Ensemble-based multiobjective optimization of on/off control devices under geological uncertainty

Fonseca, R. M.; Leeuwenburgh, O.; Della Rossa, E; Van Den Hof, P. M.J.; Jansen, J. D.

Publication date

2015

Document Version

Final published version

Published in

SPE Reservoir Evaluation and Engineering

Citation (APA)

Fonseca, R. M., Leeuwenburgh, O., Della Rossa, E., Van Den Hof, P. M. J., & Jansen, J. D. (2015). Ensemble-based multiobjective optimization of on/off control devices under geological uncertainty. In *SPE Reservoir Evaluation and Engineering* (Vol. 18, pp. 554-563). (SPE Reservoir Evaluation and Engineering; Vol. 18, No. 4). Society of Petroleum Engineers.

Important note

To cite this publication, please use the final published version (if applicable).
Please check the document version above.

Copyright

Other than for strictly personal use, it is not permitted to download, forward or distribute the text or part of it, without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license such as Creative Commons.

Takedown policy

Please contact us and provide details if you believe this document breaches copyrights.
We will remove access to the work immediately and investigate your claim.

Green Open Access added to TU Delft Institutional Repository

'You share, we take care!' – Taverne project

<https://www.openaccess.nl/en/you-share-we-take-care>

Otherwise as indicated in the copyright section: the publisher is the copyright holder of this work and the author uses the Dutch legislation to make this work public.

Ensemble-Based Multiobjective Optimization of On/Off Control Devices Under Geological Uncertainty

R.-M. Fonseca, Delft University of Technology; O. Leeuwenburgh, TNO; E. Della Rossa, Eni;
P. M. J. Van den Hof, Eindhoven University of Technology, and J.-D. Jansen, Delft University of Technology

Summary

We consider robust ensemble-based (EnOpt) multiobjective production optimization of on/off inflow-control devices (ICDs) for a sector model inspired by a real-field case. The use of on/off valves as optimization variables leads to a discrete control problem. We propose a reparameterization of such discrete controls in terms of switching times (i.e., we optimize the time at which a particular valve is either open or closed). This transforms the discrete control problem into a continuous control problem that can be efficiently handled with the EnOpt method. In addition, this leads to a significant reduction in the number of controls that is expected to be beneficial for gradient quality when using approximate gradients. We consider an ensemble of sector models where the uncertainty is described by different permeability, porosity, net/gross ratios, and initial water-saturation fields. The controls are the ICD settings over time in the three horizontal injection wells, with approximately 15 ICDs per well. Different optimized strategies resulting from different initial strategies were compared. We achieved a mean 4.2% increase in expected net present value (NPV) at a 10% discount rate compared with a traditional pressure-maintenance strategy. Next, we performed a sequential bi-objective optimization and achieved an increase of 9.2% in the secondary objective (25% discounted NPV to emphasize short-term production gains) for a minimal decrease of 1% in the primary objective (0% discounted NPV to emphasize long-term recovery gains), as averaged over the 100 geological realizations. The work flow was repeated for alternative numbers of ICDs, showing that having fewer control options lowers the expected value for this particular case. The results demonstrate that ensemble-based optimization work flows are able to produce improved robust recovery strategies for realistic field-sector models against acceptable computational cost.

Introduction

Recently, there has been an increased focus on the application of different model-based optimization techniques for optimal control to achieve improved reservoir-management strategies. Most of these studies have used relatively simple models, whereas a limited number of studies such as Bailey et al. (2005), Sarma et al. (2008), Alhuthali et al. (2009), Chaudhri et al. (2009), Forouzanfar et al. (2013), van Essen et al. (2010), and Raniolo et al. (2013) have used realistic field-scale or sector models. Most of the studies that use realistic real field models concerned single-objective optimization on a single geological realization, with the exception of Alhuthali et al. (2009) and Raniolo et al. (2013), who performed single-objective optimization by use of an ensemble of geological realizations. In the present study we use a sector model inspired from a real field case.

There are two reasons for suboptimality of strategies resulting from the use of a single-objective, single-model optimization.

First, the geomodeling process used to create the dynamic models is fraught with uncertainties in data and interpretations. Because the true geological description of a reservoir is never known with certainty, it is imperative to include geological uncertainty within the optimization framework. Second, strategies resulting from a single (long-term) objective are often in conflict with short-term objectives. The short-term objectives (on a time scale of days to months) are driven by operational criteria, facility constraints, production-sharing contracts, and so forth. Thus, including the short-term goals within the optimized operational strategy is also important to correct for compromises made during the single-objective (long-term) optimization. Finding strategies that account for multiple objectives may achieve solutions that are more acceptable to a decision maker.

Jansen et al. (2009), among others, have shown that significantly different operational strategies can lead to very-similar objective-function values, which was reasoned to be the result of the ill-posedness of the problem resulting in nonuniqueness of the solution. A similar nonuniqueness in minimizing the mismatch between measured and simulated data during computer-assisted history matching was demonstrated by Oliver et al. (2008). van Essen et al. (2011) proposed a hierarchical framework to exploit these redundant degrees of freedom in the control space to solve a multiobjective optimization problem. One of the approaches they introduced uses a switching function to optimize multiple objectives, which will be used in this study.

Yeten et al. (2003) described an approach to account for geological uncertainty during well-location optimization with the aid of multiple models. van Essen et al. (2009, 2011) used the adjoint formulation for gradient-based optimization while incorporating geological uncertainty into the optimization framework; Jansen (2011) provides an overview of the adjoint formulation and further references. The adjoint approach is computationally very efficient, because it requires only one forward and one adjoint simulation to compute the gradient, irrespective of the number of controls. However, implementation of the adjoint formulation is tedious and requires access to the simulation source code, which is the method's biggest drawback because this is not possible when working with commercial simulators. To overcome this difficulty, there has been an increase in the application of stochastic gradient-based techniques. One such stochastic technique, ensemble-based optimization (EnOpt), introduced by Chen et al. (2009), has proved to be a good practical alternative to the adjoint because it is relatively easy to implement and is independent of the reservoir simulator, albeit computationally inferior. Chen (2008) and Chen and Oliver (2010), starting from certain theoretical assumptions, proposed a framework for robust optimization with EnOpt that is computationally attractive. Chen et al. (2009) reported a successful application of this approach to the SPE Brugge benchmark case, whereas Raniolo et al. (2013) and Li et al. (2012) have investigated the applicability of approximate gradient techniques for life-cycle robust waterflooding optimization. Recently, Fonseca et al. (2014) suggested a modified robust gradient formulation and also investigated the applicability of EnOpt for robust multiobjective optimization on a synthetic test case. They showed that EnOpt is a computationally competitive alternative when the adjoint is unavailable, especially for robust

optimization. In Appendix A we give an overview of the essential theory of EnOpt as used in our study.

