Computers & Geosciences 95 (2016) 123–139

Contents lists available at ScienceDirect

Computers & Geosciences

journal homepage: www.elsevier.com/locate/cageo

Research paper

Optimisation of decision making under uncertainty throughout field lifetime: A fractured reservoir example

Dan Arnold [a,*], Vasily Demyanov [a], Mike Christie [a], Alexander Bakay [b], Konstantin Gopa [a]

[a] Institute of Petroleum Engineering, Heriot-Watt University, Edinburgh, UK
[b] Gazpromneft Science & Technology Centre, Saint-Petersburg, Russia

ARTICLE INFO

Article history:
Received 6 August 2015
Received in revised form
20 July 2016
Accepted 21 July 2016
Available online 25 July 2016

Keywords:
Uncertainty quantification
History matching
Fractured reservoirs
Model classification
Optimisation
Field development

ABSTRACT

Assessing the change in uncertainty in reservoir production forecasts over field lifetime is rarely undertaken because of the complexity of joining together the individual workflows. This becomes particularly important in complex fields such as naturally fractured reservoirs. The impact of this problem has been identified in previous and many solutions have been proposed but never implemented on complex reservoir problems due to the computational cost of quantifying uncertainty and optimising the reservoir development, specifically knowing how many and what kind of simulations to run.

This paper demonstrates a workflow that propagates uncertainty throughout field lifetime, and into the decision making process by a combination of a metric-based approach, multi-objective optimisation and Bayesian estimation of uncertainty. The workflow propagates uncertainty estimates from appraisal into initial development optimisation, then updates uncertainty through history matching and finally propagates it into late-life optimisation. The combination of techniques applied, namely the metric approach and multi-objective optimisation, help evaluate development options under uncertainty. This was achieved with a significantly reduced number of flow simulations, such that the combined workflow is computationally feasible to run for a real-field problem.

This workflow is applied to two synthetic naturally fractured reservoir (NFR) case studies in appraisal, field development, history matching and mid-life EOR stages. The first is a simple sector model, while the second is a more complex full field example based on a real life analogue. This study infers geological uncertainty from an ensemble of models that are based on the carbonate Brazilian outcrop which are propagated through the field lifetime, before and after the start of production, with the inclusion of production data significantly collapsing the spread of P10-P90 in reservoir forecasts. The workflow links uncertainty estimation with the appropriate optimisation at appraisal, development and reservoir management stages to maximise oil recovery under uncertainty.

1. Introduction

A key challenge in reservoir simulation is to know the appropriate number of models to run in order to make a good decision given estimates of uncertainty in reservoir description. Much work has been published on topics such as history matching, uncertainty quantification, sensitivity analysis and optimisation to tackle this challenge across 4 key stages for reservoir modelling & development:

1. Reservoir appraisal: Sensitivity analysis is (commonly) used to tell us the preproduction estimate of uncertainty.
2. Initial reservoir development planning: The best development options are explored to maximise value given the uncertainty.
3. Reservoir model history matching: Production data is integrated to improve the estimates of uncertainty.
4. Mid-late life Reservoir management: Additional mid/late life development decisions are optimised given the uncertainty, which has been updated from history matching.

Stages 1 and 3 are concerned with reservoir forecasting given the uncertainty in the subsurface model. Stages 2 and 4 are concerned with identifying new development opportunities to maximise reservoir value given a range of different options, well locations and other engineering trade-offs (e.g. maximising oil vs minimising water). The stages represent a set of interconnected ill-posed inverse problems, where many possible models/solutions may exist given the data yet provide a range of forecasts.

Stochastic optimisation methods are often employed to solve the inverse problem in reservoir simulation, and Bayesian

* Corresponding author.
E-mail address: dan.arnold@pet.hw.ac.uk (D. Arnold).

http://dx.doi.org/10.1016/j.cageo.2016.07.011
0098-3004/

formalism can be applied to the results of history matching to estimate the uncertainty. Alternatively, Park et al. (2013) for instance estimated uncertainty using kernel probability density estimation, a different formalism to the Bayesian way we employ in this paper.

Uncertainty analysis is rarely however propagated from one stage to another in a rigorous yet practical way which is the aim of the workflow described in this paper. The aim is to create an improved workflow that would fully couple stages 1–4 and propagate the uncertainty through each stage of forecasting and optimisation. The reservoir development will be optimised across the full range of possible reservoir descriptions, given the considered uncertainty, to provide the best development option (or set of options to choose from). This could be described as optimisation under uncertainty.

