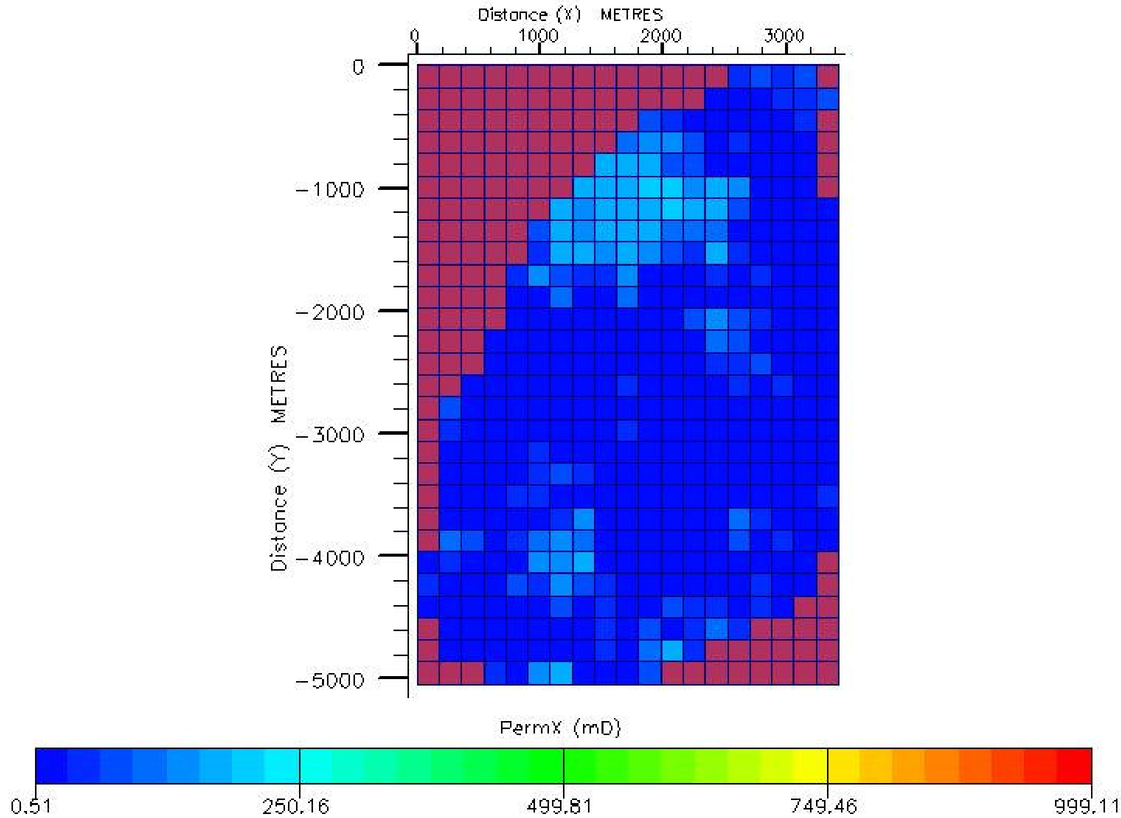

Fig. 9 – Truth case horizontal permeability for PUNQ reservoir – layer 2.

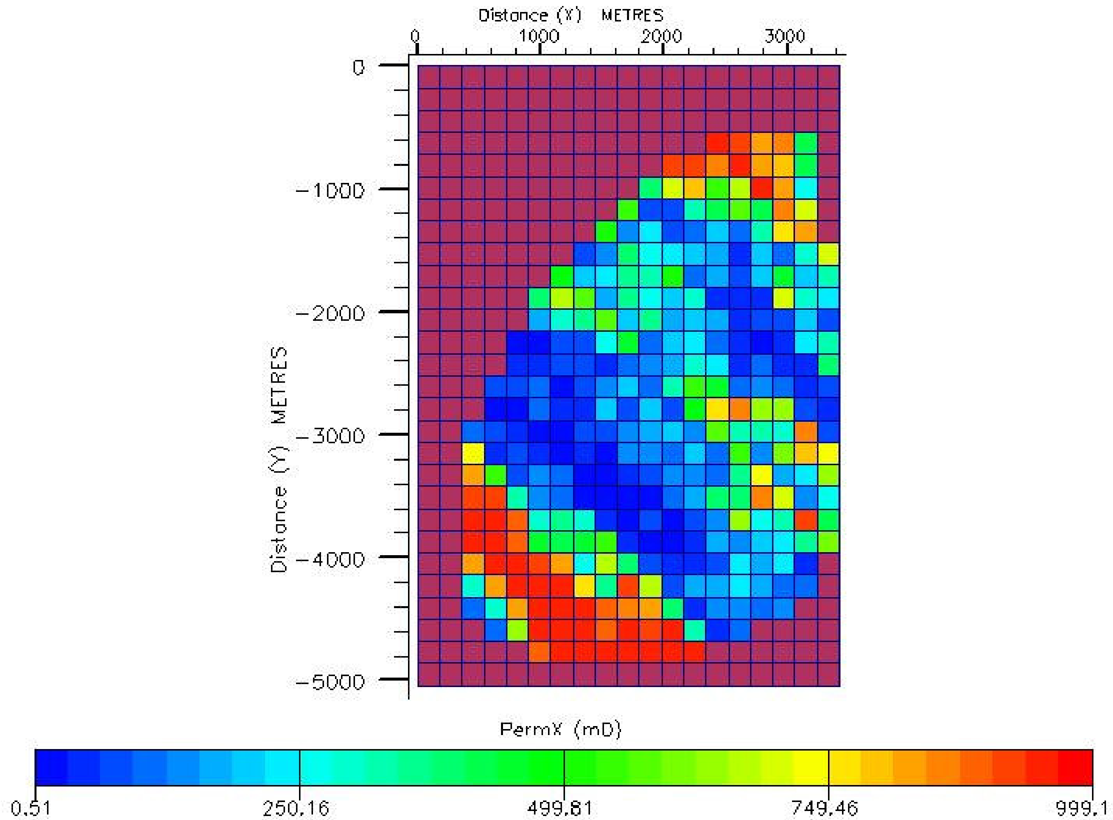

Fig. 10 – Truth case horizontal permeability for PUNQ reservoir - layer 3.

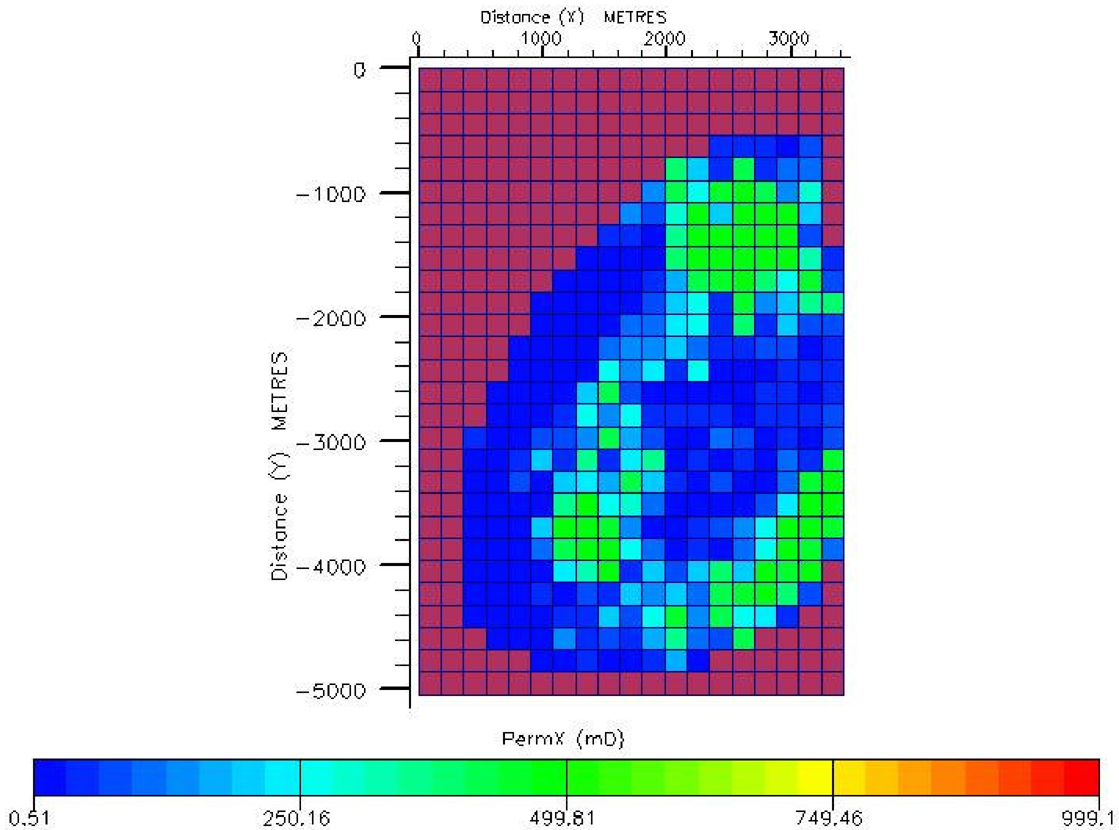

Fig. 11 – Truth case horizontal permeability for PUNQ reservoir - layer 4.

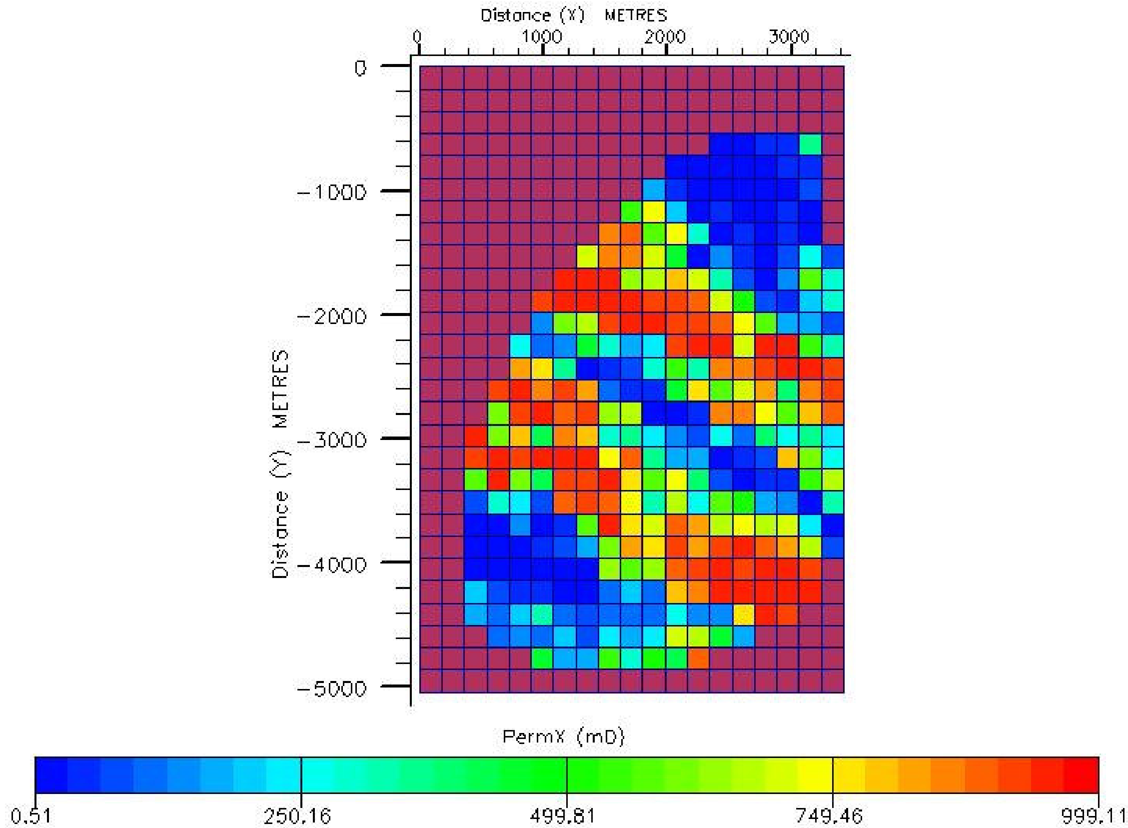

Fig. 12 – Truth case horizontal permeability for PUNQ reservoir - layer 5.

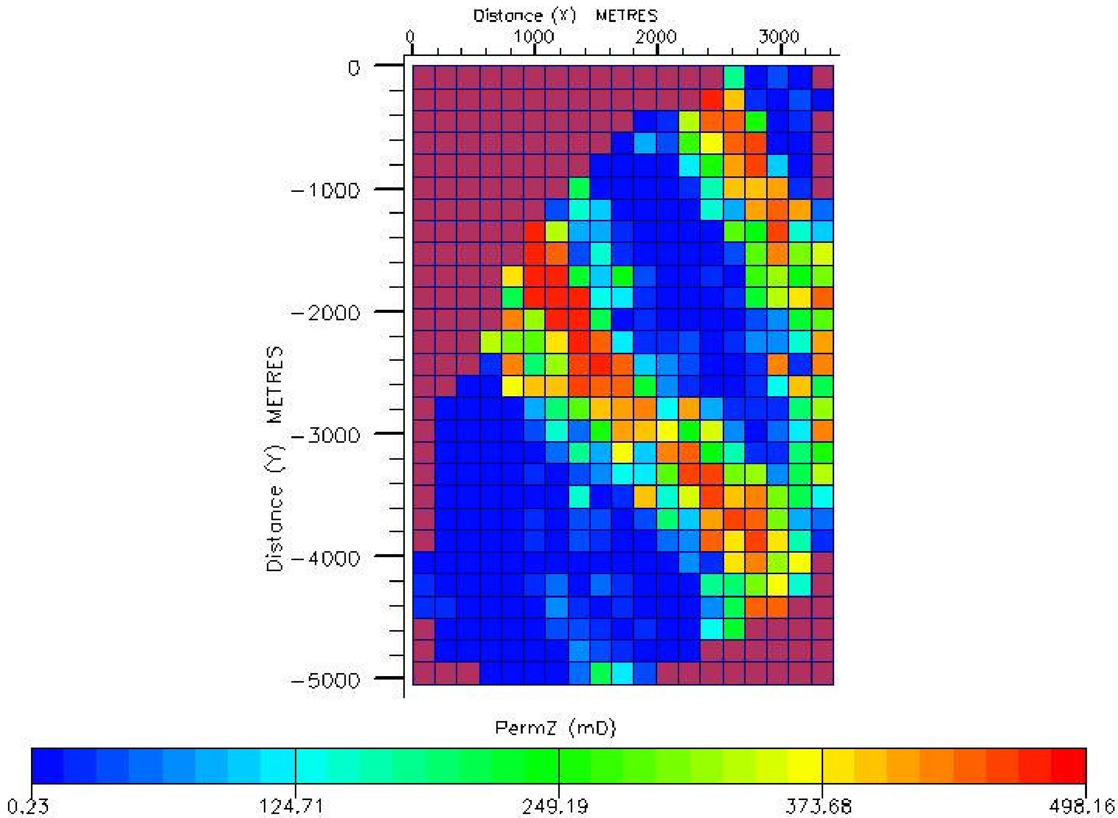

Fig. 13 – Truth case vertical permeability for PUNQ reservoir - layer 1.