In this study we optimize the settings of on/off inflow-control devices—sometimes also referred to as inflow-control valves—which have discrete open or closed settings, as represented by the numerical values of 1 and 0, respectively. However, EnOpt, like other gradient-based techniques, cannot efficiently handle discrete control problems. Thus, we use a reparameterization of the controls into switching times, which will be discussed in detail in this study. By use of this technique, we investigate the applicability of an ensemble-based robust hierarchical multiobjective optimization framework to optimize an ensemble of sector models inspired from a real field case. The uncertainty in the ensemble of sector models is characterized by differing permeability, porosity, net/gross ratios, and initial water saturation fields.

Theory

In this work we have chosen to use an approximate gradient method, ensemble-based optimization (EnOpt), for the optimization instead of a derivative-free technique because the computational costs for derivative-free methods are usually higher. For our example, which incorporates geological uncertainty in the form of 100 realistic reservoir models, a derivative-free technique would be computationally extremely challenging. In addition, in our limited experience, approximate gradient techniques usually have fewer tuning parameters compared with derivative-free methods.

Objective Function. The objective function is the usual expression for (simple) net present value as used in life-cycle optimization studies. It is defined as

$$J = \sum_{k=1}^K \left(\frac{\{[(q_{o,k}) \cdot r_o - (q_{wp,k}) \cdot r_{wp}] - [(q_{wi,k}) \cdot r_{wi}]\} \cdot \Delta t_k}{(1+b)^{t_k/\tau_i}} \right), \quad \dots \dots \dots (1)$$

where $q_{o,k}$ is the oil-production rate in bbl/day, $q_{wp,k}$ is the water-production rate in bbl/day, $q_{wi,k}$ is the water-injection rate in bbl/day, r_o is the price of oil produced in USD/bbl, r_{wp} is the cost of water produced in USD/bbl, r_{wi} is the cost of water injected in USD/bbl, Δt_k is the difference between consecutive timesteps in days, b is the discount rate expressed as a fraction per year, t_k is the cumulative time in days corresponding to timestep k , and τ_i is the reference time period for discounting, typically 1 year (365.24 days), which is in line with common oil-industry project-economics practices.

Update Rules. The EnOpt methodology results in an approximate gradient $\mathbf{g}$ of the objective function J with respect to a vector of controls $\mathbf{u}$. This approximate gradient, as computed from Eq. A-11, can be used in any gradient-based optimization algorithm. In our study we used a simple steepest ascent scheme:

$$\mathbf{u}^{\ell+1} = \mathbf{u}^{\ell} + \alpha^{\ell} \mathbf{g}^{\ell}, \quad \dots \dots \dots (2)$$

where the superscript ℓ is the iteration counter, and α^{ℓ} is a step length in the direction of the gradient. Note that if $\mathbf{u}$ and $\mathbf{g}$ both contain elements with different dimensions (e.g., when the elements of $\mathbf{u}$ are pressures and rates), an additional scaling of the gradient elements may be required. We scaled the gradient by its infinity norm and then for each iteration used an initial step size $\alpha = 1$. Thereafter, we allowed for a maximum of three back-tracking steps, each time reducing the step size with a factor of one-half.

Robust Ensemble Optimization. Geological modeling and interpretation is an inherently uncertain process. Incorporating these uncertainties into the optimization framework reduces the uncertainty involved with the optimization (Yeten et al. 2003). van Essen et al. (2009) provide a list of references of nonpetroleum engineering applications that incorporated uncertainty within the modeling-and-control framework. From these, they imported the

terminology “robust optimization” into the petroleum literature (i.e., optimization over an ensemble of geological realizations that represents the geological uncertainty). In an example they showed that this technique increased the expected value and reduced the variance of the optimized strategy applied to the different geological realizations in comparison with a reactive strategy. van Essen et al. (2009) reported significant improvements in the expected objective-function values by use of adjoint-based optimization. Thereafter, Chen et al. (2012) performed robust optimization by use of the adjoint formulation to solve a multiobjective optimization problem. A robust version of EnOpt was introduced by Chen (2008) and Chen and Oliver (2010), wherein they use two ensembles, one of controls and another of geological models, which consist of M members each. Intuitively in such a scenario we would require M^2 function evaluations for a gradient estimate, which is computationally not attractive. Thus, to make the method computationally more efficient, Chen (2008) provided an argumentation for the possibility of evaluating only M samples to approximate the robust EnOpt gradient. Thus, to estimate a “robust gradient,” Chen (2008) coupled one member from the control ensemble with one member of the geological ensemble in a 1:1 ratio, which will be referred to as the “original formulation” in this study. Recently, there has been an increase in the application and understanding of the ensemble-based robust gradient estimate, which is briefly discussed in Appendix A.

Multiobjective Optimization. Most real-world problems have multiple objectives that need to be satisfied. Usually these objectives are in conflict with each other, and one must accept decreases in one objective to achieve increases in another objective. In theory, there exist many methods to solve a multiobjective problem and recently there has been an increased focus on finding methods to solve multiobjective problems in the reservoir-simulation community. These objectives are usually defined as long-term (life-cycle) objectives from a reservoir-engineering viewpoint and short-term objectives from a production-engineering/operational-constraints viewpoint. van Essen et al. (2009) showed that these two objectives are in conflict with each other and suggested the use of a hierarchical framework for multiobjective optimization. In this study we use the hierarchical-switching method with EnOpt for the optimization. Details of the implementation are given in Fonseca et al. (2014). This switching method optimizes the objectives alternately and stays within a maximum-allowed very-small decrease ε in the primary objective. An alternative to hierarchical biobjective optimization (in which the primary objective is considered much more important than the secondary objective) is regular biobjective optimization, in which there is no predefined preference for one of the objectives. Isebor and Durlowsky (2014) and Liu and Reynolds (2014) have introduced methodologies to plot the “Pareto front” for a regular biobjective reservoir-optimization problem (i.e., a set of objective-function pairs corresponding to optimal controls for many possible weighted combinations of the two objective functions, which gives the decision maker a range of possible solutions).

Control Parameterization. Recent advances in technology and the need for improved controllability of the oil-recovery process have led to the use of inflow-control devices (ICDs), also known as inflow-control valve. Although many variants of such devices exist, in our study we have used ICDs that can be individually activated by use of electric line or coiled tubing through relatively simple rigless intervention techniques. This type of ICD is currently commercially available with up to 16 ICDs per well, and is used in field applications. In the present study we optimize the settings of these ICDs over the producing life of a field. ICDs can have settings that vary continuously between 0 and 1 or could be restricted to settings of either 0 or 1; that is, either fully open or fully closed. In our study the ICDs belong to the latter class, and have settings of either 0 or 1. Thus the optimization problem is now discrete in nature. To solve such discrete problems, integer-programming techniques are usually used (Isebor et al. 2014).

Fig. 1—Permeability field for Layer 14 of Realization 65.

Fig. 2—Oil saturation of Layer 14 for Realization 65 after a time of 20 years, indicating a high water production.