One of the most statistically robust approaches to estimating the uncertainty in an inverse problem is to apply Monte Carlo (MC) techniques to infer unknown distribution of the uncertain model parameters. Bayesian approaches to uncertainty quantification often use Markov Chain Monte Carlo (MCMC) to quantify uncertainty. In all cases (particularly for MCMC as many samples are rejected in the process of running the algorithm) Monte Carlo methods are computationally expensive with many hundreds of thousands or millions of iterations being required for convergence in high dimensional problems.

Building a complete field lifetime workflow using these MCMC techniques aiming to propagate the uncertainty through the workflow would create an unmanageable computational cost, unless surrogate models are used, however these trade-off speed of solution against additional errors in uncertainty quantification. The resulting combinatorial explosion in, for instance, finding the optimal solution for drilling a new well given a large ensemble of possible reservoir descriptions could easily lead to millions of model runs – an unfeasibly large number for practical engineering problems.

This work also aims to reduce the computational workload of estimating and propagating uncertainties effectively to improve on the industry norm. The reality for many practicing engineers, given typical constraints on time and the computer power, is something equivalent to the following steps:

- Simulations are run to assess the uncertainty at the initial stage, when the uncertainty is theoretically at its greatest. These simulations are more commonly developed as Min–Max ranges rather than statistically significant confidence intervals.
- One single “most likely”/“base case” is carried forward for development planning and optimisation at the initial stage. Sometimes the min and max cases are tested against the optimal development plan for the “base case”.
- The base case model is compared to actual production data and history matched to the production data where possible. The history matching process can often involve adding parameters such as multipliers to the simulation model in order to add flexibility to the model and make matching easier. Where the base case cannot match, another model from the appraisal stage or a newly generated model may be used to improve the history match.
- The best history matched model is graduated to the next stage and used in subsequent reservoir development decisions, history matching and optimisation phases, being the “most likely” model.

The problem with the above approach is that at each stage, the uncertainty is not estimated with any reasonable level of statistical robustness (such as you would get from MCMC) and the uncertainty estimates are not propagated into the optimisation stages (Stages 2 and 4). Therefore, the “optimum” is only the optimal for the most likely reservoir description, rather than in respect to the reservoir uncertainty.

In this paper we propose a solution to the problem of propagating estimates of uncertainty throughout field lifetime with a minimum number of simulation runs to provide accurate estimates of uncertainty. This is attempted by developing a full field lifetime workflow that:

- propagates the uncertainty from one stage to the next such that we optimise decisions given the level of uncertainty and no information is lost between each stage;
- calculates and updates statistical estimates of uncertainty at all stages;
- is computationally efficient in minimising the number of stochastic iterations required to assess uncertainty and optimise the development.

The objective is to create a complete, field lifetime workflow for uncertainty propagation that is efficient enough to be attempted on a real field scenario. Recent papers, such as Shirangi and Durlofsky (2015), demonstrate approaches for this type of (Closed-loop) workflow where uncertainty analysis and optimisation are coupled but required large numbers of iterations in to achieve a result (220,000 flow simulations) thus reducing the simulation model requirement would be an advantage for more complex field problems. Maucec et al. (2011) used a Bayesian MCMC approach on a proxy model to estimate geological uncertainties then MDS to dynamically rank the realisations, however they did not carry these uncertainties into an optimisation step thus well covered steps 1 and 3.

This paper describes a novel combination of existing techniques to facilitate an optimisation under uncertainty process (carried out both before and after history matching) to maximise the value of the field throughout its lifetime.

The first section of this paper describes the set of techniques employed in this new workflow that enables accurate estimation of the uncertainty with a minimal cost in terms of simulation time. This is a combination of three key technologies: Multi-Objective optimisation, posterior NA-Bayes inference and model clustering/classification using Multi-dimensional scaling. These techniques are used in different combinations at each stage to best capture the uncertainty while minimising the computational cost. There is a brief description of each technique used.

The second section of the paper demonstrates variations of the new workflow on 2 synthetic case study examples, (1) a simple sector models and (2) a more complex fractured field example. The second case study is a relatively complex but realistic (70,740 cells) synthetic fractured field example developed using fracture analogue data from Van Eijk (2014), to create a set of Discrete Fracture Network (DFN) models, which were then upscaled to create representative simulation grids. The field is 3 phase, with a gas cap and a relatively thin oil bearing section, making optimisation of well placement a challenge given the quantified range of uncertainty and model complexity.

To be consistent throughout the paper, a predefined set of nomenclature is used to explain important terminology around modelling. These terms are described as follows:

- Scenario – a high level geological idea that is to be tested. Different geological scenarios describe differences in the geological conceptual model.
- Realisation – this is the outcome from a stochastic process (which can be either stochastic simulation of the static model or the models produced from stochastic optimisation (history matching or optimisation). We can produce many realisations of a scenario