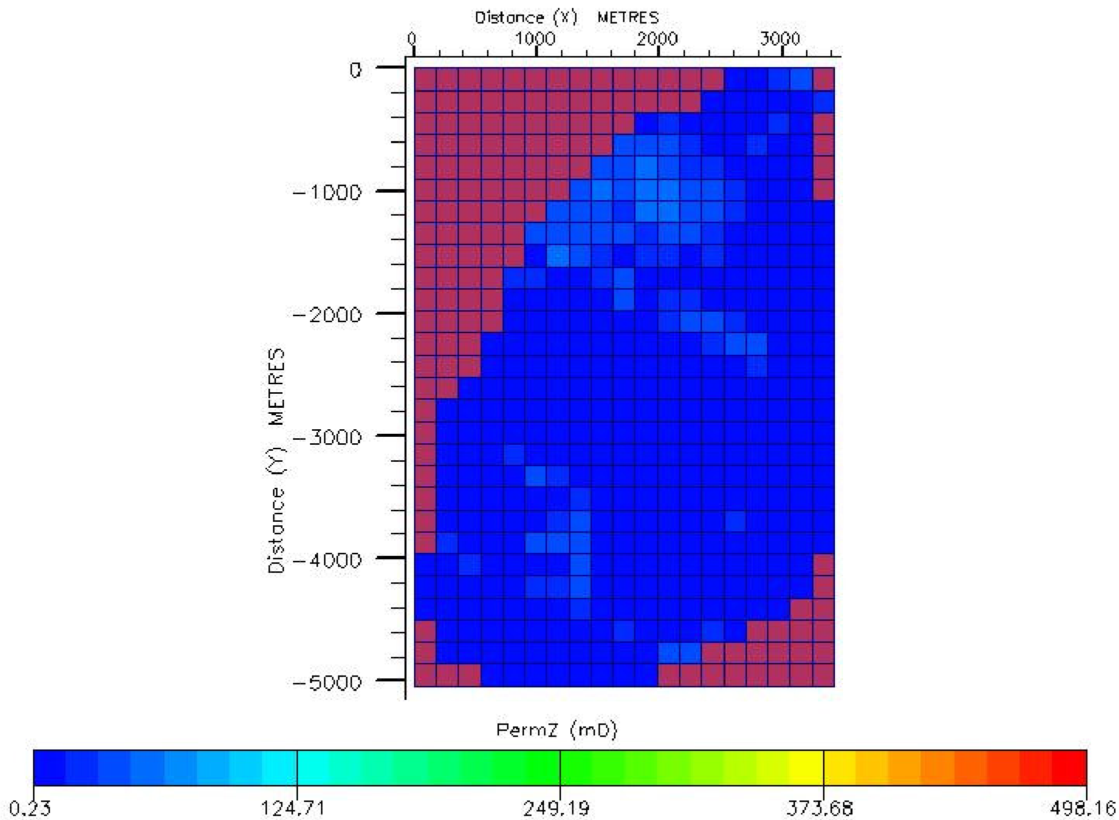

Fig. 14 – Truth case vertical permeability for PUNQ reservoir - layer 2.

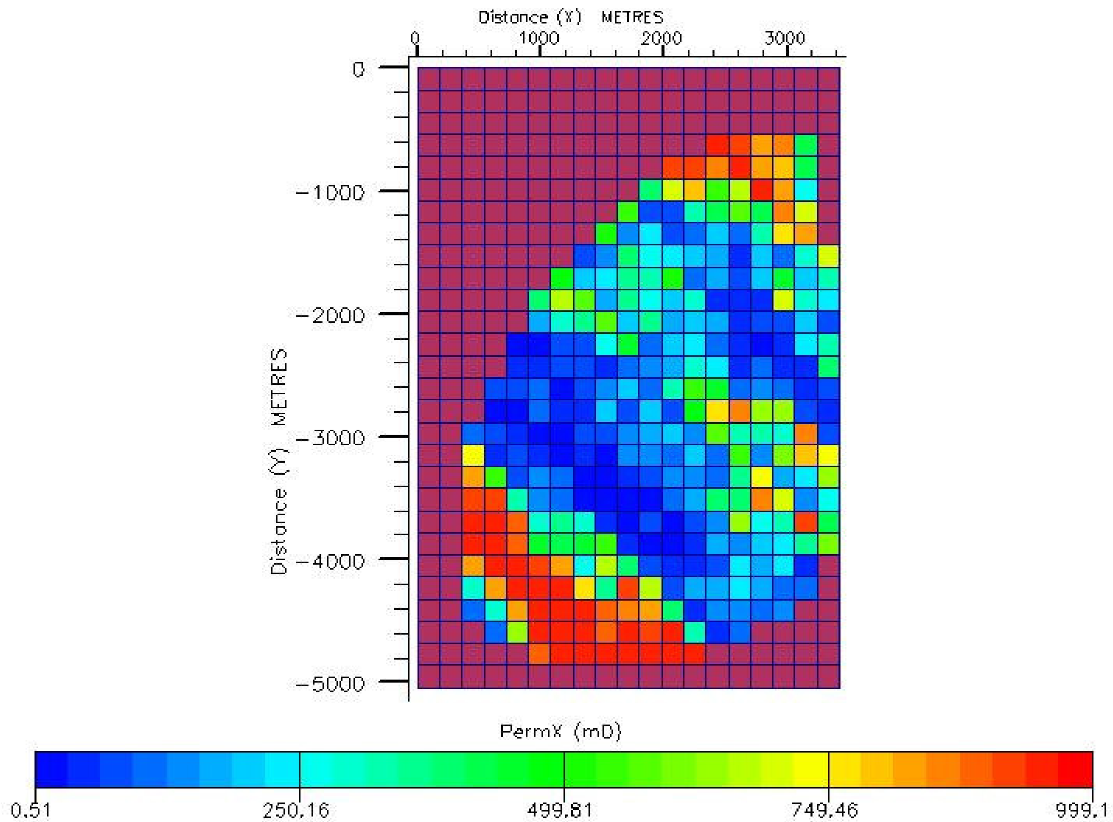

Fig. 15 – Truth case vertical permeability for PUNQ reservoir - layer 3.

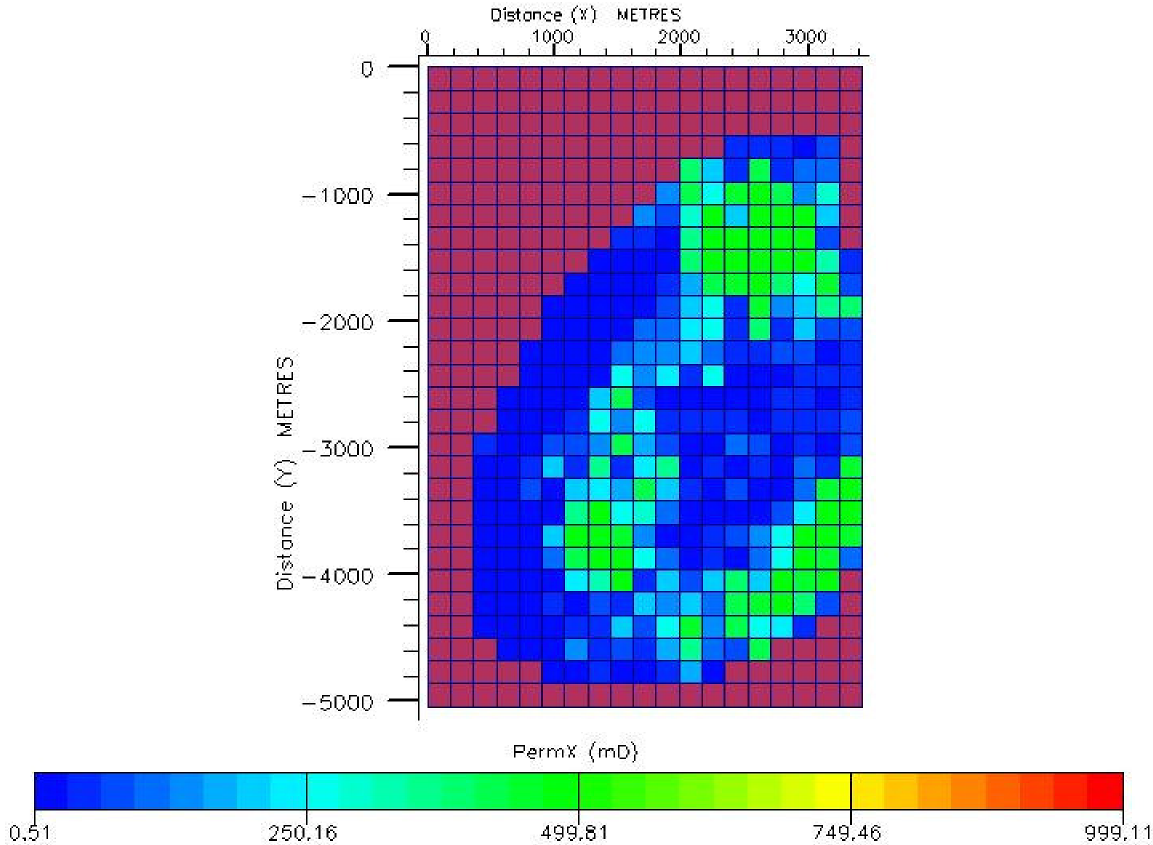

Fig. 16 – Truth case vertical permeability for PUNQ reservoir - layer 4.

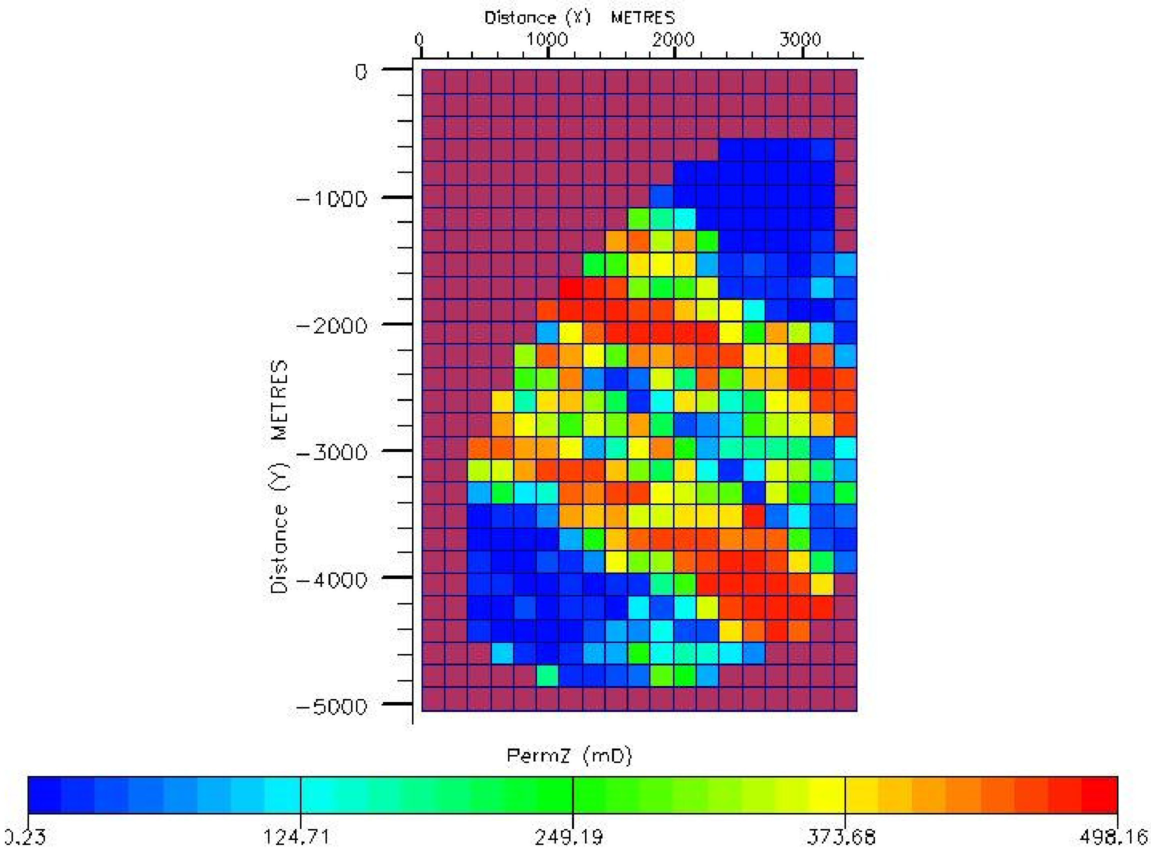

Fig. 17 – Truth case vertical permeability for PUNQ reservoir - layer 5.

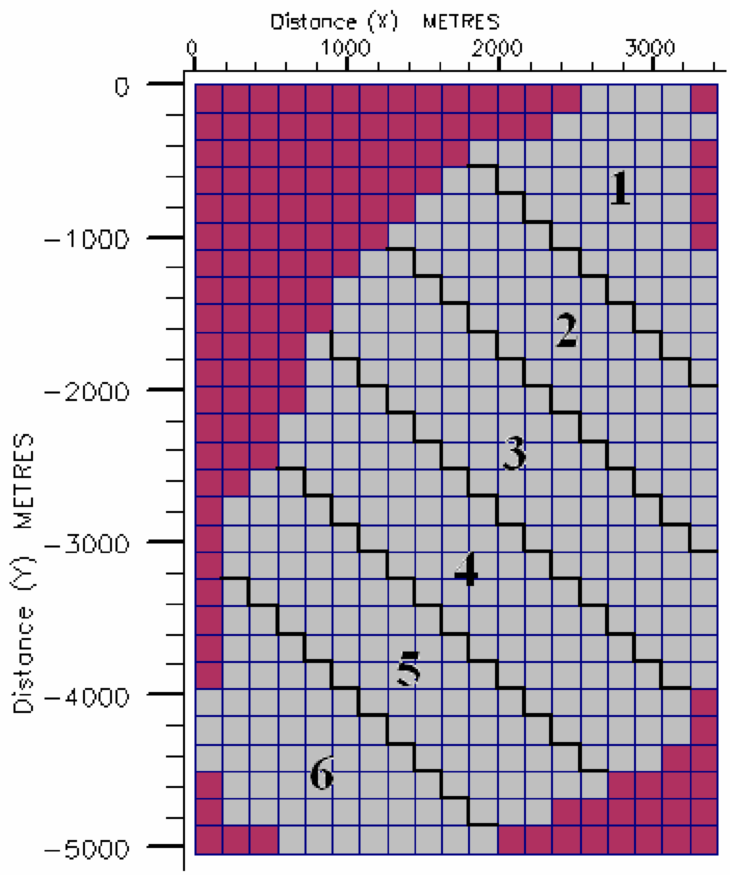

Fig. 18 – Synthetic test multiplier regions. The multiplier regions used to parameterize PUNQ reservoir for the synthetic test.