Gradient-based methods like EnOpt, which have been successfully applied to problems with continuous variables, are not suitable for discrete control problems. To use EnOpt to solve this particular optimization problem, a parameterization of the controls into continuous variables is necessary. Sudaryanto and Yortsos (2000) suggested the use of switching times as control variables for a production-optimization exercise to study the behavior of bang-bang (i.e., on/off) controls. Zandvliet et al. (2007) investigated the theoretical aspects of bang-bang control problems and provided a list of references from different engineering applications where switching times have been used as controls. However, they did not use switching times as the controls; instead, they used continuous control variables in combination with a cutoff criterion to mimic the discrete controls. We have also investigated this approach; however, we did not achieve solutions that were better than a switching-times approach. Hasan and Foss (2013) recently used switching times as controls for adjoint-based waterflooding optimization with the halving-time-interval method to update the controls, and Namdar Zanganeh et al. (2014) used a switching approach to optimize a surfactant (foam)-injection application. In this study we use the switching-time-interval-based parameterization of the controls, modeled after the switching-time-optimization method provided by Sudaryanto and Yortsos (2000). Whereas Sudaryanto and Yortsos (2000) use explicit times to define a switch, we use time intervals as controls. There are two distinct advantages for the use of such a parameterization:

- It transforms a discrete control problem into a continuous control problem that can be efficiently handled by EnOpt.
- It leads to a possibly significant reduction in the number of control variables. This could be particularly important when using stochastic approximate gradient-based techniques such as EnOpt because the gradient quality may deteriorate for increasing numbers of optimization variables.

For example, in this case, if we choose fixed control timesteps of 1 year and use 48 ICD settings to be optimized at each control timestep, the use of amplitude-based controls for a 20-year period would result in a control vector $\mathbf{u}$ containing $48 \times 20 = 960$ elements. In the switching-times approach, the user must predefine the number of allowable switches during the simulation time for each ICD. In our case, we use five switching times for each ICD over the producing life of the reservoir, thus leading to $48 \times 5 = 240$ controls (i.e., a factor-of-four reduction in the number of controls). Note that we work with time intervals to define the times at which a particular ICD must be either switched on or off depending on its previous setting. We do not explicitly find the time a control is switched; rather, we find the interval after which a control setting should be changed. Also note that five switching times is a maximum limit and it is possible to achieve an optimized strategy with fewer switches. In case the final switching-time interval exceeds the end of the producing time, the interval is simulated only until the end time.

Numerical Example

Reservoir Model. The reservoir model used in this study is a sector model inspired by a real field case. The reservoir formation

is unconsolidated sandstone at a depth of approximately 4,000 ft, with a net thickness ranging from 30 to 45 ft, and is not significantly faulted. 50,000 active gridblocks are used in this sector model. Model dimensions cannot be disclosed for confidentiality reasons. The reservoir rock is of good quality with porosities ranging from 20 to 35% and net/gross ratios of approximately 50–90%. The permeability distribution is not very heterogeneous with values ranging from 100 to 700 md, usually approximately 350 md. Permeability modeling has been performed following industry standard practices, modeled after facies and petrophysical modeling. In addition, an integration of core data coming from some of the wells with computer-processed interpretation logs has been incorporated. A porosity/permeability correlation has been identified and grids have been populated by use of a standard geostatistical algorithm. After this, well-test data have been integrated and the initial ensemble of permeabilities have been conditioned to historical data to obtain an updated ensemble, which is used in this study. The reservoir is operated by use of a line-drive strategy with horizontal wells with lengths of 4,000–10,000 ft, with inflow-control devices (ICDs) installed in some of the injectors. The field has unconventional reservoir and fluid properties compared with other fields in the vicinity. In particular, the field temperature is lower and the oil is heavier and more viscous, compared with neighboring fields.

Fig. 1 is an illustration of the permeability field of Layer 14 from Realization 65 out of an ensemble of 100 realizations. The sector model shown is produced by use of a line-drive strategy with three injectors and three producers. The injectors are equipped with ICDs that can be mechanically operated by use of a coiled-tubing unit and can only be either “open” or “closed,” with values of either 0 or 1. Injector 1 has 15 ICDs along its length, whereas Injectors 2 and 3 have 16 and 17 ICDs, respectively. The settings of these ICDs are the control variables for this optimization study. The wells are modeled as multisegment wells in a commercial fully implicit finite-difference black-oil simulator (Eclipse 2011). Because the ICDs cannot be modeled to be fully closed—i.e., 0—because of numerical limitations, in this exercise the valve settings have a minimum value of 10^{-4} . To effectively capture the effect of geological uncertainty within the optimization framework, we use a set of 100 different geological (sector) models. The realizations vary in terms of their permeability fields, porosity fields, varying net/gross ratios, and initial water saturations.

The relative homogeneity in the permeability fields, coupled with the specific properties of the oil, suggest that for this model the scope for optimization lies in the reduction of the volumes of water injected and produced. Fig. 2 illustrates the oil saturation of Layer 14 from Realization 65 after 20 years of production. The field water cut for a 20-year horizon is approximately 89%, and this increase in the volumes of water produced and injected makes the problem interesting for optimization. Hence, we have chosen an optimization time horizon of 20 years.

Life-Cycle Optimization. In this model, the injectors and producers are operated on bottomhole-pressure constraints, whereas the ICD settings are allowed to vary over the producing life (20

Fig. 3—Mean objective (undiscounted NPV) -function value for the 20-year-simulation time period. The red line is the mean value of the objective value, whereas the blue lines are the objective-function values for the individual ensemble members through the iteration process.

years) of the reservoir. On the basis of the engineering judgment, we allow for only five switching times per ICD throughout the 20 years, and because we have 48 ICDs per control time, the control vector $\mathbf{u}$ has $N = 5 \times 48 = 240$ elements. An optimal life-cycle strategy of ICD settings in the injection wells is obtained by optimizing net present value (NPV), as described in Eq. 1 with $r_o = 90$ USD/bbl, $r_{wp} = 8$ USD/bbl, and $r_{wi} = 5$ USD/bbl. The discount rate b is chosen to be either 0 or 10%. The initial strategy (starting point) of the life-cycle optimization is a control vector with all switching-time intervals equal to 0; i.e., there are no switches allowed, which implies that all the ICDs are fully open at all control times. **Fig. 3** illustrates the optimization process, where the blue lines represent the evolution of the objective function values for the 100 different geological realizations during the iterations, whereas the red line is the expected value of the ensemble. The optimized expected objective-function value is approximately USD 695 million. Because of a lack of significant further change in the objective-function value, the optimization process was terminated after 30 iterations.

We observe, for an undiscounted NPV, an approximately 12% increase in mean objective-function value at the end of the optimization compared with the starting point of the optimization (that is, compared with a “do-nothing” strategy). Although this result is encouraging, an analysis of the cumulative oil and water rates shows an approximately 10% mean decrease in cumulative oil production for a corresponding 47.5% mean decrease in cu-

Fig. 5—10% discounted NPV, optimization started from end of Fig. 3 (i.e., the optimized strategy).