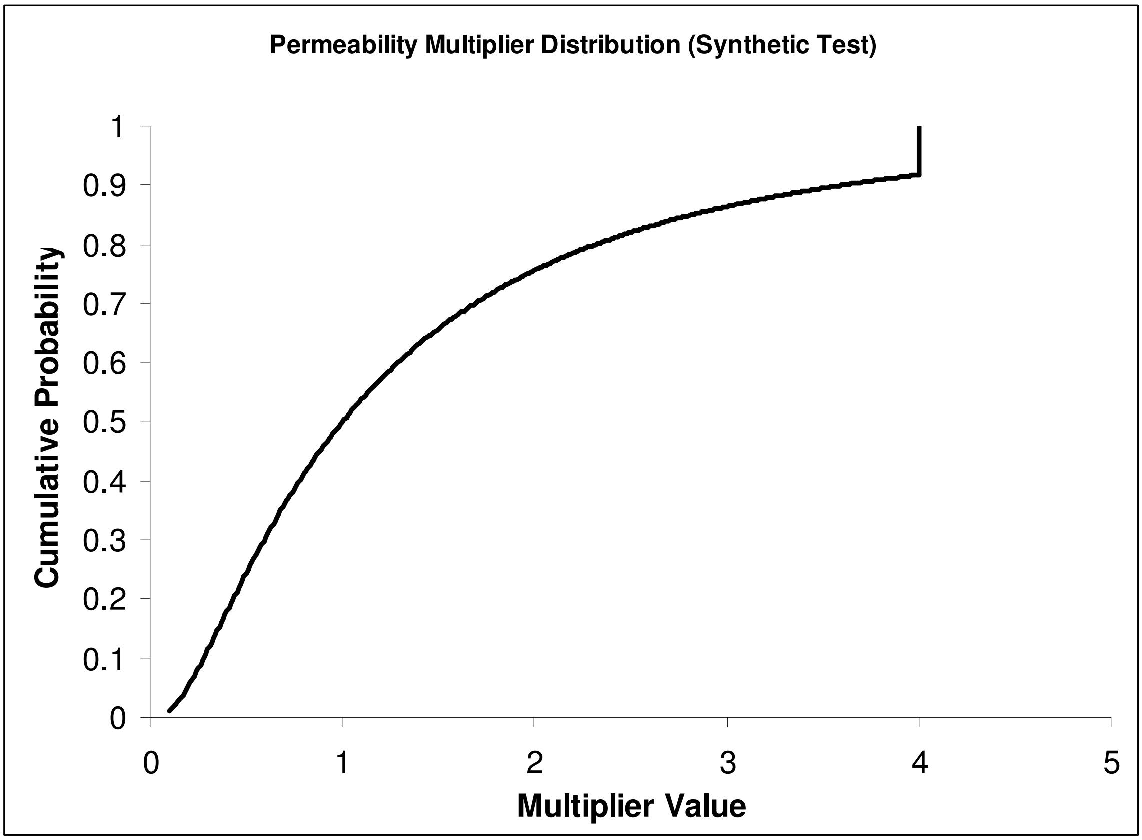

Fig. 19 – Synthetic test permeability multiplier distribution. The distribution used for the permeability (both vertical and horizontal) multiplier parameters in synthetic test. This is a capped lognormal distribution with a mean of 1, a standard deviation of 1, capped at 4.0.

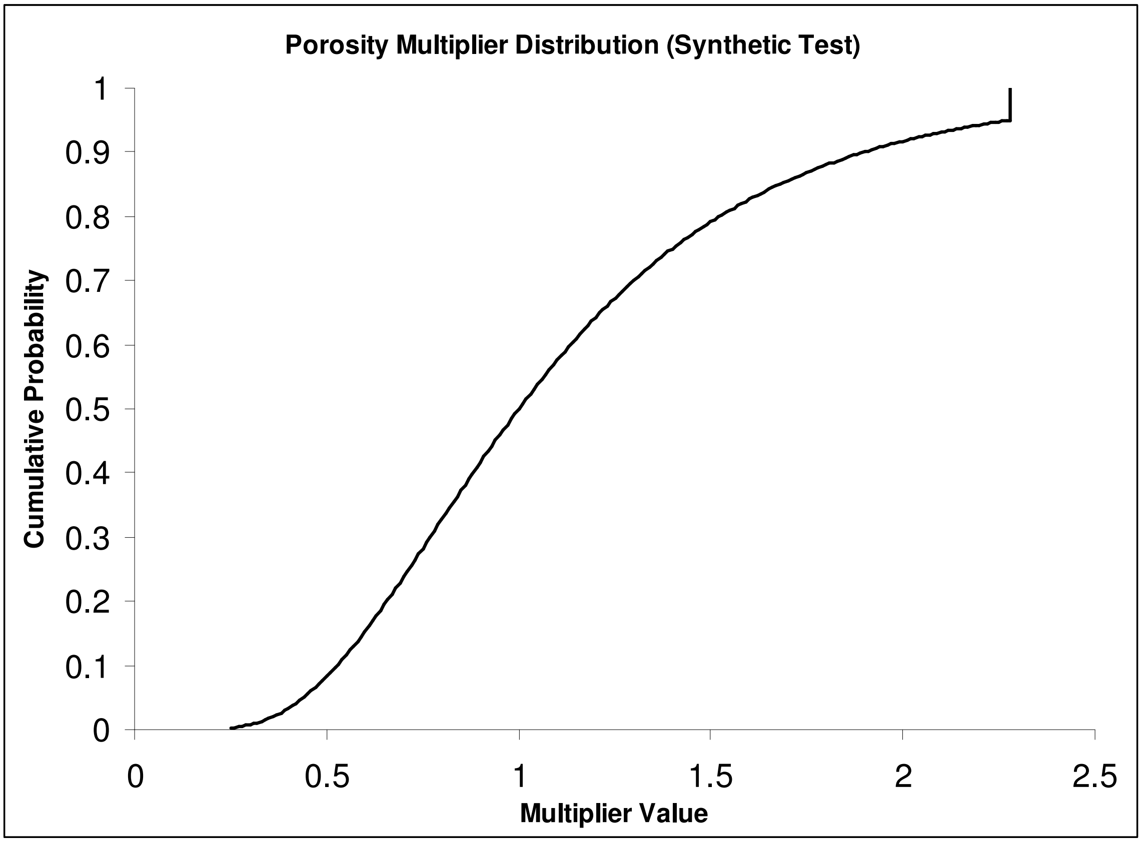

Fig. 20 – Synthetic test porosity multiplier distribution. The distribution used for porosity multiplier parameters in synthetic test. This is a capped lognormal distribution with a mean of 1, a standard deviation of 0.5, capped at 2.28.

Parameter Space Search

As mentioned above, 45,000 simulation runs were made corresponding to a five year period in the reservoir's life. Rather than run in real time, in order to continuously simulate over a significant percentage of the reservoir's life while conducting the test in a timely fashion, time was "accelerated" so that 750 simulation runs correspond to a month in the life of the reservoir. This effectively means that every simulation run maps to a point in the reservoir's life. For instance runs 1-750 were run in January of Year 4 and were matched against any data available at that time. Similarly, runs 9000 to 9750 map to January of Year 5 and were matched against the data available at that point in time. Stepping through the historical data one month (or 750 runs) at a time, 45000 runs of the PUNQ-S3 reservoir were made, replicating five years of continuous simulation. **Fig. 21** shows the cumulative number of runs made versus the producing time of the reservoir.

The data set used in the objective function is the same used in previously published work and is summarized in **Table 2**. We can see that new data is available roughly every half a year. This additional data was included in the objective function calculation at the corresponding point in time in the reservoir's life. Since the objective function is used in the GA for selection and adding new data essentially changes the objective function definition, care must be taken in when data is added in order to avoid disrupting the GA. So that the objective function used within a given GA generation is identical and directly comparable between all runs, data was only added at the beginning of a month (which corresponds to the start of new a generation).

Fig. 22 shows the objective function values for all runs made in this test, listed by time in the reservoir's life when they were run. We see that there are several points in the process where the objective function values shift. These shifts are caused by adding new data and thus changing the objective function definition. Because the objective function value changes with time, care must be taken when making comparisons between runs made at different points in time. This is especially an issue when combining individual runs into probabilistic forecasts. In order to address this forecasts are only created from runs made during a set period of time and thus evaluated with comparable objective functions.

Table 2 – Synthetic Test Observed Data. (After Gu and Oliver[16])

Time (days)	WBHP	WGOR	WWCT
1.01	6	-	-
91	6	-	-
182	6	-	-
274	6	-	-
366	6	-	-
1461	6	-	-
1642	-	6	-
1826	6	6	-
1840	6	-	-
1841	-	6	-
2008	-	6	-
2192	6	6	-
2206	6	-	-
2373	-	6	-
2557	6	6	-
2571	6	-	-
2572	-	-	1
2738	-	6	-
2922	6	6	6
2936	6	-	-

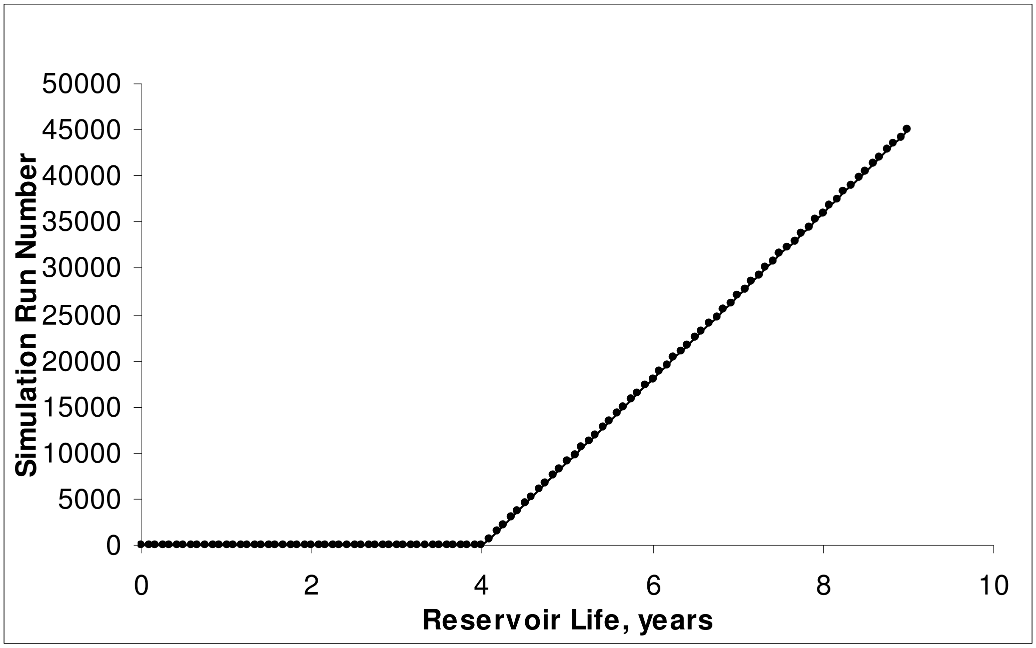

Fig. 21 – Synthetic test run number by time. Run number versus point in reservoir life for the synthetic test.

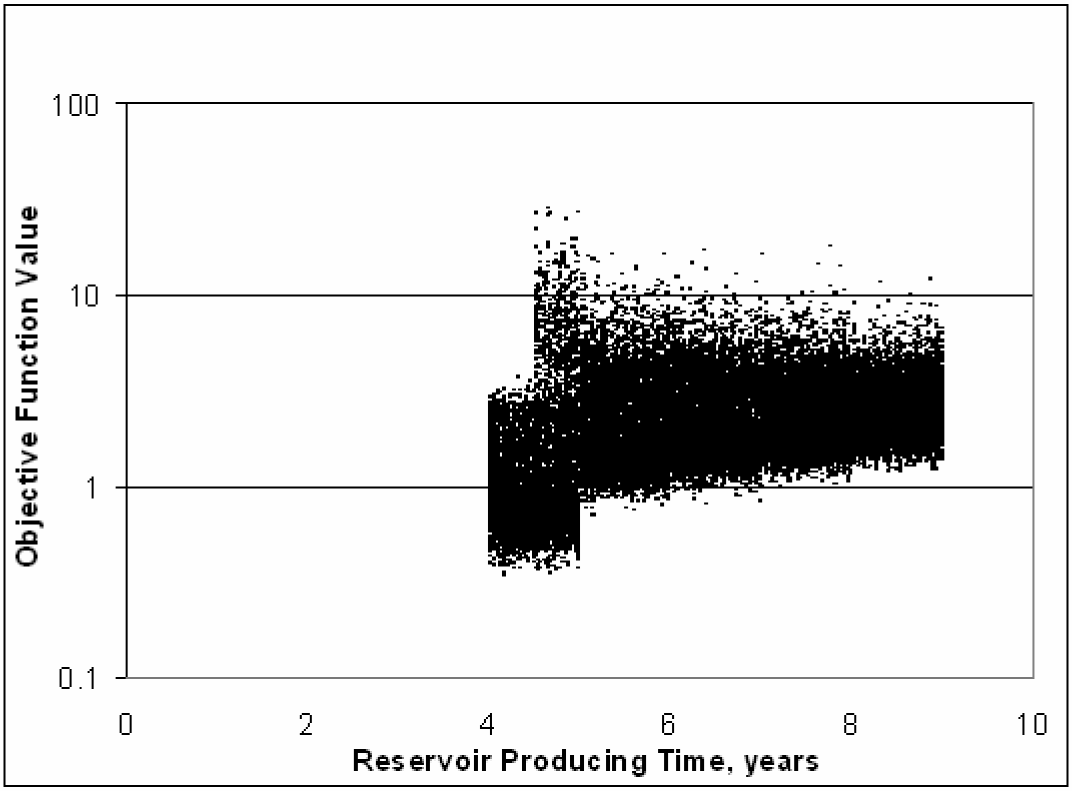

Fig. 22 – Synthetic test objective function. Objective function values for all runs made during the synthetic test. We see several shifts in the magnitude of the objective function, notably at 4.5 and 5 years. These shifts correspond to the introduction of additional data points to the objective function.

Forecasts

At discrete points throughout the simulation process probabilistic forecasts were generated which represent the forecast that would have been available at a given point in the life of the field. Forecasts were made at 4.5 years, 5 years, 6 years, 7 years, and 8 years. These probabilistic forecasts were created by taking all the runs made over the past year (or half year at 4.5 years) with an objective function value below a fixed cutoff of 1.7 as done in Barker *et al.*[7] In addition, the runs were filtered to remove runs with nearly identical parameter values. The purpose of this filtering is to prevent the forecast from being biased due to over sampling a particular region of the parameter space. After filtering, each run was given equal weight in the forecast. These forecasts are shown individually in **Figs. 23-28** and the cumulative distributions of these forecasts are shown together in **Fig. 29**.

These probabilistic forecasts are shown together in **Fig. 30** along with the PUNQ-S3 forecasts published in Barker *et al.*[7] and forecast ranges created using the EnKF by Gu and Oliver.[16] It should be noted that all the published forecasts, including the EnKF ranges, were created using the full 8 years of production history. We see that the uncertainty predicted by most of the previous work falls within the range predicted by the new forecasts. Also, the mean value of these forecasts lies very close to the truth case. Capturing the truth case as well as the range of uncertainty predicted by previous studies provides anecdotal evidence that the approach taken in this study quantifies the uncertainty in forecast at least as well as other methods. While differences in parameterization and simulators make direct comparisons precarious, we can draw some general conclusions from this figure. First, we see that the uncertainty ranges in our forecasts are wider than most published forecasts. This is likely a result of the wide uncertainty we considered in the parameterization this reservoir, which in turn we were able to explore thanks to the large number of runs enabled by the continuous simulation process. In light of Capen's[4] work, which demonstrated the tendency to vastly underestimate uncertainty and the fact that several published forecasts miss the truth case, perhaps it is desirable that our uncertainty quantification be wide enough to ensure they reliably predict the truth case. In addition, we see a lot of scatter in the ranges of the published forecasts which qualitatively suggests a lot of uncertainty surrounding these forecasts. The fact that most of the published ranges lie within our ranges indicates that our approach does a good job in quantifying this uncertainty.

In addition, we see in **Fig. 30** that as time progresses the forecast distributions narrow and shift slightly. This is the behavior we would expect as additional information (i.e. new production and pressure data) should alter our assessments if the data are of any value. We see that this narrowing and shifting is most dramatic in the early forecasts. Again this seems reasonable as early on the data set used in the objective function is smaller and each new data point will have a larger impact on the objective function value. As time progresses, the size of the data set grows. In turn the relative impact of any individual data point decreases and the narrowing continues, but appears less dramatic. Stated more generally, as more information about an event becomes available the uncertainty around that event should decrease.

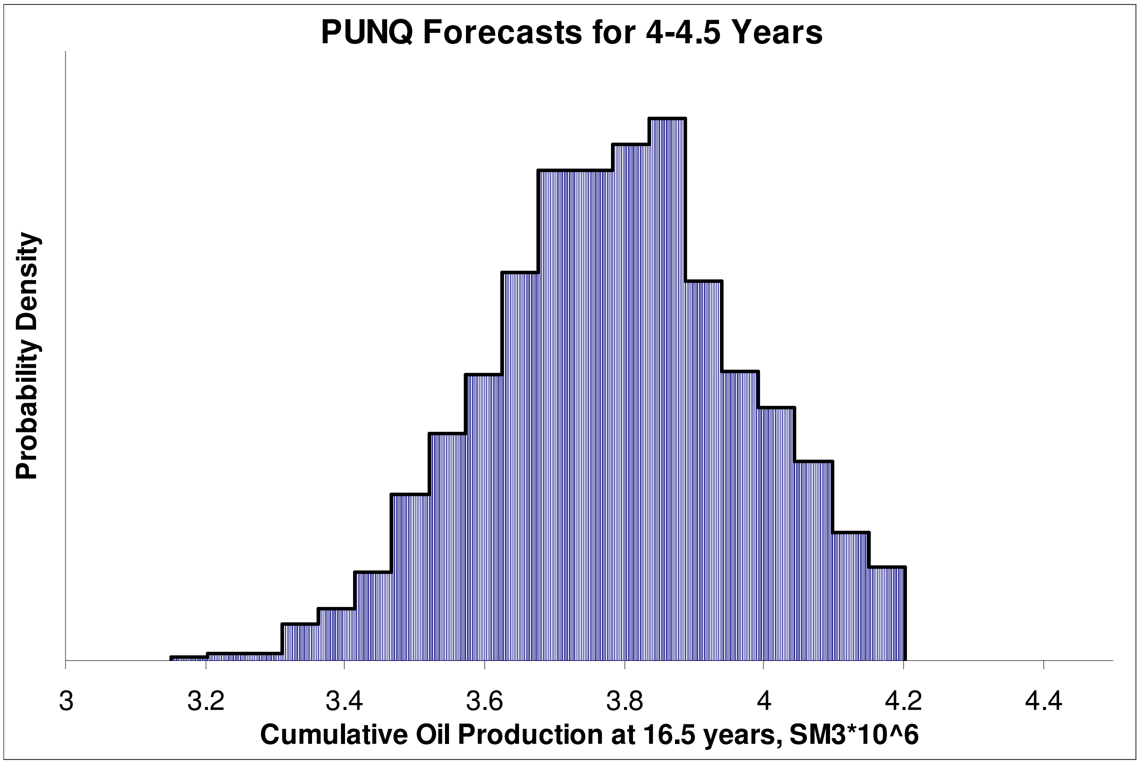

Fig. 23 – Synthetic test forecast – 4 to 4.5 years. The forecast generated from runs 1 – 4500.

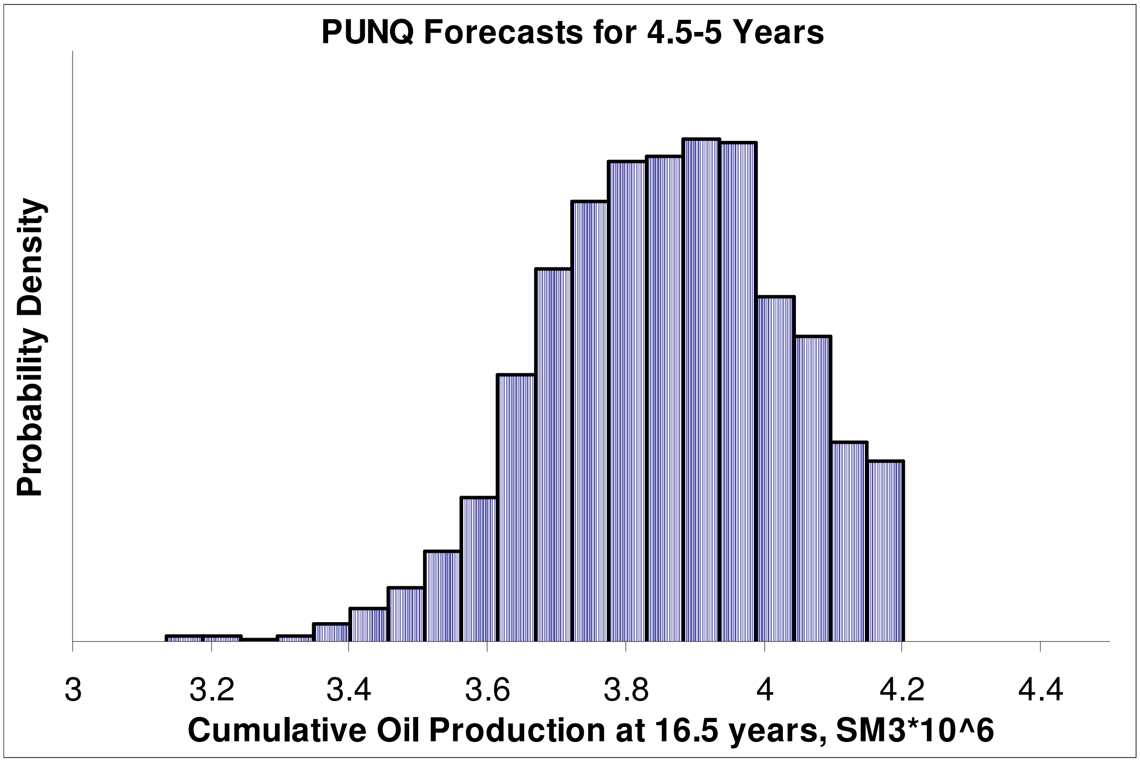

Fig. 24 – Synthetic test forecast – 4.5 to 5 years. The forecast generated from runs 4501-9000.

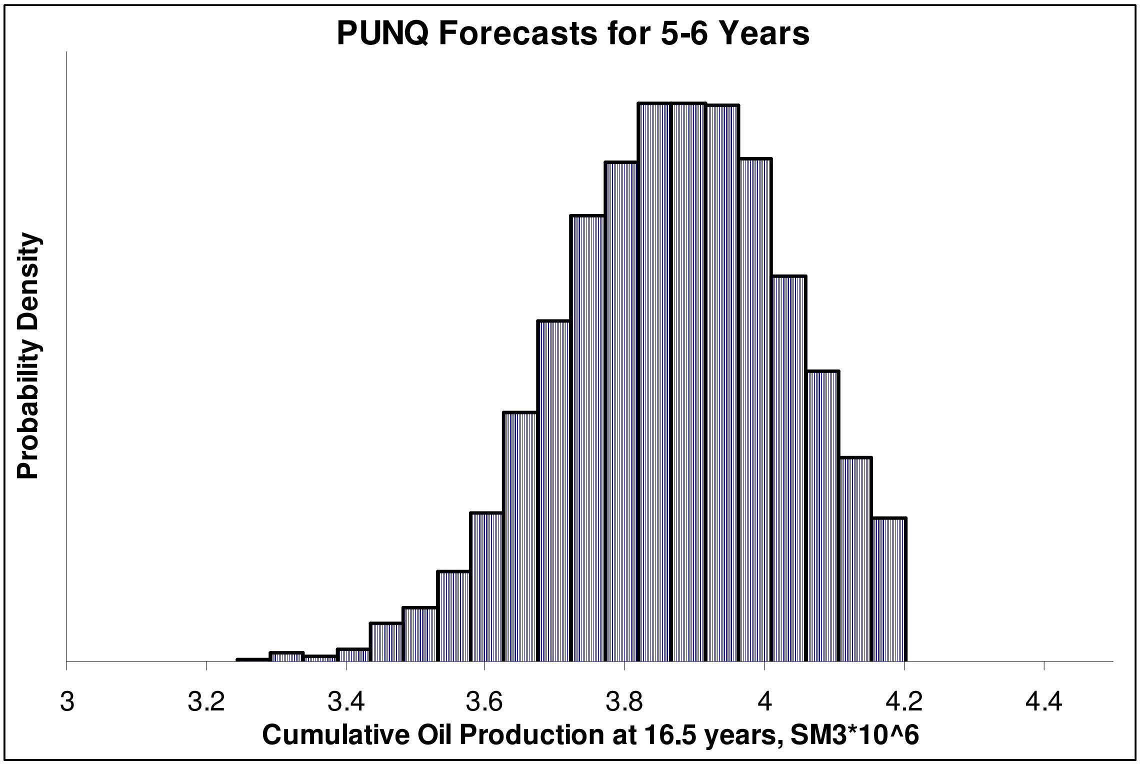

Fig. 25 – Synthetic test forecast – 5 to 6 years. The forecast generated from runs 9001-18000.

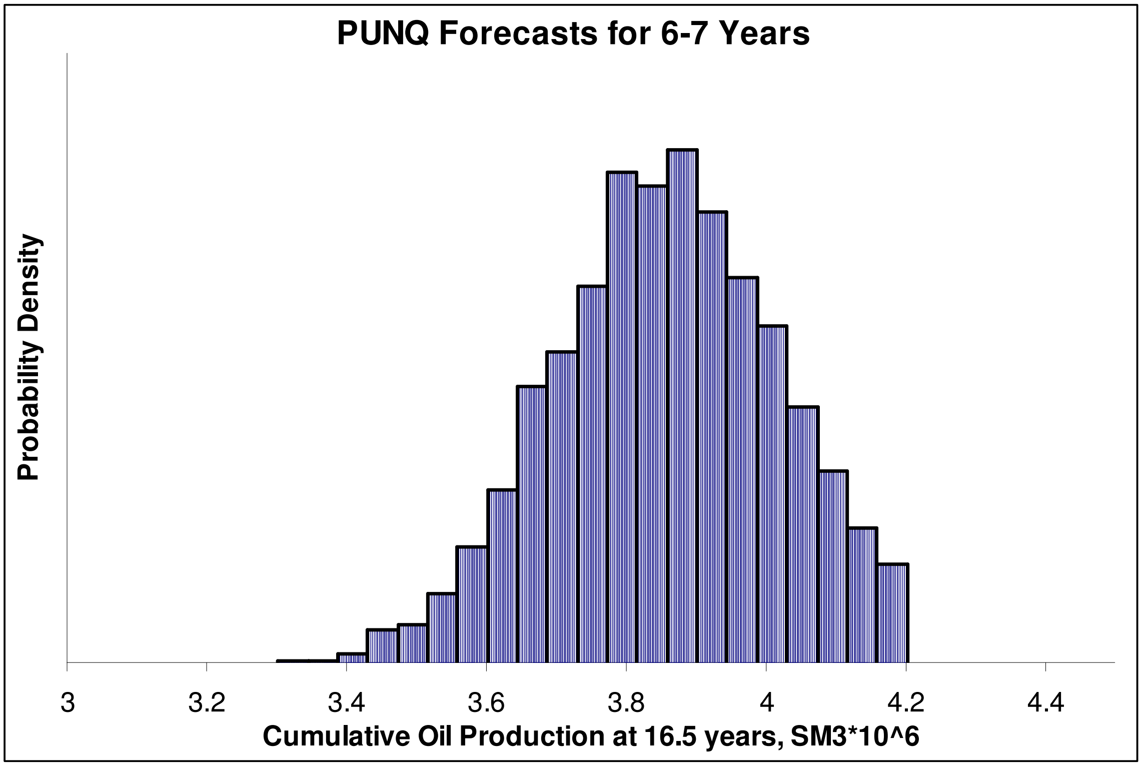

Fig. 26 – Synthetic test forecast – 6 to 7 years. The forecast generated from runs 18001 - 27000.

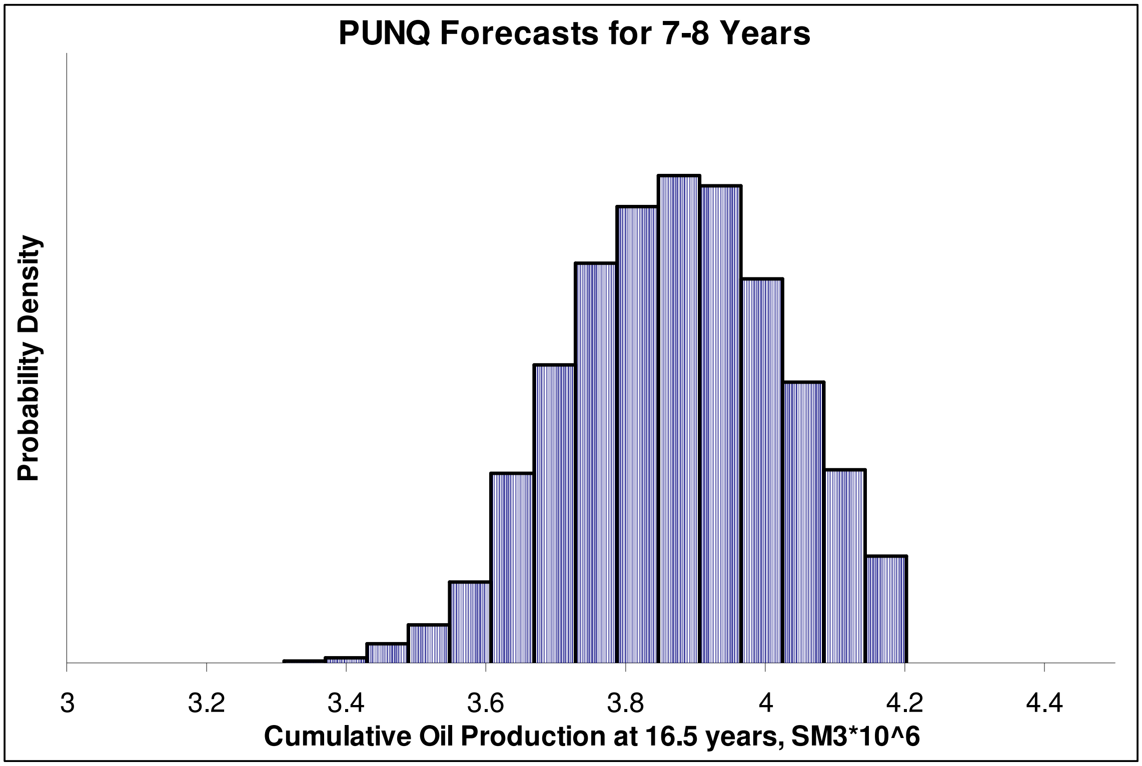

Fig. 27 – Synthetic test forecast – 7 to 8 years. The forecast generated from runs 27001 – 36000.