Fig. 4—10% discounted NPV, optimization started from “do-nothing” strategy.

mulative water injected and 70% mean decrease in cumulative water produced. Although the volumes of water injected and produced have reduced significantly, this corresponds to a large decrease in cumulative oil production, which may not be the most-attractive strategy.

Effect of Discount Factor. In reality, a discounted economic objective is traditionally used in the analysis of any project. Thus we also use discounted NPV as the objective function with a 10% discount factor. First, as in Fig. 3, we start the optimization from the “do-nothing” strategy, which is a strategy where all the ICDs are open throughout the life of the field. A comparison of this result with a simulator-handled pressure-maintenance strategy shows a 2.8% mean increase in NPV (**Fig. 4**). The mean cumulative oil volumes obtained with this strategy are higher than with the pressure-maintenance strategy. However, the volumes of water injected and produced are also higher. Because the starting point of the optimization is rather aggressive (in terms of water-injection volumes), we find an optimized strategy that reflects the same behavior. This strategy is hereafter referred to as Optimization Strategy 1.

Following this result, we perform an optimization experiment from a different initial strategy, or a different point in the control space. The initial strategy is the optimized strategy from the undiscounted life-cycle result, shown in Fig. 3. We observe that starting from a different initial strategy has led us to a solution with a higher NPV compared with Fig. 4. In addition, compared with the pressure-maintenance strategy, this optimized strategy (**Fig. 5**), hereafter referred to as Optimization Strategy 2, achieves a mean increase of 4.2%. Thus the optimization is fairly sensitive to the initial starting point of the optimization. Optimization Strategy 2 starts from a significantly less-aggressive strategy (in terms of volumes of water injected) compared with Optimization Strategy 1 and thus achieves a lower mean cumulative oil production but also injects much-lower volumes of water. However, irrespective of the starting point, the optimized strategy always achieves better solutions in terms of NPV compared with the simulator-handled pressure-maintenance strategy.

If we perform an optimization exercise while requiring the simulator to enforce pressure maintenance, we observe that the optimization is not successful, because of interference of the simulator-handled control with the optimization (**Fig. 6**).

Table 1 highlights the key differences between the two different optimized strategies when using a discounted objective function. We observe that although Optimization Strategy 2 achieves a higher NPV compared with Optimization Strategy 1, the amount of cumulative oil produced is marginally lower than with the simulator-handled pressure maintenance. However, there is a significant reduction in cumulative volumes of water needed for injection as well as water produced. On the other hand,

Fig. 6—10% discounted NPV with pressure maintenance handled by the simulator.

Optimization Strategy 1 achieves not only a higher NPV but also a 3% mean increase in cumulative oil produced at the cost of a higher cumulative water-injection volume, and thus a higher volume of water produced. The difference for the two strategies stems from the different initial starting points. Although it is difficult to conclude in favor of either strategy, decision makers now have a choice depending on their objectives.

Reactive Control. The economic water cut derived from the economic parameters used to calculate the objective function is 91%. However, after 20 years of production the average water cut is

89%, whereas none of the wells has yet reached the 91% value. Thus we have not yet reached a true reactive control strategy, and the reactive control strategy corresponds to the “do-nothing” strategy, the NPV of which was the starting point of the optimization.

Optimized Control Strategy. Fig. 7 is an illustration of the comparison between an optimized strategy (red) of 16 ICDs installed in Injector 2 with the reactive strategy. Recall that we defined a maximum of five allowable switches for each ICD throughout the 20-year optimization horizon. We observe that for the optimized strategy, for most cases, the ICDs do not need five switches; instead, the optimized strategy consists of mainly one or two switches per ICD. Many life-cycle optimization studies have obtained optimal control sets that are nonsmooth in nature; that is, they display frequent adjustments to the control settings, which is practically undesirable and probably not feasible to implement. However, an optimized strategy as illustrated in Fig. 7 would be more appealing to implement as an operational strategy because of its smooth behavior. The settings for the ICDs in Injectors 1 and 3 also showed very similar behavior to that observed in Fig. 7.

Comparison of Different Gradient Formulations. Many alternative formulations to the standard robust ensemble-gradient estimate have been proposed. Because of the computational complexity of the model, on the basis of the findings of Fonseca et al. (2014, 2015), we investigate the application of the three robust gradient formulations discussed in the theory section. Fig. 8 depicts the optimization process obtained by applying the different formulations. We observe that the original formulation (1:1 ratio) achieves a solution superior to the traditional pressure maintenance, “do-nothing” strategy; however, the “selected-model formulation” (blue curve) leads to a solution that achieves an expected value 5% higher than

Property	Pressure Maintenance	Optimization Strategy 1	Optimization Strategy 2
NPV at 10%	3.8363e ⁸	3.9432e ⁸ (+2.8%)	3.9971e ⁸ (+4.2%)
Cumulative oil (STB)	1.3712e ⁵	1.4130e ⁵ (+3.0%)	1.3604e ⁵ (−0.8%)
Cumulative water produced (STB)	3.2245e ⁵	3.3985e ⁵ (+5.4%)	2.8454e ⁵ (−13%)
Cumulative water injected (STB)	4.6053e ⁵	4.8332e ⁵ (+5.0%)	4.2054e ⁵ (−9.5%)

Table 1—Comparison of different optimization strategies.

Fig. 7—Comparison of Optimization Strategy 2 (red) for Injector 2 over the 20-year horizon with the “do-nothing”/reactive strategy (blue). In most cases, the ICDs make one or two switches over the 20-year period.

Fig. 8—Comparison of the optimization performance for different gradient formulations.

the original formulation (green curve). It should be noted that there is always ambiguity with respect to how the models used in this approach are chosen and what impact a different set of chosen models would have on the optimization. The modified robust gradient formulation (red curve) achieves the highest-expected objective-function value, which is approximately 5.5% higher than the value for the original formulation and 0.5% higher compared with the selected-model formulation. The modified-gradient formulation, like the original formulation, uses the entire ensemble of geological realizations; i.e., it accounts for all the uncertainty available to estimate the robust gradient, unlike the selected-model approach.

Reduction in Number of ICDs. Because ICDs are expensive to install and operate, reducing the number of ICDs could be economically beneficial. However, the impact of having fewer ICDs could result in a reduction in controllability. The grouping of the ICDs is primarily dependent on the geological perspective and well-path design. van Essen et al. (2010) showed that a grouping of ICDs dependent on dynamic results instead of a geology-based grouping may lead to better results in terms of the objective-function value. The well-path design used in the present study is undulating. Compared with the example in van Essen et al. (2010), in which the horizontal wells are completely horizontal, we expect that dynamic grouping will not lead to improved results. The dynamic grouping methodology proposed in van Essen et al. (2010) is modeled after a visualization of the optimized control strategy. Visual inspection of the optimal control set illustrated Fig. 7 does not suggest any apparent dynamic grouping possibilities. Fig. 9 is a comparison of the optimization procedure with a significant (approximately a factor of four) reduction (blue curve) in the number of ICDs dependent on geological insight compared with the base case (red curve). We observe that having fewer ICDs results in a loss of controllability and thus an optimized strategy with a lower NPV. The difference in the objective-function values between the two cases is approximately 3%, where the cost reduction caused by a reduced number of ICDs has not been taken into account. We do not have data about the actual costs (installation) of a single ICD. On the basis of discussions with field specialists, the total costs saved as a result of installation of fewer ICDs used for this case is estimated, on the higher side, to be two orders of magnitude lower than the values of NPV shown in Fig. 9. Thus, for this example, the costs saved by the installation of fewer ICDs does not offset the gain in NPV achieved by use of a higher number of ICDs.