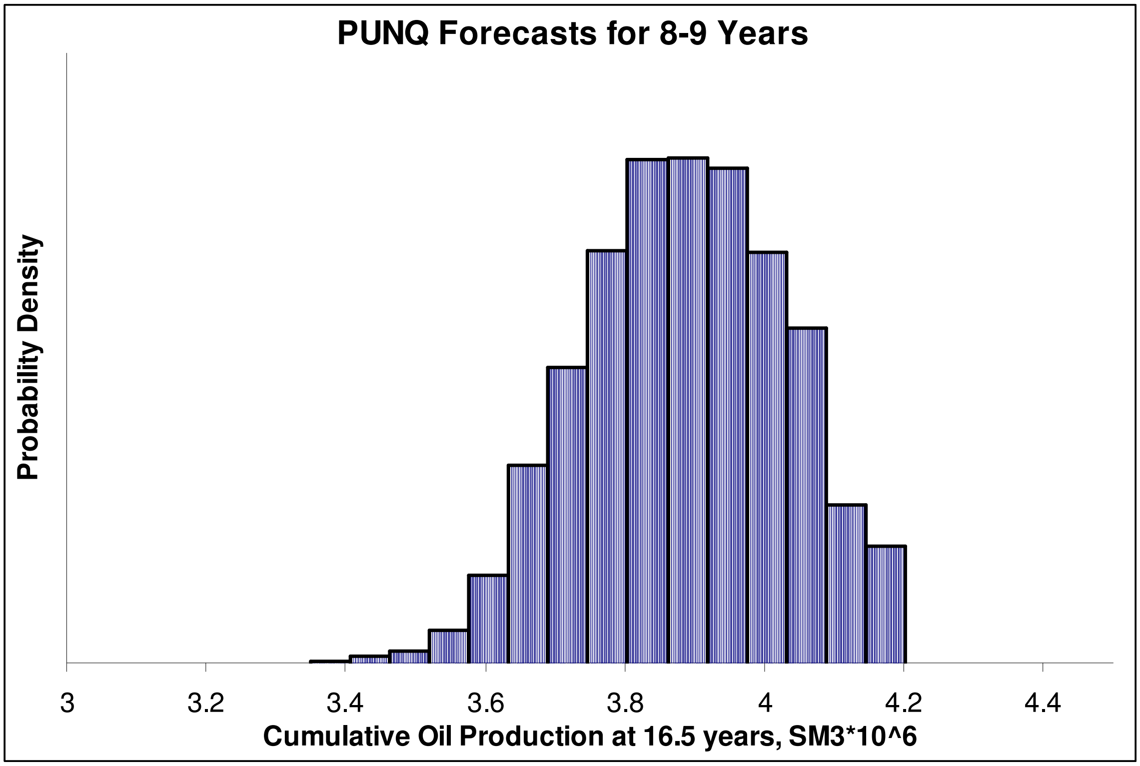

Fig. 28 – Synthetic test forecast – 8 to 9 years. The forecast generated from runs 36001 – 45000.

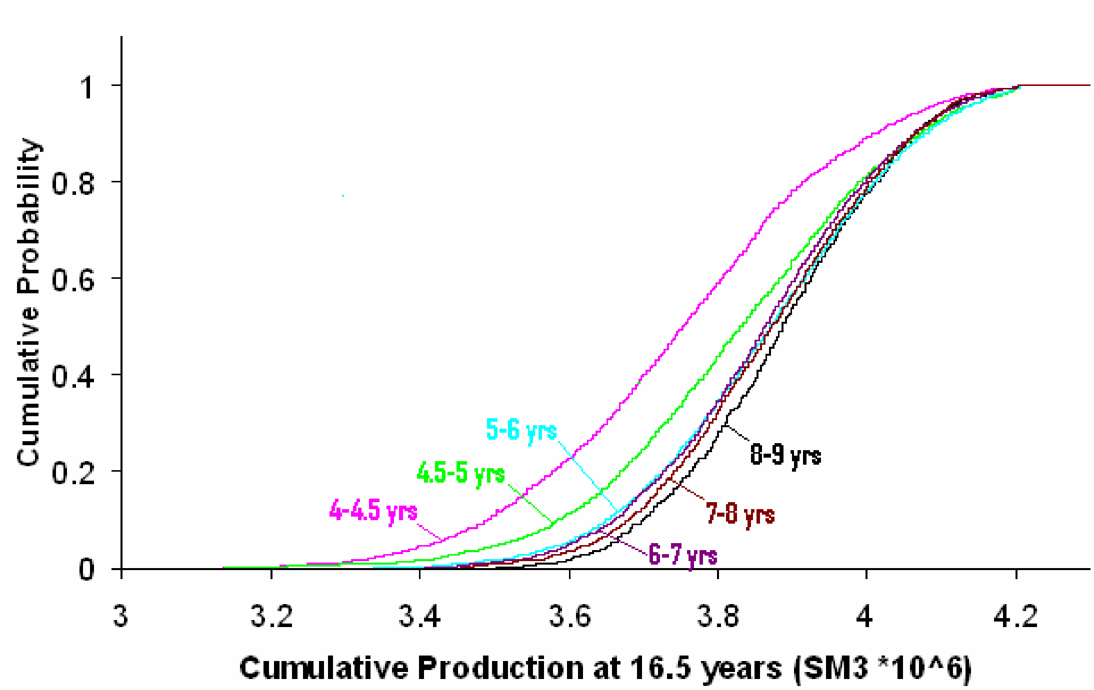

Fig. 29 – Synthetic test forecast CDFs. A comparison of the cumulative distribution functions for the various forecasts made during the synthetic test.

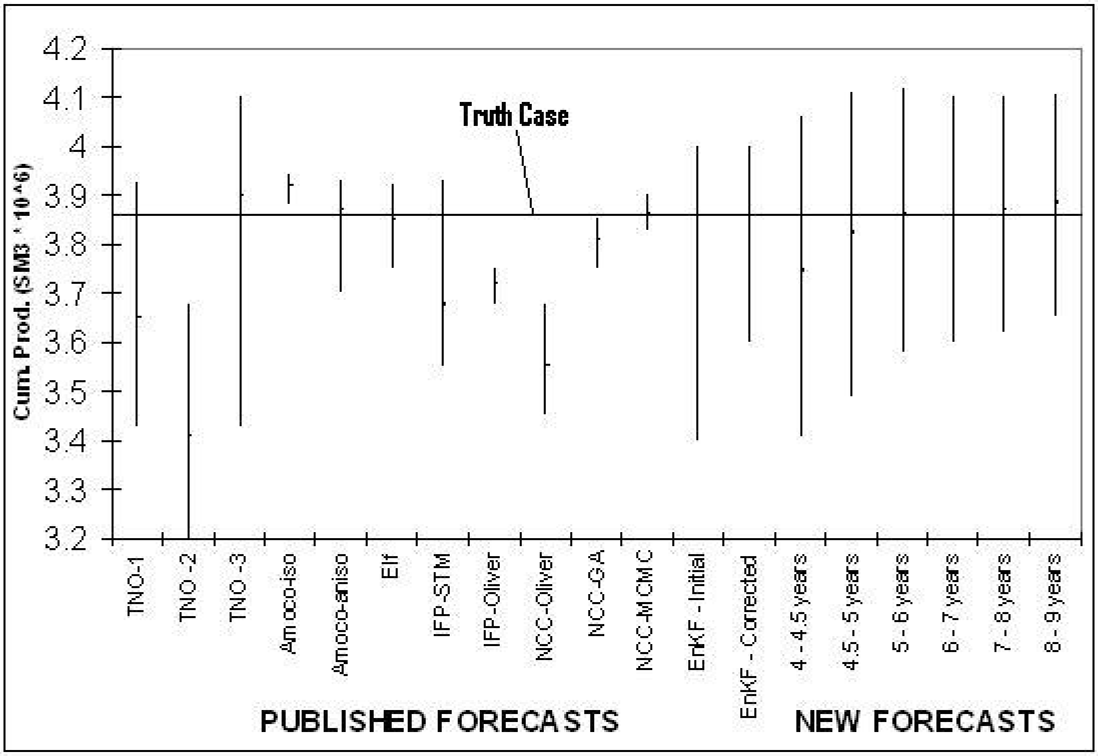

Fig. 30 – Synthetic test forecasts compared. A comparison of forecasts from the synthetic test to published forecast for the PUNQ reservoir.
(Published forecasts after Barker *et al*.[7] and Gu and Oliver[16])

In addition to providing probabilistic production forecasts, it was expected that this process would provide a probabilistic assessment of reservoir properties. Such information could be valuable in routine reservoir management tasks, such as infill drilling. Probabilistic assessments of reservoir properties were generated from the same set of runs used in forecasting. Assessments were created by combining the parameter values into a distribution with equal weighting. In layers and regions in which wells were completed the distributions of parameters varied significantly from the prior parameter distributions. For parameters in places where wells were not completed, however, there seems to be little deviation from the prior distribution. Recalling the prior distribution of the permeability multiplier (**Fig. 19**), we can see an example of this in **Fig. 31,** which shows the prior distribution, the distribution of the horizontal permeability multiplier in layer 4, region 4 where a well is completed and the distribution in layer 1, region 4 where there is not a well. In this figure we see that the revised distribution for layer 4, region 4 deviates significantly from the prior. Meanwhile, the revised assessment of the horizontal permeability multiplier in layer 1, region 4 is quite similar to the underlying prior distribution. This behavior is typical of the other regions in regions in the reservoir. Thus, this process seems to allow us to only narrow our assessments of reservoir properties in certain regions by providing us with revised distributions for our parameters.

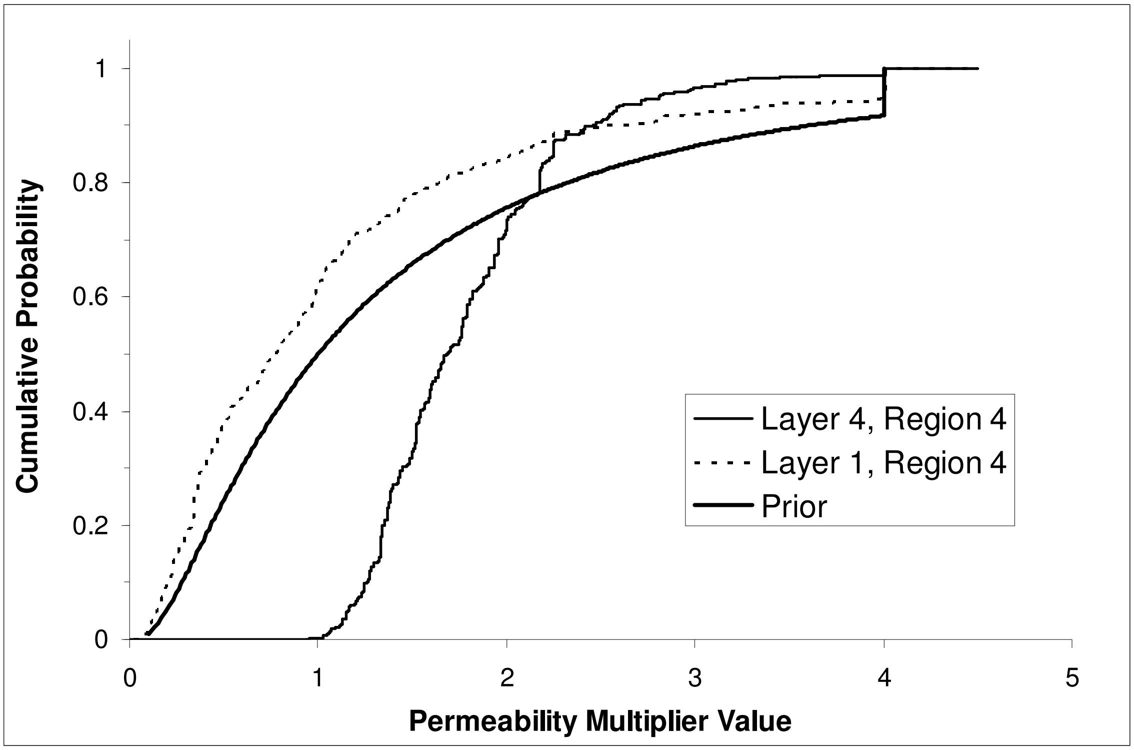

Fig. 31 – Synthetic test revised permeability multiplier assessments.

Summary of Results

This synthetic test demonstrates the value of the continuous simulation process. By making many runs, we are able to consider a wide range of uncertainty in our parameterization. Also, this test demonstrates how we can generate reliable probabilistic forecasts early in the life of the reservoir and narrow our uncertainty ranges over time. In addition, the synthetic test shows that we can make improvements to our assessments of reservoir properties.

LIVE RESERVOIR STUDY

Description of Field and Simulation Model

The second test of the continuous reservoir simulation process is a three month, 77,500 simulation run test conducted on a live field. Contractual obligations restrict the use of the field's real name and it will henceforth be referred to by the pseudonym "Field A." Field A is a mature, layered, domestic tight gas field. The wells in Field A are equipped with real-time monitoring systems that report flow rates and tubing head pressures, amongst other information. Over a three-month period Field A was continuously simulated while receiving daily updates of real-time data.

Field A, the subject of this test, is not a traditional clearly delineated reservoir. Instead it is a subsection of a much larger tight gas field. The edge of the subsection is treated as a no-flow boundary. **Fig. 32** shows a structure map of this subsection. In **Fig. 33** we see the well locations in the simulation model. Within the subsection, as well as in the surrounding areas of the larger field, well spacing and reservoir quality are fairly consistent and there are no known faults or flow barriers. Given these properties there is no reason to expect significant net flow across the subsection boundary. Therefore, it is reasonable to model this subsection in such a way.

Field A was modeled using a 13,824 cell single-phase Eclipse simulation model. This model is laid out as a 6-layer 48 by 48 grid. As mentioned above, Field A is a layered reservoir. There are six major producing layers, all separated by breaks and, to the best of our knowledge, not naturally in communication. There is, however, limited

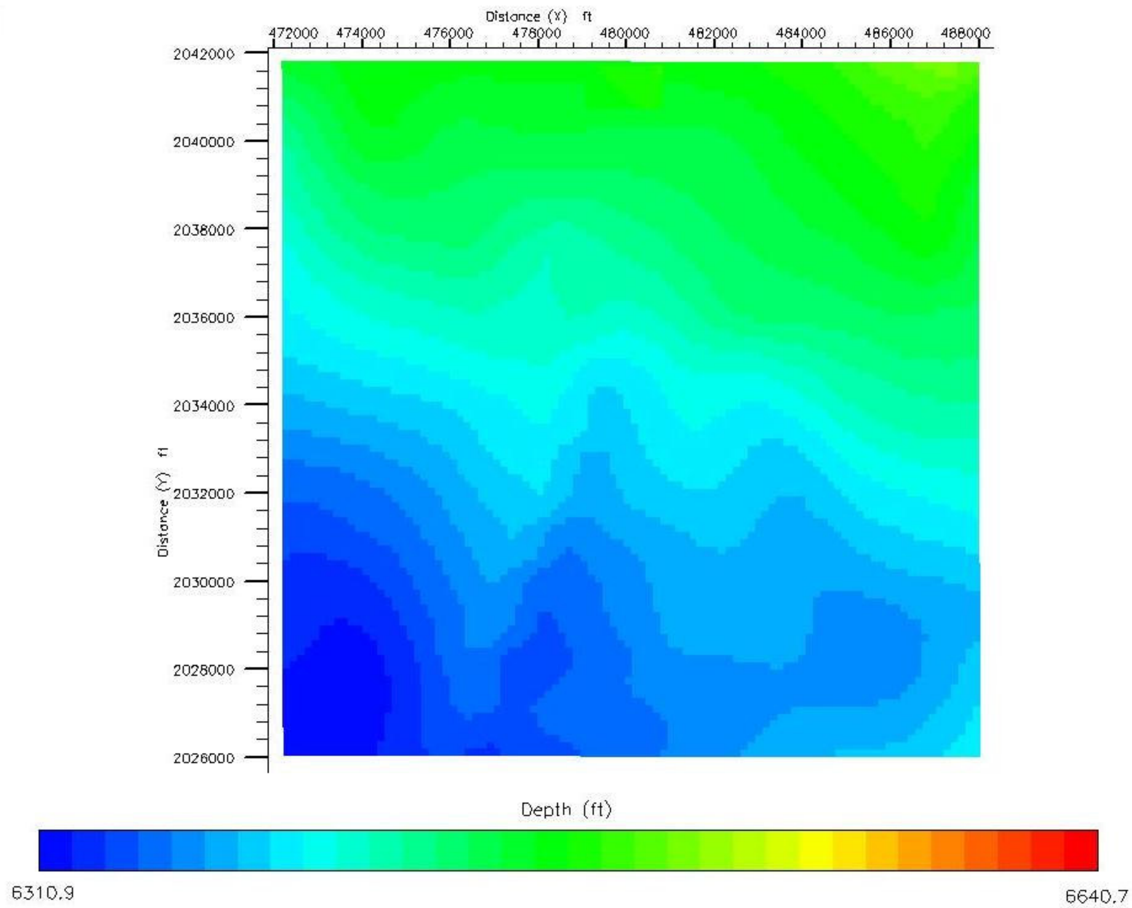

Fig. 32 – Field A structure map.

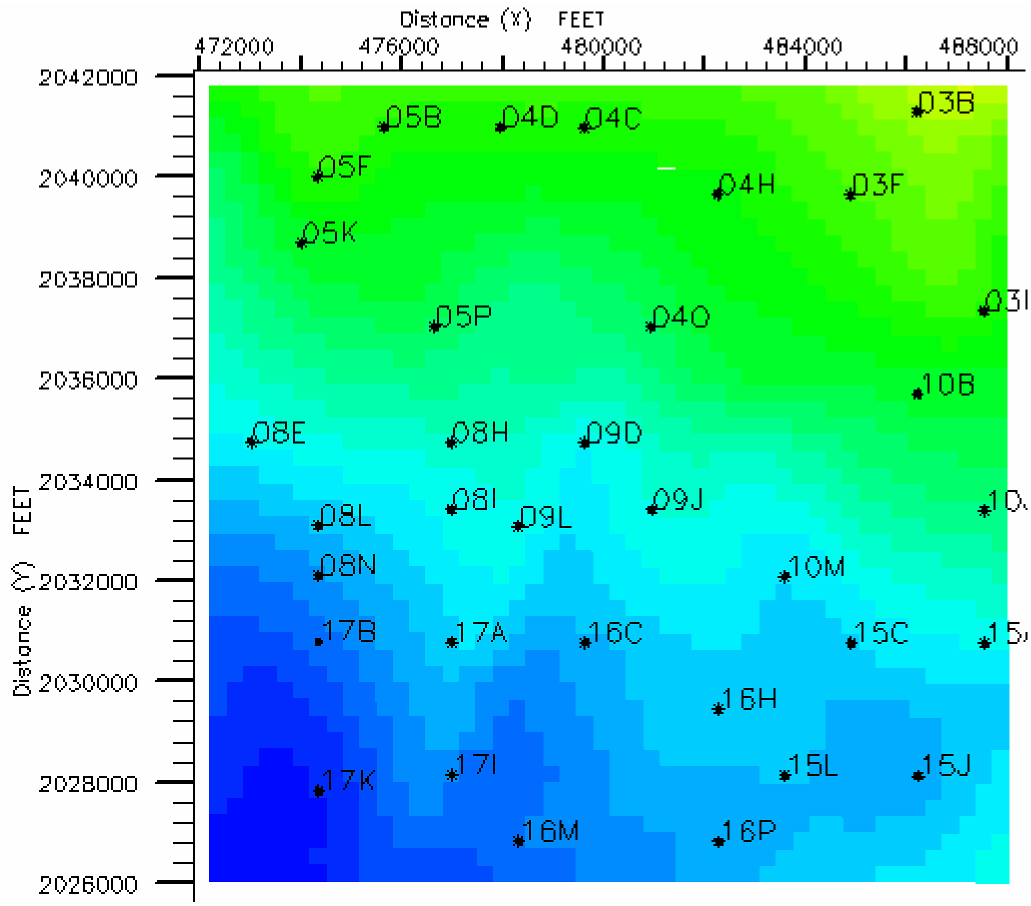

Fig. 33 – Field A well locations. This map shows the location of producing wells in the Field A simulation model.

communication via wellbores and this communication is modeled in the simulator. The layers are shown in the cross-section in **Fig. 34**. From this figure we see that the thickest layers are layers 3 and 6.

Field A is a mature field that first began producing 59 years ago. It was initially developed on 640-acre spacing. Later, well spacing was reduced to 320-acre and then 160-acre spacing. Presently, 80-acre infill wells are being drilled in some parts of the field with plans for additional 80-acre infill wells in the coming years. Thus, despite the field's age, Field A remains attractive for future development and stands to benefit from the use of this system. Cumulative production for the field is shown in **Fig. 35**. We can see that the field has produced approximately 70 million Mscf of gas during its productive life.

The quality and quantity of data available for model construction and history matching is limited due to the age of the field and the several changes in ownership and operators. This lack of data would make conducting a deterministic reservoir study extremely problematic; however, in the context of uncertainty quantification additional uncertainty in the data will simply result in more uncertainty in the forecasts. The geologist at the current operator has mapped both the structure of the field and the thickness of individual layers. These maps were used directly and served as the basis for the simulation model.

Porosity and permeability were not as well defined as very little measured porosity and permeability data exists. Therefore, it is impossible to apply geostatistical techniques to describe the spatial distribution of these properties. As with the synthetic test, constant values of porosity and horizontal permeability, based on the geologist's estimates, were assigned to individual layers. As mentioned above, the individual layers are separated by shale breaks and so vertical permeability is not applicable. The base porosity and horizontal permeability values are given in **Table 3**.

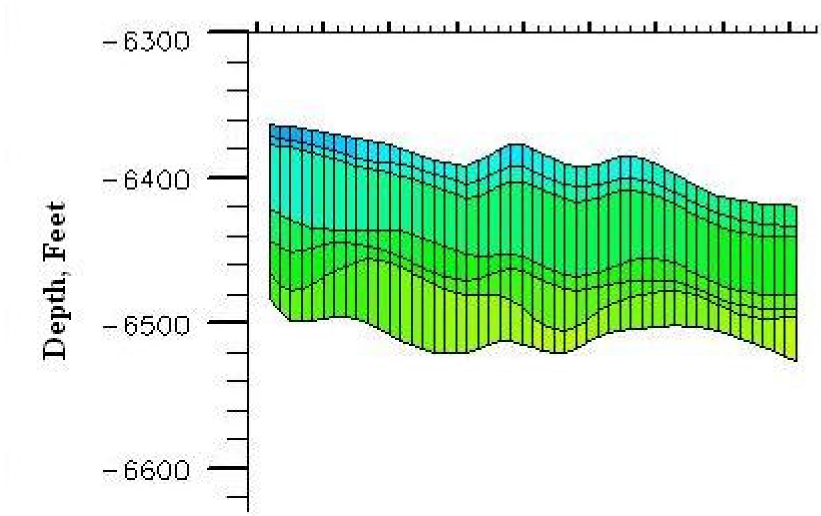

Fig. 34 - Field A cross section. Cross section from Field A showing the six producing layers.

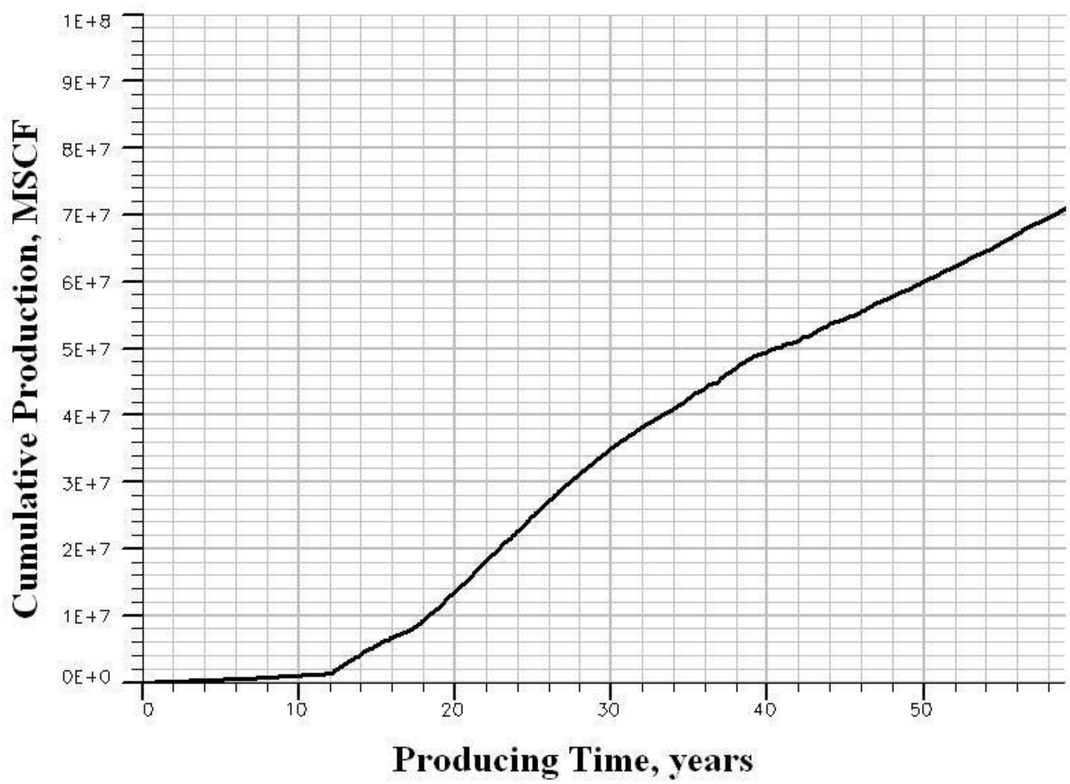

Fig. 35 – Field A historical gas production. This figure shows cumulative gas production to the start of the test of live field test. We see that cumulative production is just above 70 million mscf.

Table 3 - Field Test Base Porosity and Permeability.

Layer	Porosity	Horizontal Permeability (md)
1	5 %	0.7
2	5 %	0.7
3	3 %	0.1
4	5 %	0.7
5	3 %	0.1
6	3 %	0.1

Parameter Space Search

To generate individual models in the history matching and uncertainty quantification process Field A was parameterized much like the synthetic test. The properties adjusted in the GA were porosity and horizontal permeability. To generate an individual model, porosity and permeability multipliers were applied to individual regions within each layer. In this case, nine square regions were used in each of the six layers. These regions are shown on the grid in **Fig. 36**. Two properties times six layers times nine regions gives a total of 108 parameters to be adjusted by the GA. Given the large amount of uncertainty in the values of porosity and permeability, the ranges used for these parameters are accordingly wide. Like in the synthetic test, we should be able to handle this large number of parameters because we are able to make a large number of runs. Prior multipliers for both properties were described using a capped log-normal distribution and are shown in **Figs. 37** and **38**. We note that these distributions are much wider than those used for the synthetic test. These wider multiplier distributions represent the additional uncertainty (due to limited data) in this live field test.

Like the data used in model construction, the data available for history matching are mixed. The objective function contains two types of data: pressure and production. For production data from past years, annual gas production for each individual well is used. For the current, partial year, year-to-date gas production for each well is used. Two types of pressure data were available. First, a number of calculated bottomhole shut-in pressures were available for various wells. These pressure data date from roughly 1970 through the early 1990’s. In computing the objective function these pressures were

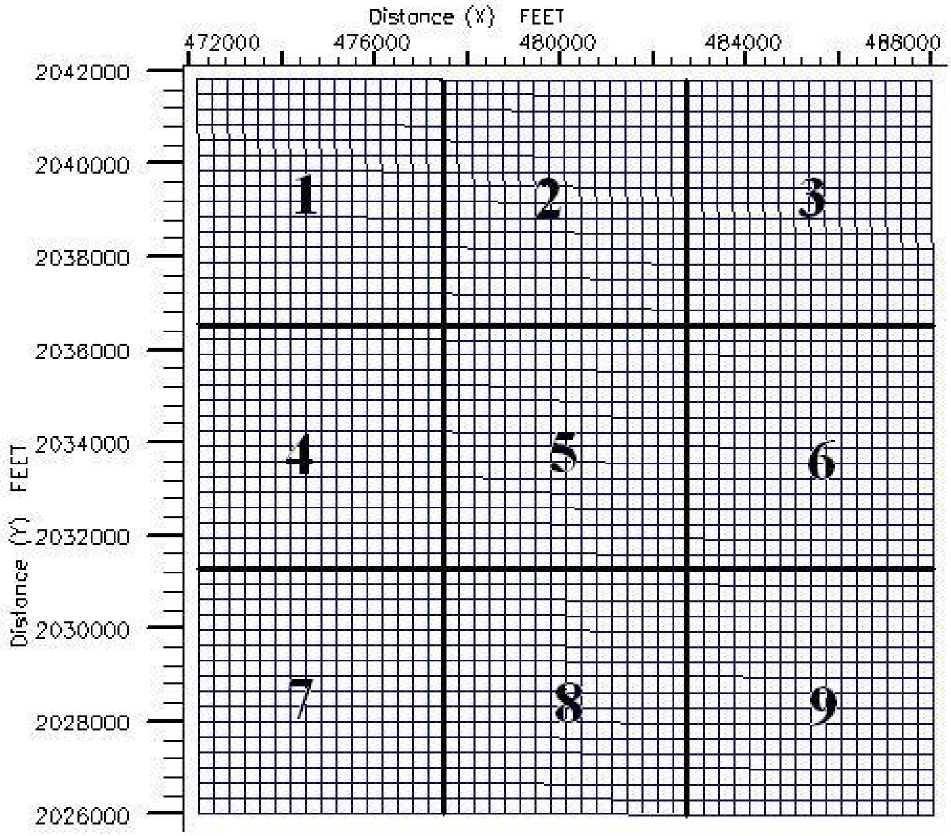

Fig. 36 – Field A multiplier regions. The multiplier regions used to parameterize the live field test.