Hierarchical-Switching Optimization. The hierarchical (bio-jective) switching-optimization method is used to achieve multi-objective optimization under uncertainty, as illustrated in Fig. 10. Following van Essen et al. (2011), the primary objective is undiscounted NPV, as shown in Fig. 3, whereas the secondary objective is a highly (25%) discounted NPV to account for the short-term gains. We observe a mean increase of approximately 9.2% in the secondary objective function compared with a maximum-

Fig. 9—Comparison of the impact of the use of a reduced number of ICDs on the basis of geology and well-path designs. A reduced set of ICDs (blue) results in reduced controllability and thus a lower NPV.

allowed mean decrease ε of 1% in the primary objective function. The switching optimization begins from the optimized solution achieved by the modified robust gradient formulation for life-cycle optimization. The same modified formulation is also used for this hierarchical optimization. In addition, the mean increase in cumulative oil produced over 20 years is marginally (2%) higher because of a higher (4.5%) increase in cumulative water injected compared with the solution achieved for life-cycle optimization. This confirms the general trend observed for this sector model (i.e., higher volumes of water injected will result in higher volumes of oil produced). The results illustrate the capacity of ensemble-based multiobjective optimization under geological uncertainty to achieve results of practical importance.

Note that we have terminated the optimization for the secondary objective function after 15 iterations, which translates to 1,500 reservoir simulations for the gradient estimate and 2,100 simulations for evaluation of the updated control set (i.e., a total of 3,600 reservoir simulations). Although the total number of simulations is not high, because of the complex nature of the models, a forward simulation lasts approximately 15 minutes. Thus, to obtain these results in a sequential manner would take roughly 37 days. An inherent advantage of ensemble-based optimization is that the gradient can be estimated by use of distributed computing, as has been performed for the present study. The speedup achieved in this case, by use of 25 cores, is roughly a factor of nine (i.e., the robust hierarchical optimization was performed in approximately 4 days).

Fig. 11 depicts a comparison of the mean cumulative cash flow over time for the optimized solutions achieved by the switching algorithm (blue), life-cycle optimization (green), and the “do-nothing” strategy (red). It is observed that after 2,000 days, the cumulative cash flow of the multiobjective optimization (blue curve) is approximately 12% higher compared with the life-cycle optimization (green curve), which will enable the project to achieve the break-even point faster, whereas the ultimate NPV of the two strategies is nearly equal. Similar to the results obtained in Fonseca et al. (2014), the “do-nothing”/reactive control strategy gives the best short-term performance at a price of the worst long-term performance. Note that in the “do-nothing” scenario, none of the production wells has reached the economic water-cut limit of 91%, and the drop in NPV after approximately 6,500 days is caused by water-injection costs.

Discussion

We have not considered the commonly used weighted-sum method in this work because it is impossible to know a priori which weight combination will give the desired results. The advantage of the hierarchical approach is that the user decides the

Fig. 10—Multiobjective optimization with an initial strategy dependent on life-cycle 0% discounted NPV, showing a mean 9.2% increase in the secondary objective function (25% discounted NPV) for a 1% mean decrease in primary objective function.

maximum-allowable decrease in the primary objective. This feature is not known a priori when using the weighted-sum method. We would need to perform a trial-and-error set of experiments, which would be computationally demanding.

Also, we did not use a full Pareto-curve approach, in which a large number of weighted-sum simulations is performed with different weight combinations. Generating such a full Pareto curve would provide a decision maker a range of possible solutions from which to choose. However, computing the full Pareto curve is computationally much more intensive and is outside the computational limits of our study. The choice of $\varepsilon = 1\%$ in our hierarchical-switching approach is simply a choice; a user is free to decide his own choice depending on how much a user values the long-term targets. We do not claim that this is the correct choice, because, as with any multiobjective approach, especially those that aim to generate a Pareto front, the idea is to obtain a range of possible solutions from which to choose. No single solution is necessarily the correct one. van Essen et al. (2011), who first introduced the hierarchical-switching method, have included an illustration of the principle and execution of this hierarchical-switching method.

By request of one of the reviewers, we repeated the optimization for a much-lower oil price of USD 50/bbl while keeping the water costs the same. **Fig. 12** depicts the results. Compared with

Fig. 3, we now increase the objective-function value much more (approximately 35% against approximately 12% in Fig. 3), which illustrates that at lower oil prices the effect of reducing water production and injection becomes more important.

Conclusions

1. An ensemble-based robust multiobjective optimization workflow tested on a sector model inspired from a real field case shows results of practical value against acceptable computational cost.
2. Parameterization of on/off-type controls by use of a switching-time-interval method is efficient when working with stochastic gradient techniques, such as ensemble-based optimization.
3. The hierarchical-switching algorithm leads to an approximately 9% mean increase in the secondary objective function (short-term targets) against a mean decrease of 1% in the primary objective function (life-cycle targets).
4. The main scope of optimization for this example lies in the reduction of the volumes of water injected and produced.
5. The modified robust gradient formulation performs better than other formulations described in the literature.
6. A reduction in the number of inflow-control devices (ICDs) results in a loss of controllability and thus lower objective-

Fig. 11—Comparison of the mean cash flow over time for the entire ensemble of geological realization for the different optimization strategies.

Fig. 12—Mean increase in objective-function value over the ensemble of realizations for an oil price of USD 50/bbl.

function values. However, it may still be a better strategy if the cost of the ICDs is incorporated into the objective function.

- The optimization is sensitive to the initial starting point. Two different optimized strategies have been provided for the decision maker to choose from.

Nomenclature

a = step size
 b = discount rate
 $\mathbf{c}_{ul}$ = ensemble cross-covariance vector
 $\mathbf{C}$ = distribution-covariance matrix
 $\mathbf{C}_{uu}$ = ensemble-covariance matrix
 $\mathbf{g}$ = gradient vector
 $\mathbf{j}$ = vector of mean-shifted objective-function values
 J = objective-function value
 $\bar{J}$ = mean objective-function value
 k = timestep counter
 K = total number of timesteps
 ℓ = iteration counter
 M = number of ensemble members
 N = number of control variables
 q = flow rates
 r = price per unit volume
 t = time
 u = control variable
 $\mathbf{u}$ = vector of control variables
 $\bar{\mathbf{u}}$ = ensemble mean
 $\bar{\mathbf{U}}$ = matrix of ensemble-mean-shifted control vectors
 ε = maximum allowed decrease in primary-objective-function value
 τ = reference time for discounting

Subscripts

o = oil
 w = water
 wi = injected water
 wp = produced water

Acknowledgments

This research was performed within the context of the Integrated Systems Approach to Petroleum Production (ISAPP) knowledge center. ISAPP is a joint project of TNO, Delft University of Technology, Eni, Statoil, and Petrobras. The authors acknowledge Schlumberger for providing multiple academic Eclipse licenses for this work. The authors also acknowledge Stefano Raniolo for assistance in setting up the simulation models.