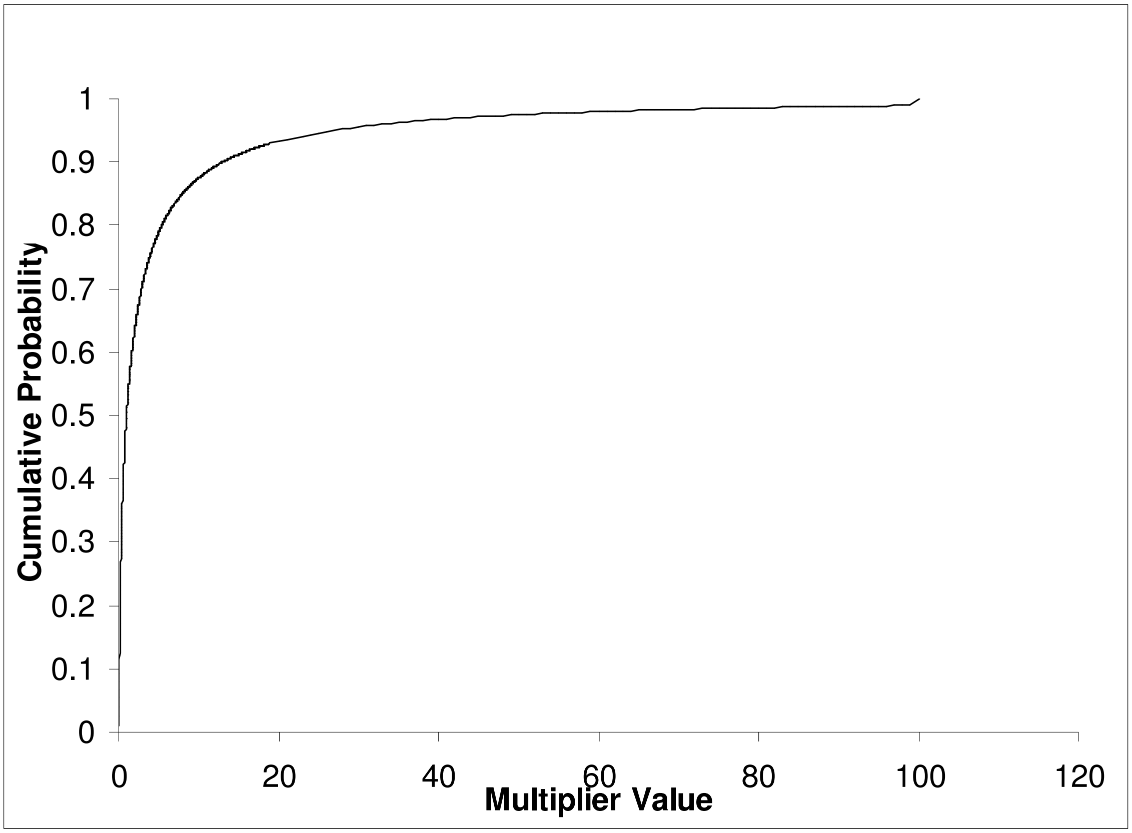

Fig. 37 – Field A permeability multiplier distribution. Prior distribution used for the permeability multiplier parameter in the live field test.

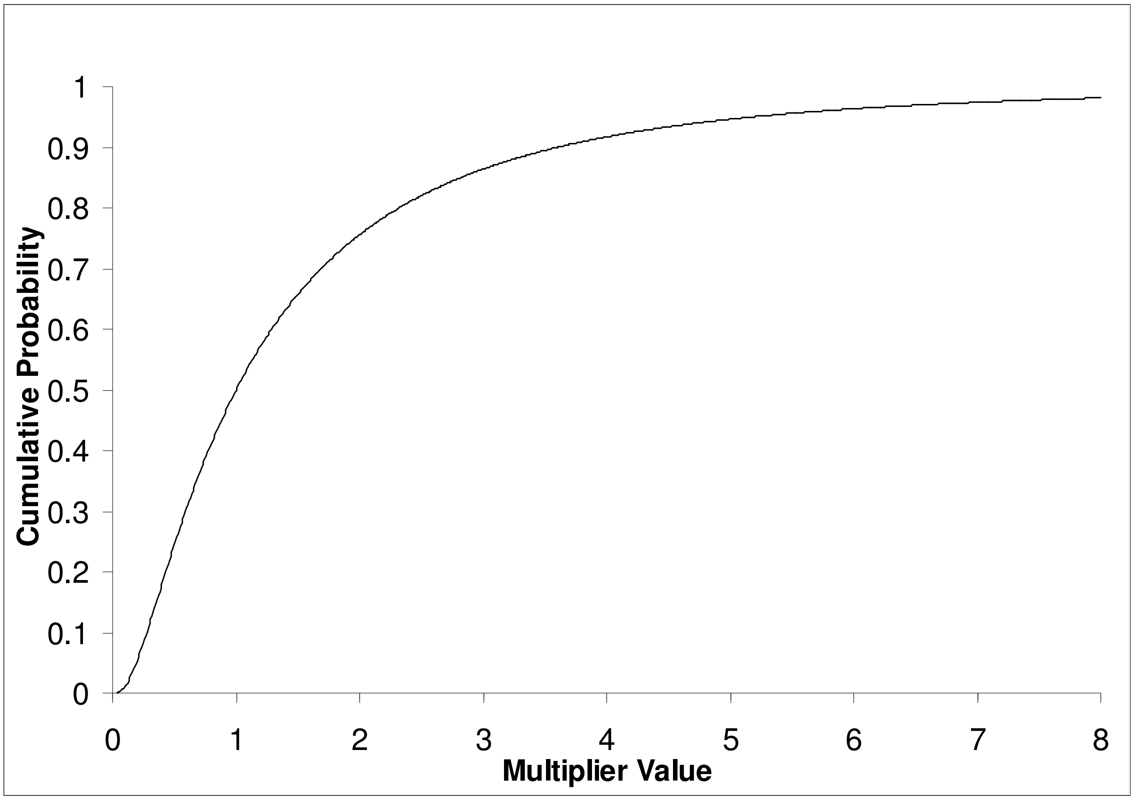

Fig. 38 – Field A porosity multiplier distribution. Prior distribution used for the porosity multiplier parameter in the live field test.

compared to the average simulated pressure of the group of cells in which the well was completed. In order to compare an observed shut-in well pressure to a simulated cell block pressure, the observed pressures were corrected using the correction technique described by Peaceman.[21] Peaceman provides an equation to convert an observed shut-in pressure to a simulation well-block pressure based on reservoir properties, well-block size and shut-in time. The second type of pressure data available is flowing tubing pressures from the real-time data acquisition systems and is described below.

As mentioned previously, all wells in the Field A are equipped with real-time data acquisition systems. These monitoring systems are able to record and report data about five times a second. Constraining and matching a full field simulation model, however, does not require data at this frequency and so averaged daily data are gathered from the field. Of the data received, production rates are the most reliable as they are used for sales contracts. Given their reliability they are updated daily both in the simulation constraints and in the objective function. Adding production data to the constraint file simply involves adding an additional time-step corresponding to a particular day and constraining wells with the actual rate observed on that day. To assess the match to current year production rates, the objective function contains a data point for the year-to-date production in each well. When new real-time data are added, the value of this year-to-date production point in the objective function is incremented to reflect the most recent value of year-to-date production.

Relative to production data, tubing pressures can sometimes be unreliable. For this reason it is useful to have a human review of the pressure data before adding it to the objective function. In addition to ensuring data quality, there are other reasons for not automatically adding new pressure points to the objective function. As mentioned above adding more production data to the objective function simply involves incrementing the year-to-date production data point for a given well. Therefore, adding new production data does not entail adding additional data points to the objective function. When you add new pressure data, however, you are adding additional points to the objective function, which has the potential to significantly shift the magnitude of the objective function for a given model. In the extreme case, too many points for any single type of data will result in reservoir models predominately conditioned to that type of data as opposed to models that match a set of data representative of the field's entire history. While a shift may occur regardless of how and when the pressure data are added, it can be better monitored and managed if the data are added manually and less frequently. In addition, given the low permeability in this field, pressures do not change rapidly and so the pressure one day is highly dependent on the previous day. Although pressure data are available daily, given these concerns it does not make sense to add pressure data every day. Therefore, once a month a tubing pressure data point for each of the fifteen flowing wells is added to the objective function. The system began running without tubing pressure data and the first data was included beginning with the 8000^{th} run (approximately ten days into the test).

The model contains 40 wells, most of which are still producing. Of these producing wells, all but 15 are on artificial lift. Since the pressure data received from the field is tubing head pressure there must be some modeling of wellbore pressure drop in order to utilize this data for history matching. To accomplish this, the simulation model contains a wellbore model which can only model pressure drop in the 15 free flowing wells. Using this wellbore model, calculated tubing head pressures are generated for the appropriate wells by the simulator at times for which observed data are available. These calculated pressures are then compared to the observed pressures in calculating the objective function. The real-time tubing pressure data are particularly valuable given the lack of any recent shut-in pressure data.

Forecasts

History matching against these pressure and production data, Field A was simulated over a period of three months. During this three-month period 77,500 simulation runs were made. These 77,500 runs were divided among 310 Genetic Algorithm generations, each with 250 members. The forecasts for this field were run through January 1, 2026, and all the forecasts reported here were made through that date.

For the purposes of generating forecasts only runs with an acceptable objective function value were considered. The threshold used for this test was 3.0. This objective function threshold was deemed to represent an adequate match to the observed shut-in pressure, production, and tubing pressure data. **Fig. 39** shows the objective function values for the runs included in the forecast, sorted by GA generation number. We can see a discontinuity in these values at the 32nd generation. This discontinuity was the result of introducing additional data, specifically tubing pressures, into the objective function at this point. Because this appears to be a major discontinuity, two sets of forecasts were created: one for the first 32 generations (first 8000 runs) and one for the remaining generations (runs 8001-77,500). Of the first 8000 runs, 3473 runs were below the cutoff and therefore included in the forecast. Obtaining a suitable match became significantly more difficult with the inclusion of tubing pressure data. Of the remaining 69500 runs, only 5311 were below the cutoff and used for forecasting. The equally weighted forecasts for the field test are shown in **Figs. 40-41**. The forecast in **Fig. 41** predicts cumulative production of roughly 80-90 Bscf of gas by 2026. This represents an additional 10-20 Bscf of production beyond the current cumulative of roughly 70 million

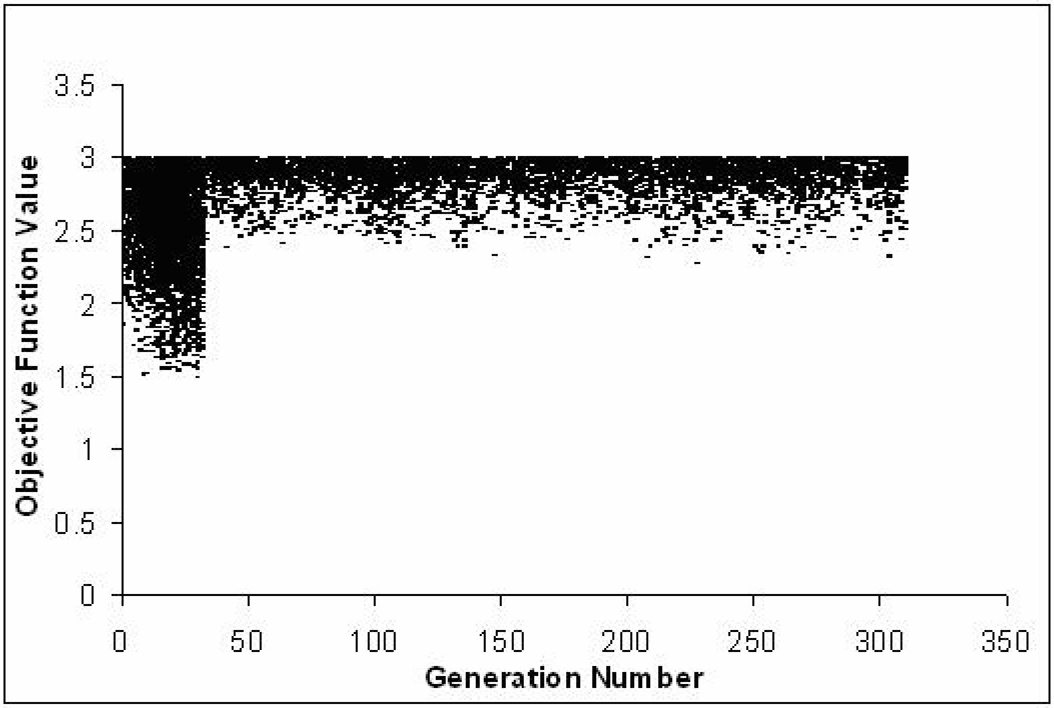

Fig. 39 – Field A objective function values. Live field test objective function values for runs used in forecasts.

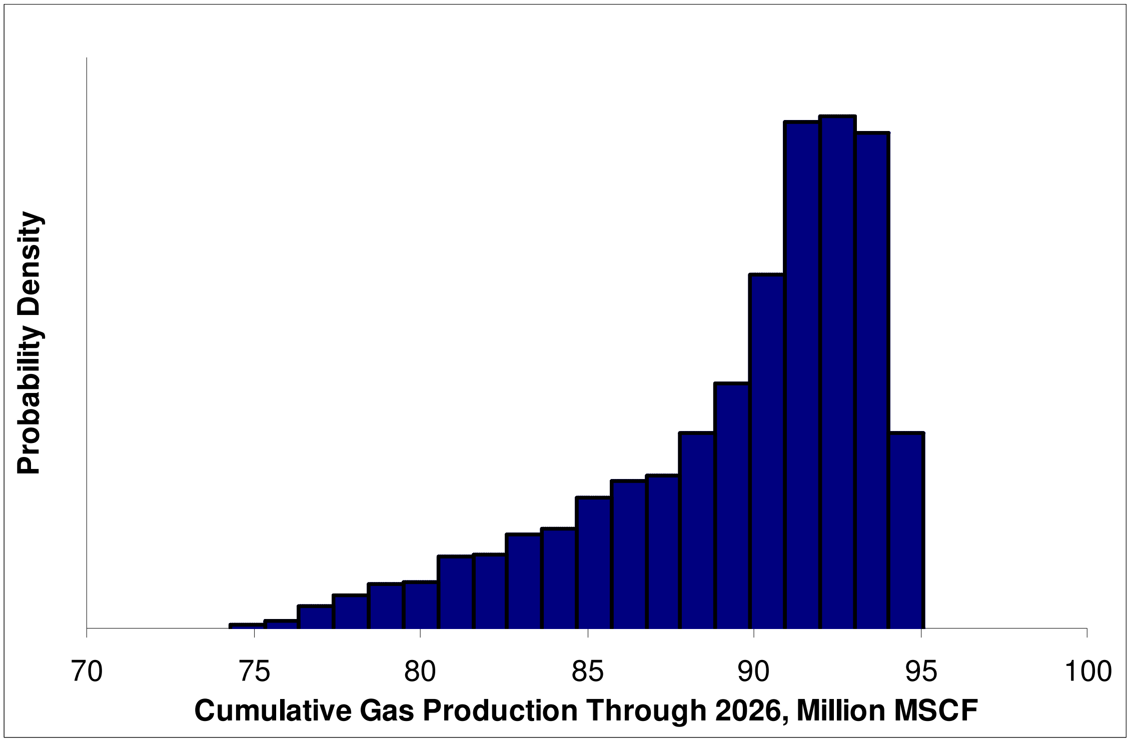

Fig. 40 – Field A forecast without tubing pressure. Live field test forecast from early runs made without tubing pressure data, runs 1 - 8000.

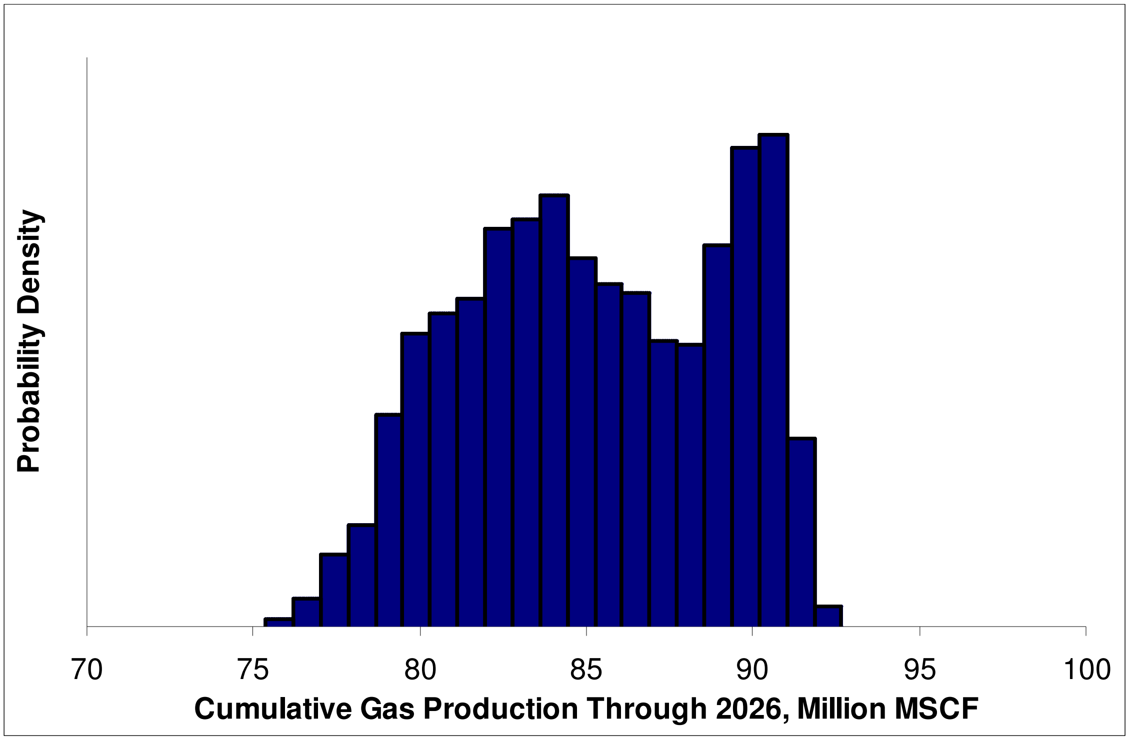

Fig. 41 – Field A forecast with tubing pressure. Live field test forecast from runs made with tubing pressure data, runs 8001 – 77,500.

mscf. We can see that the distribution in the forecast made matching with tubing pressures are significantly narrower than those made without tubing pressures data. I suspect that the addition of tubing pressure eliminated many models with unrealistically high pore volume thus causing this shift.

The forecasts from this live field test differ noticeably from the synthetic test. Most notably we see in **Fig. 39**, with the exception of the discontinuity at the 31st generation, the objective function values are relatively stable with no noticeable narrowing or shifting. This is likely due to the limited time period over which this process was run (3 months for Field A vs. 4 years for the synthetic test) as well as the age of the field when the process was initiated (59 years for Field A vs. 3 years for the synthetic test). At the end of the synthetic test the process had been running for roughly half of the life of the field, compared to less than 0.5% of the life of the field in the live field test. The live field test, however, does demonstrate the feasibility of performing the continuous process on a live field. Furthermore, although forecasts remain stable over time, the live field test has demonstrated that it is possible to use production uncertainty quantification techniques to generate probabilistic forecasts for a real field.

As with the synthetic test, it was hoped that the results of the simulation runs in the live field test could be used to narrow parameter probability distributions. **Fig. 42** shows the prior distribution for the porosity multiplier, the distribution for the porosity multiplier in layer 1, region 1 obtained from acceptable runs, and the distribution for the porosity multiplier in layer 3, region 4 obtained from acceptable runs. We see the distribution for layer 1, region 1 is significantly different from the prior distribution. Thus, this provides us with an updated assessment of porosity in this part of the reservoir. Unfortunately, as with the synthetic test, assessments can be narrowed only in some regions. **Fig. 42** also shows that the revised distribution for the porosity multiplier in layer 3, region 4 barely deviates from the prior. I believe the distribution in layer 3, region 4 remains constant because the observed data are not as sensitive to the reservoir properties in this region. Thus, we are able to revise our assessments of parameters in certain parts of the reservoir though not everywhere.

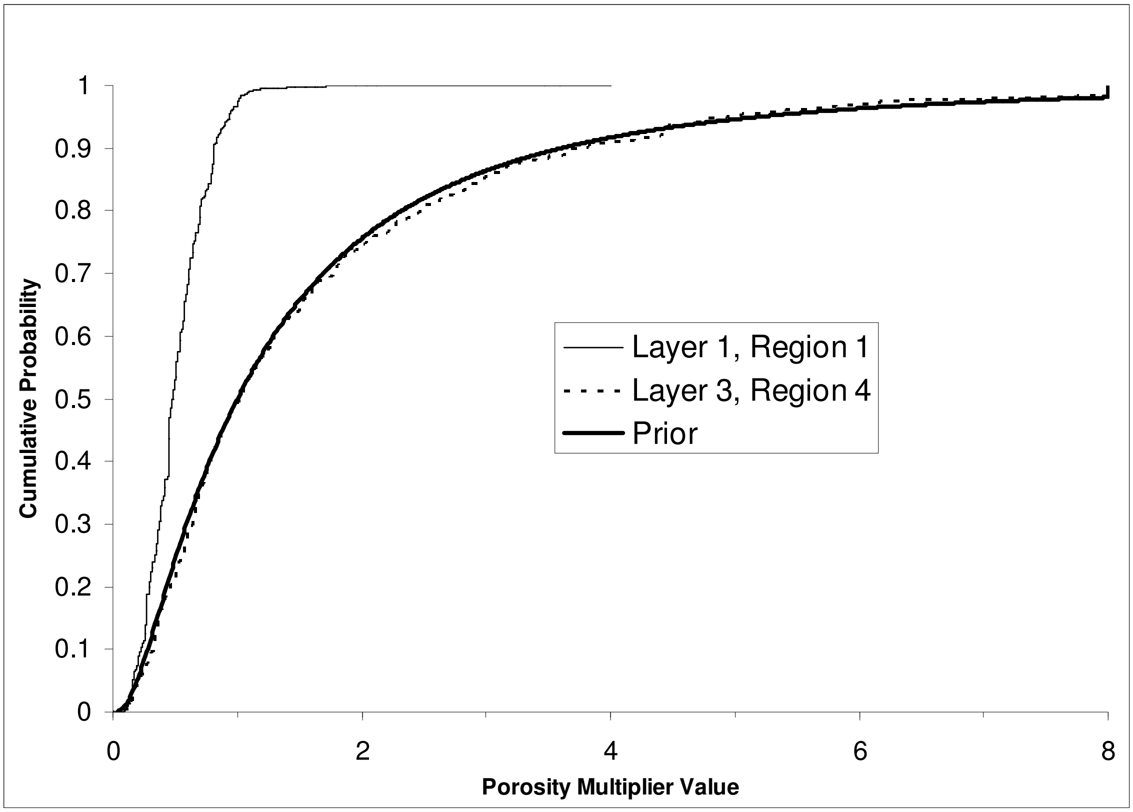

Fig. 42 – Field A revised porosity multiplier assessments.

Summary of Results

This live field test demonstrates that the continuous simulation process is feasible and practical for real reservoirs. We are able to apply this approach in real time to an actual field and generate probabilistic forecasts from the simulation results. Furthermore, like the synthetic test, the simulation results allow us to narrow our assessments of some reservoir properties.

CONCLUSIONS AND RECOMMENDATIONS

Conclusions

Continuous simulation is a promising tool for reservoir management. We can draw the following conclusions from the tests described above:

1. The synthetic test demonstrates that the continuous simulation process can quantify uncertainty in forecasts by making simulation runs throughout the life of a field, thus allowing a more thorough exploration of the parameter space than previously possible. Furthermore, the synthetic test shows that uncertainty quantification in forecasts improves as more data are acquired.

2. The live reservoir test demonstrates that the continuous simulation process is feasible and practical for use on actual fields and could be applied immediately to real world problems.

3. Both tests demonstrate that continuous simulation can allow reservoir engineers to narrow assessments of reservoir properties. These improved assessments of reservoir properties can be utilized in future reservoir management tasks (i.e. infill drilling). Thus, the continuous simulation process provides benefits beyond production forecasts.

Recommendations for Future Work

To the best of my knowledge, the work presented here is the first implementation of this continuous reservoir simulation process. In conducting this study several additional areas for future work were identified, specifically:

1. In this work the search of the parameter space was guided by a Genetic Algorithm. In theory, however, this continuous approach should work with other techniques for searching the parameter space. It would be useful to examine the behavior of the continuous approach using other search techniques, for example the Markov Chain Monte Carlo.
2. So far, the only data incorporated during the continuous process were dynamic data (rates and pressures). This approach should be expanded to handle the inclusion of new static data, as such data will almost certainly be acquired over the life of a reservoir.
3. One of the main premises of this work is that improved uncertainty quantification obtained through the use of the continuous approach should improve the decision making process. It would be useful to test this system in association with specific reservoir management decisions (for instance, drilling a new well).

4. The size (number of grid blocks) of the simulation models used here were fixed for the duration of the continuous simulation process. Over the life of an actual reservoir technology and objectives will change, likely requiring models of varying resolution. Further work is required to investigate the reusability of previous simulation results when switching to more (or less) detailed models.

5. Making simulation runs continuously produces a large quantity of data in the form of simulator output. Approximately 280 gigabytes of simulation results were accumulated during the course of the three-month test on Field A. Techniques for storing and managing these data will be necessary in order to run this process for longer periods of time.

NOMENCLATURE

c = Normalization constant

f = Objective function

f_{prior} = Objective function prior term

L = Likelihood function

P_i = Parameter

y_i^{calc} = Simulated data point

y_i^{obs} = Observed data point

x_i = Experimental data point

μ = Normal distribution mean

μ_p = Expected value of parameter

σ = Standard deviation of normal distribution

σ_i = Standard deviation of error in observed data

σ_p = Standard deviation of parameter distribution

REFERENCES

1. Wiggins, R.L. and Startzman, R.A.: "An Approach to Reservoir Management," Paper SPE 20747 presented at the 1990 SPE Annual Technical Conference and Exhibition, New Orleans, 23-26 September.

2. Thakur, G.C.: "What is Reservoir Management," *JPT* (June 1996) 520.

3. Thomas, G.W.: "The Role of Reservoir Simulation in Optimal Reservoir Management," paper SPE 14129 presented at the 1986 SPE International Meeting on Petroleum Engineering, Beijing, 17-20 March.

4. Howard, R.A.: Decision Analysis, class notes, Department of Management Science and Engineering, Stanford University, Stanford, CA (2005).

5. Capen, E.C.: "The Difficulty of Quantifying Uncertainty," *JPT* (August 1976) 843.

6. Floris, F.J.T. *et al.*: "Methods for Quantifying the Uncertainty of Production Forecasts: A Comparative Study," *Petroleum Geoscience* (2001) **7**, S87.

7. Barker, J.W., Cuypers, M. and Holden, L.: "Quantifying Uncertainty in Production Forecasts: Another Look at the PUNQ-S3 Problem," *SPEJ* (December 2001) 433.

8. Bos, C.F.M. *et al.*: "Production Forecasting with Uncertainty Quantification (PUNQ), Final Report," TNO-report NITG 99-255 A, December 1999

9. Sen, M.K. *et al.*: "Stochastic Reservoir Modeling Using Simulated Annealing and Genetic Algorithms," *SPEFE* (March 1995), 49.

10. Romero, C.E. *et al.*: "A Modified Genetic Algorithm for Reservoir Characterisation," paper SPE 64765 presented at the 2000 SPE International Oil and Gas Conference and Exhibition in China, Beijing, 7-10 November.

11. Goldberg, D.E.: *Genetic Algorithms in Search, Optimization, and Machine Learning,* Addison-Wesley Professional, New York (1989).

12. Raza, S.H.: "Data Acquisition and Analysis for Efficient Reservoir Management," *JPT* (April 1992) 466.

13. Mochizuki, S. *et al*.: "Real Time Optimization: Classification and Assessment," paper SPE 90123 presented at the 2004 SPE Annual Technical Conference and Exhibition, Houston, 25-29 September.

14. Barden, D.: "Coupling Reservoir Simulation Technology with the Internet to Provide Realtime Reservoir Management," paper SPE 66392 presented at the 2001 SPE Reservoir Simulation and Symposium, Houston, 11-14 February.

15. Naevdal, G. *et al*.: "Reservoir Monitoring and Continuous Model Updating Using Ensemble Kalman Filter," paper SPE 84372 presented at the 2003 SPE Annual Technical Conference and Exhibition, Denver, Colorado, 5-8 October.

16. Gu, Y. and Oliver, D.S.: "History Matching of the PUNQ-S3 Reservoir Model Using the Ensemble Kalman Filter," paper SPE 89942 presented at the 2004 SPE Annual Technical Conference and Exhibition, Houston, 26-29 September.

17. "Simulation Software Manuals," 2005A, Schlumberger, Houston (2005).

18. Vose, D.: *Risk Analysis: A Quantitative Guide,* second edition, John Wiley & Sons LTD, New York (2000) 237.

19. "PUNQ-S3," TNO-NITG, www.nitg.tno.nl/punq/cases/ounqs3/, August 21, 2006.

20. Craig, D.P. *et al*.: "Permeability, Pore Pressure, and Leakoff-Type Distributions in Rocky Mountain Basins," *SPEPF* (February 2005) 48.

21. Peaceman, D.W.: "Interpretation of Well-Block Pressures in Numerical Reservoir Simulation," *SPEJ* (June 1978), 183.

VITA

Name:	Jay Cuthbert Holmes
Address:	Harold Vance Department of Petroleum Engineering 3116 TAMU College Station, TX 77843-3116
Education:	B.A., Computer Science, Rice University M.S., Petroleum Engineering, Texas A&M University