References

- Alhuthali, A. H. H., Datta-Gupta, A., Yuen, B. et al. 2009. Field Applications of Waterflood Optimization via Optimal Rate Control With Smart Wells. Presented at the SPE Reservoir Simulation Symposium, The Woodlands, Texas, 2–4 February. SPE-118948-MS. <http://dx.doi.org/10.2118/118948-MS>.
- Bailey, W. J., Couet, B. and Wilkinson, D. 2005. Field Optimization Tool for Maximizing Asset Value. *SPE Res Eval & Eng* **8** (1): 7–21. SPE-87026-PA. <http://dx.doi.org/10.2118/87026-PA>.
- Chaudhri, M. M., Phale, H. A. and Liu, N. 2009. An Improved Approach For Ensemble Based Production Optimization: Application To A Field Case. Presented at the EUROPEC/EAGE Conference and Exhibition, Amsterdam, 8–11 June. SPE-121307-MS. <http://dx.doi.org/10.2118/121307-MS>.
- Chen, C., Li, G. and Reynolds, A. 2012. Robust Constrained Optimization of Short- and Long-Term Net Present Value for Closed-Loop Reservoir Management. *SPE J.* **17** (3): 849–864. SPE-141314-PA. <http://dx.doi.org/10.2118/141314-PA>.
- Chen, Y. 2008. *Efficient Ensemble-Based Reservoir Management*. PhD dissertation, University of Oklahoma, Norman, Oklahoma (2008).
- Chen, Y. and Oliver, D. S. 2010. Ensemble-Based Closed-Loop Optimization Applied to Brugge Field. *SPE Res Eval & Eng* **13** (1): 56–71. SPE-118926-PA. <http://dx.doi.org/10.2118/118926-PA>.
- Chen, Y., Oliver, D. S. and Zhang, D. 2009. Efficient Ensemble-Based Closed-Loop Production Optimization. *SPE J.* **14** (4): 634–645. SPE-112873-PA. <http://dx.doi.org/10.2118/112873-PA>.
- Conn, A. R., Scheinberg, S. and Vicente, L. N. 2009. *Introduction to Derivative-Free Optimization*. Philadelphia, Pennsylvania: Society for Industrial and Applied Mathematics.
- Dehdari, V. and Oliver, D. S. 2012. Sequential Quadratic Programming for Solving Constrained Production Optimization—Case Study From Brugge Field. *SPE J.* **17** (3): 874–884. SPE-141589-PA. <http://dx.doi.org/10.2118/141589-PA>.
- Do, S. T. and Reynolds, A. C. 2013. Theoretical Connections Between Optimization Algorithms based on an Approximate Gradient. *Computat. Geosci.* **17** (6): 959–973. <http://dx.doi.org/10.1007/s10596-013-9368-9>.
- Eclipse. 2011. <http://www.slb.com/services/software/reseng/eclipse>.
- Fonseca, R. M., Stordal, A. S., Leeuwenburgh, O., et al. 2014. Robust Ensemble-Based Multi-Objective Optimization. Oral presentation given at the 14th European Conference on the Mathematics of Oil Recovery (ECMOR XIV), Catania, Italy, 8–11 September.
- Fonseca, R. M., Kahrobaei, S., Van Gastel, L. J. T., et al. 2015. Quantification of the Impact of Ensemble Size on the Quality of an Ensemble Gradient Using Principles of Hypothesis Testing. Presented at the 2015 SPE Reservoir Simulation Symposium, Houston, 22–25 February. SPE-173236-MS. <http://dx.doi.org/10.2118/173236-MS>.
- Forouzanfar, F., Della Rossa, E., Russo, R., et al. 2013. Life-Cycle Production Optimization of an Oil Field with an Adjoint-Based Gradient Approach. *J. Pet. Sci. Eng.* **112** (December): 351–358. <http://dx.doi.org/10.1016/j.petrol.2013.11.024>.
- Golub, G. and van Loan, C. 1989. *Matrix Computations*, 4th ed. The John Hopkins University Press, Baltimore, USA.
- Hasan, A. and Foss, B. 2013. Optimal Wells Scheduling of a Petroleum Reservoir. Oral presentation given at the European Control Conference (ECC), Zurich, Switzerland, 17–19 July.
- Isebor, O. J. and Durlafsky, L. J. 2014. Biobjective Optimization for General Oil Field Development. *J. Pet. Sci. Eng.* **119** (July): 123–138. <http://dx.doi.org/10.1016/j.petrol.2014.04.021>.
- Isebor, O. J., Echeverría Ciaurri, D. and Durlafsky, L. J. 2014. Generalized Field-Development Optimization With Derivative-Free Procedures. *SPE J.* **19** (5): 891–908. SPE-163631-PA. <http://dx.doi.org/10.2118/163631-PA>.
- Jansen, J. D. 2011. Adjoint-based Optimization of Multi-Phase Flow Through Porous Media—A review. *Comput. Fluids* **46** (1) 40–51. <http://dx.doi.org/10.1016/j.complfluid.2010.09.039>.
- Jansen, J. D., Douma, S. G., Brouwer, D. R., et al. 2009. Closed Loop Reservoir Management. Presented at SPE Reservoir Simulation Symposium, The Woodlands, Texas, 2–4 February. SPE-119098-MS. <http://dx.doi.org/10.2118/119098-MS>.
- Li, L., Jafarpour, B. and Mohammad-Khaninezhad, M. R. 2012. A Simultaneous Perturbation Stochastic Approximation Algorithm for Coupled Well Placement and Control Optimization Under Geologic Uncertainty. *Computat. Geosci.* **17** (1): 167–188. <http://dx.doi.org/10.1007/s10596-012-9323-1>.
- Liu, X. and Reynolds, A. C. 2014. Gradient-Based Multi-Objective Optimization with Applications to Waterflooding Optimization. Oral presentation given at the 14th European Conference on the Mathematics of Oil Recovery (ECMOR XIV), Catania, Italy, 8–11 September.
- Lorentzen, R. J., Berg, A., Naevdal, G., et al. 2006. A New Approach for Dynamic Optimization of Waterflooding Problems. Presented at the Intelligent Energy Conference and Exhibition, Amsterdam, 11–13 April. SPE-99690-MS. <http://dx.doi.org/10.2118/99690-MS>.
- Namdar Zanganeh, M., Kraaijevanger, J. F. B. M., Buurman, H. W., et al. 2014. Challenges in Adjoint-Based Optimization of a Foam EOR Process. *Computat. Geosci.* **18** (3–4): 563–577. <http://dx.doi.org/10.1007/s10596-014-9412-4>.
- Nwaozo, J. 2006. *Dynamic Optimization of a Waterflooding Reservoir*. Master's thesis, University of Oklahoma, Norman, Oklahoma (2006).
- Oliver, D. S., Reynolds, A. C. and Liu, N. 2008. *Inverse Theory for Petroleum Reservoir Characterization and History Matching*. Cambridge, UK: Cambridge University Press.

- Pajonk, O., Schulze-Riegert, R., Krosche, M., et al. 2011. Ensemble-Based Water Flooding Optimization Applied to Mature Fields. Presented at the SPE Middle East Oil and Gas Show and Conference, Manama, Bahrain, 25–28 September. SPE-142621-MS. <http://dx.doi.org/10.2118/142621-MS>.
- Raniolo, S., Dovera, L., Cominelli, A., et al. 2013. History Match and Polymer Injection Optimization in a Mature Field Using the Ensemble Kalman Filter. *Proc.*, 17th European Symposium on Improved Oil Recovery, St. Petersburg, Russia, 16–18 April. <http://dx.doi.org/10.3997/2214-4609.20142642>.
- Sarma, P., Durlafsky, L. J. and Aziz, K. 2008. Computational Techniques for Closed-Loop Reservoir Modeling with Application to a Realistic Reservoir. *Pet. Sci. Tech.* **26** (10–11): 1120–1140. <http://dx.doi.org/10.1080/10916460701829580>.
- Stordal, A. S., Szklarz, S. P. and Leeuwenburgh, O. 2014. A Closer Look at Ensemble Based Optimization in Reservoir Management. Oral presentation given at the 14th European Conference on the Mathematics of Oil Recovery (ECMOR XIV), Catania, Italy, 8–11 September.
- Strang, G. 2006. *Linear Algebra and its Applications*, fourth edition. Boston, Massachusetts: Brooks Cole.
- Sudaryanto, B. and Yortsos, Y. C. 2000. Optimization of Fluid Front Dynamics in Porous Media Using Rate Control. *Phys. Fluids* **12** (7): 1656–1670. <http://dx.doi.org/10.1063/1.870417>.
- van Essen, G. M., Jansen, J. D., Brouwer, R., et al. 2010. Optimization Of Smart Wells in the St. Joseph Field. *SPE Res Eval & Eng* **13** (4): 588–595. SPE-123563-PA. <http://dx.doi.org/10.2118/123563-PA>.
- van Essen, G. M., Van den Hof, P. M. J. and Jansen, J. D. 2011. Hierarchical Long-Term and Short-Term Production Optimization. *SPE J.* **16** (1): 191–199. SPE-124332-PA. <http://dx.doi.org/10.2118/124332-PA>.
- van Essen, G. M., Zandvliet, M. J., Van den Hof, P. M. J., et al. 2009. Robust Waterflooding Optimization of Multiple Geological Scenarios. *SPE J.* **14** (1): 202–210. SPE-102913-PA. <http://dx.doi.org/10.2118/102913-PA>.
- Yeten, B., Durlafsky, L. J. and Aziz, K. 2003. Optimization of Nonconventional Well Type, Location, and Trajectory. *SPE J.* **8** (3): 200–210. SPE-86880-PA. <http://dx.doi.org/10.2118/86880-PA>.
- Zandvliet, M. J., Bosgra, O. H., Van den Hof, P. M. J., et al. 2007. Bang-Bang Control and Singular Arcs in Reservoir Flooding. *J. Pet. Sci. Eng.* **58** (1–2): 186–200. <http://dx.doi.org/10.1016/j.petrol.2006.12.008>

Appendix A—Ensemble-Based-Optimization (EnOpt) Theory

Deterministic Formulation. Chen (2008) and Chen et al. (2009) gave systematic descriptions of the EnOpt method as it is currently used. Predecessors to EnOpt, modeled after the ensemble Kalman filter, were introduced by Lorentzen et al. (2006) and Nwazo (2006). The following is a short overview of the method, largely following the description in Fonseca et al. (2015). A single control vector is defined as

$$\mathbf{u} = [u_1, u_2, \dots, u_N]^T, \quad \dots \quad (\text{A-1})$$

where N is the total number of control variables (e.g., bottomhole pressures, well rates, or valve settings over time). Thus $\mathbf{u}$ is a vector with N elements, where N usually equals the number of control timesteps multiplied by the number of control variables per timestep. In EnOpt, a multivariate, Gaussian-distributed ensemble $[\mathbf{u}_1, \mathbf{u}_2, \dots, \mathbf{u}_M]$ of size M is generated with a distribution mean $\bar{\mathbf{u}}$ and a predefined distribution-covariance matrix $\mathbf{C}$, which is kept constant throughout the optimization process, whereas $\mathbf{u}$ is updated until convergence. To estimate the gradient, a mean-shifted ensemble matrix is defined as

$$\mathbf{U} = [\mathbf{u}_1 - \bar{\mathbf{u}}, \mathbf{u}_2 - \bar{\mathbf{u}}, \dots, \mathbf{u}_M - \bar{\mathbf{u}}], \quad \dots \quad (\text{A-2})$$

where

$$\bar{\mathbf{u}} = \frac{1}{M} \sum_{i=1}^M \mathbf{u}_i, \quad \dots \quad (\text{A-3})$$

is the ensemble mean (i.e., the sample mean that is an estimator of the distribution mean $\mathbf{u}$). [Note that in earlier publications we used the transposed version of $\mathbf{U}$. We modified our notation to bring it in line with that of publications such as Conn et al. (2009)]. Similarly, a mean-shifted objective-function vector is defined as

$$\mathbf{j} = [J_1 - \bar{J}, J_2 - \bar{J}, \dots, J_M - \bar{J}]^T, \quad \dots \quad (\text{A-4})$$

where values of J_i correspond to the simulated response to control vectors $\mathbf{u}_i$, and where

$$\bar{J} = \frac{1}{M} \sum_{i=1}^M J_i. \quad \dots \quad (\text{A-5})$$

The approximate gradient as proposed by Chen (2008) and Chen et al. (2009) is given by

$$\mathbf{g} = \mathbf{C}_{uu}^{-1} \mathbf{c}_{uJ}, \quad \dots \quad (\text{A-6})$$

where

$$\mathbf{C}_{uu} = \frac{1}{M-1} (\mathbf{U} \mathbf{U}^T), \quad \dots \quad (\text{A-7})$$

and

$$\mathbf{c}_{uJ} = \frac{1}{M-1} (\mathbf{U} \mathbf{j}), \quad \dots \quad (\text{A-8})$$

are ensemble (sample)-covariance and cross-covariance matrices, respectively. (Note that $\mathbf{c}_{uJ}$ is a 1D matrix; i.e., a vector.) For the usual case where $M < N$, matrix $\mathbf{C}_{uu}$ is rank-deficient, and Chen (2008) and Chen et al. (2009) therefore propose not to use Eq. A-6 but to instead use

$$\mathbf{g}' = \mathbf{c}_{uJ}, \quad \dots \quad (\text{A-9})$$

or

$$\mathbf{g}'' = \mathbf{C}_{uu} \mathbf{c}_{uJ}, \quad \dots \quad (\text{A-10})$$

and both can be interpreted as “smoothed” gradients. In the present study, we use a straight gradient—Eq. A-6—computed as the underdetermined least-squares solution,

$$\mathbf{g} = (\mathbf{U} \mathbf{U}^T)^{-1} \mathbf{U} \mathbf{j} = \mathbf{U}^\dagger \mathbf{j}, \quad \dots \quad (\text{A-11})$$

where the superscript $\dagger$ indicates the Moore-Penrose pseudoinverse (Golub and van Loan 1989), which is conveniently computed by use of a singular-value decomposition (Strang 2006). This formulation was also described by Dehdari and Oliver (2012), whereas Do and Reynolds (2013) recently demonstrated that it is akin to what is known as a “simplex gradient” (Conn et al. 2009).

Robust Formulation. Original Robust Formulation. Chen (2008) and Chen and Oliver (2010) introduced a technique to include the concept of robust optimization within the EnOpt framework. They suggested the use of two ensemble sets: one ensemble of controls and another ensemble of geological models. Intuitively, if the two ensembles consist of M members each, then the use of Eq. A-11 would require M^2 function evaluations (i.e., reservoir simulations) for a gradient estimate, which is computationally not attractive. Chen (2008) provided an argumentation for the possibility of evaluating only M samples (by coupling one geological-ensemble member with one control-ensemble member) to approximate the robust EnOpt gradient. Note that the size of the two ensembles must be the same. The use of such a 1:1 ratio makes EnOpt computationally very attractive for robust optimization. Thus, the original formulation will be similar to Eqs. A-2, A-4, and (A-11) when using the 1:1 ratio. In Stordal et al. (2014), it is shown that this implementation converges to the correct gradient when the

ensemble size becomes very large. However, the original formulation with the 1:1 ensemble ratio described previously has not always lead to satisfactory results (Raniolo et al. 2013; Fonseca et al. 2014). Various studies have suggested variations to the original formulation to improve the robust gradient estimate. In Fonseca et al. (2014, 2015), descriptions of the different variations are provided and their results compared, and only a few methods of which are compared in this study.

Selected-Model Robust Formulation. Raniolo et al. (2013) observed that for a field-case-based study, the original formulation did not achieve results of any significant value, which they concluded to be the result of an inaccurate robust gradient estimate with the 1:1 ratio. They proposed to instead use a 1:20 ratio; that is, to couple each geological realization with 20 control samples to estimate individual gradients, the summation of which results in a single robust gradient. If an ensemble size of 100 was used, this would mean 2,000 reservoir simulations to approximate a single gradient. Because this is computationally not very efficient, they proposed the use of a subset of five models that would approximately capture the geological uncertainty. With this approach, they achieved better results than with the original formulation, which suggests that better gradient estimates can be obtained by changing the ratio. Note that the set of selected models remained fixed throughout the optimization and that a different subset of model realizations could have lead to different results.

Modified Robust Formulation. A modified formulation for EnOpt for deterministic optimization was proposed by Do and Reynolds (2013). The modification lays in the control perturbations and the resulting objective-function anomalies, which are computed with respect to their current distribution means (i.e., the control vector and its corresponding objective-function value used for the most-recent optimization iteration, rather than with respect to the sample means, as is the case in the original formulation). Fonseca et al. (2014, 2015) proposed a further modification to this formulation, wherein the objective-function anomalies are computed with respect to the individual objective-function values corresponding to the current control for each individual-model realization (and not with respect to the mean). This modification leads to a different weighting scheme for the gradient estimate wherein each realization has a weighting depending on the relative increase or decrease of its objective-function value. The 1:1 ratio still holds for this gradient formulation. Thus, the control and objective-function anomalies for the modified formulation as proposed by Fonseca et al. (2014, 2015) are given by

$$\mathbf{U} = [\mathbf{u}_1 - \mathbf{u}^\ell, \quad \mathbf{u}_2 - \mathbf{u}^\ell, \quad \dots \quad \mathbf{u}_M - \mathbf{u}^\ell], \quad \dots \quad (\text{A-12})$$

$$\mathbf{j} = [J_1 - J_1^\ell, \quad J_2 - J_2^\ell, \quad \dots \quad J_M - J_M^\ell]^T, \quad \dots \quad (\text{A-13})$$

where the superscript ℓ is the optimization iteration counter. (Thus, the “current control” $\mathbf{u}^\ell$ corresponds to iteration ℓ , whereas the perturbations $\mathbf{u}_i, i = 1, 2, \dots, M$ are used to compute the gradient for the “next control” $\mathbf{u}^{\ell+1}$.) This modified formulation has three distinct differences compared with the original formulation. First, as described previously, the subtractions in the objective-function values in Eq. A-13 are with respect to the individual objective-function values J_i^ℓ , instead of the mean objective-function value J^ℓ , as proposed by Do and Reynolds (2013). Second, for bound-control problems, $\mathbf{u}^\ell$ and $\bar{\mathbf{u}}$ may be shifted with respect to each other. Third, because we use a least-squares approach to estimate the gradient, the effect of outliers, which may strongly influence the mean value, is reduced in the modified formulation.

Rahul-Mark Fonseca is a PhD degree candidate at Delft University of Technology, working on ensemble-based production optimization and reservoir management. As of 1 September 2015, he will serve as reservoir engineer with TNO. Fonseca holds a master’s degree in petroleum engineering from Delft University of Technology.

Olwijn Leeuwenburgh is a research scientist at TNO. He previously worked for several years in the field of ocean and climate simulation and forecasting with a focus on data assimilation. Leeuwenburgh’s current interests include mathematical aspects of reservoir modeling, especially ensemble methods for assisted history matching and life-cycle optimization. He holds a master’s degree in geodetic engineering from Delft University of Technology and a PhD degree in geophysics from the University of Copenhagen.

Ernesto Della Rossa is a reservoir engineer at Eni. His main interests are reservoir uncertainty quantification and geostatistics, assisted history matching, robust life-cycle optimization with ensemble-based techniques, and surrogate-reservoir modeling. Della Rossa holds a master’s degree in applied mathematics from the University of Milan.

Paul Van den Hof is professor of control systems in the Department of Electrical Engineering at Eindhoven University of Technology. His current research interests are in system identification and model-based control and optimization, with applications in several technology domains, including petroleum-reservoir-engineering systems. Van den Hof holds master’s and PhD degrees from Eindhoven University of Technology.

Jan-Dirk Jansen is professor of reservoir systems and control and head of the Department of Geoscience and Engineering at Delft University of Technology. His current research interests are the use of systems and control theory for production optimization and reservoir management. Previously, Jansen worked for Shell International in research and operations. He holds master’s and PhD degrees from Delft University of Technology.