Partially Separated Meta-models with Evolution Strategies for Well Placement Optimization

Zyed Bouzarkouna, Didier Yu Ding, Anne Auger

Partially Separated Meta-models with Evolution Strategies for Well Placement Optimization

Z. Bouzarkouna, SPE, and D.Y. Ding, SPE, IFP Energies nouvelles, and A. Auger, INRIA.

Abstract

Finding the optimal location of non-conventional wells increases significantly the project's Net Present Value (NPV). This problem is nowadays one of the most challenging problems in oil and gas fields development. When dealing with complex reservoir geology and high reservoir heterogeneities, stochastic optimization methods are the most suitable approaches for optimal well placement. However, these methods require in general a considerable computational effort (in terms of number of reservoir simulations, which are CPU time demanding).

This paper presents the use of the CMA-ES (Covariance Matrix Adaptation – Evolution Strategy) optimizer, which is recognized as one of the most powerful derivative free optimizers, to optimize well locations and trajectories. A local regression based meta-model is incorporated into the optimization process in order to reduce the computational cost. The objective function (e.g., the NPV) can usually be split into local components referring to each of the wells: it depends in general on a smaller number of principal parameters, and thus can be modeled as a partially separable function. In this paper, we propose to exploit the partial separability of the objective function into CMA-ES coupled with meta-models, by building partially separated meta-models. Thus, different meta-models are built for each well or set of wells, which results in a more accurate modeling.

An example is presented. Results show that taking advantage of the partial separability of the objective function leads to a significant decrease in the number of reservoir simulations needed to find the "optimal" well configuration, given a restricted budget of reservoir simulations. This approach is practical and promising when dealing with a large number of wells to be located.

Introduction

Nowadays, the environment and the conditions in which oil and gas fields are discovered are more and more complex and challenging. Both existing fields and new discoveries need an optimal production scheme to be economically viable. One of the most important issues that must be addressed by reservoir engineers to maximize a given project's asset value is to optimally decide where to drill wells. Hence, the well placement optimization problem is nowadays a major focus in the petroleum industry. Many studies already investigated the well placement using different optimization methods: stochastic methods such as genetic algorithms (e.g., Bukhamsin et al., 2010, Emerick et al., 2009, Guyaguler and Horne, 2000, Montes et al., 2001, Morales et al., 2010, Yeten et al., 2003) and deterministic methods in particular adjoint methods (e.g., Sarma & Chen, 2008, Forouzanfar et al., 2010, Zandvliet et al., 2008, Vlemmix et al., 2009).

Stochastic optimization methods were shown to be an appropriate approach for the well placement problem given the multiple local optima and the non-smoothness of the objective function. A promising method is the Covariance Matrix Adaptation – Evolution Strategy (CMA-ES), a state of the art stochastic optimizer first introduced to the well placement problem by Ding (2008) and successfully applied to a well location and trajectory optimization problem by Bouzarkouna et al. (2010). The contribution of CMA-ES was also exhibited through a comparison with a genetic algorithm to conclude that CMA-ES reaches in general a higher net present value (NPV) with less reservoir simulations.

The computational cost of the optimization arising from the high number of reservoir simulations remains the main challenge when applying stochastic optimization methods in practice. In this context, CMA-ES was coupled with a local meta-model in Bouzarkouna et al. (2010). The meta-model aims at replacing some of the expensive true objective function evaluations (requiring

a reservoir simulation) with an approximate but cheap model. This approach can reduce the number of reservoir simulations by 19-25% to reach a satisfactory NPV.

In this paper, we propose a new approach to build the meta-model exploiting the so-called partial separability of the objective function. A function is partially separable if it can be split into a number of functions where each depends on a fewer number of parameters. In this new approach, we propose to build a meta-model for each well or set of wells to be placed, instead of one meta-model for all the wells. This paper is structured as follows. We begin by introducing the well placement optimization problem and its parameterization. Then, we briefly review the CMA-ES approach and its combination with local meta-models as proposed by Bouzarkouna et al. (2010). Afterwards, we define the new approach exploiting the partial separability of the objective function. Exploiting this information allows defining a partially separated approximate model. Finally, we demonstrate the contribution of the proposed approach in reducing the number of reservoir simulations on the synthetic benchmark reservoir case PUNQ-S3.

Problem definition

The problem of determining the optimal well placement is equivalent to maximizing an objective function describing the economics of the project with respect to the well locations and trajectories. The optimization problem can be formulated in a way to handle different possible configurations of multilateral wells. In this paper, we use the same parameterization as defined by Bouzarkouna et al. (2010). The mainbore is defined using 2 points, i.e., 6 parameters representing a straight line. These parameters are: the Cartesian coordinates of the heel (x_h, y_h, z_h) and the spherical coordinates of the toe (l_t, θ_t, φ_t). The laterals are defined using 4 parameters: the distance along the well between the toe and the starting point of the lateral d_l, and the spherical coordinates of the ending point of the lateral (l_l, θ_l, φ_l).

A well placement decision is evaluated according to an objective function. In this paper, the objective function is the net present value (NPV) as defined by Yeten (2003)

$$\text{NPV} = \sum_{n=1}^{Y} \left(\frac{1}{(1+\text{APR})^n} \begin{bmatrix} Q_{n,o} \\ Q_{n,g} \\ Q_{n,w} \end{bmatrix}^T \begin{bmatrix} C_{n,o} \\ C_{n,g} \\ C_{n,w} \end{bmatrix} \right) - C_d \ , \tag{1}$$

where $Q_{n,p}$ is the field production of phase p in the period n obtained using a reservoir simulation, $C_{n,p}$ is the profit or loss associated with this production. A phase p is either oil, gas or water denoted respectively o, g, w. The annual percentage rate is denoted by APR. Y is the number of discount periods (years). C_d is the drilling and completing cost for all wells defined based on the approximate formula used by Yeten (2003)

$$C_d = \sum_{k=0}^{N_{lat}} [\text{A}.d_w.\ln(l_w).l_w]_k + \sum_{m=1}^{N_{jun}} [C_{jun}]_m \ , \tag{2}$$

where $k = 0$ represents the mainbore, $k > 0$ represents the laterals, l_w is the length of the lateral (in ft), d_w is the diameter of the mainbore (in ft) and A is a constant specific to the considered field. C_{jun} is the cost of milling the junction. For this paper, the different constants used to define the NPV are defined in **Table 1**.

It is in general needed to impose different constraints on the well configuration to guarantee finding both undrillable wells and wells that violates common engineering practices. The constraints handled in the present paper are as follows:

-   a maximum length of wells: $l_w < L_{\max}$, $\forall w = 1..n_{wells}$;
-   all wells must be inside the reservoir grid.

Table 1: Constants used to define the Net Present Value (NPV).

Constant	Value
$C_{n,o}$	60 \$ / bbl
$C_{n,w}$	-4 \$ / bbl
$C_{n,g}$	0
APR	0.2
A	1000
d_w	0.1 m
C_{jun}	10^5 \$

The Optimization approach: CMA-ES with meta-models

The Covariance Matrix Adaptation – Evolution Strategy (CMA-ES)
CMA-ES (Hansen & Ostermeier, 2001) is a population based stochastic search algorithm where at each iteration a population of λ points is sampled according to a multivariate normal distribution. The objective function is then evaluated for each point of the population and the parameters of the multivariate normal distribution are updated according to the feedback obtained on the objective function, so as to reinforce the probability of finding good points in the next iteration.

Let $\mathbf{m}$ be the mean of the multivariate normal distribution, and σ and $\mathbf{C}$ respectively the step-size and the covariance matrix. Assume that $\mathbf{m}$, σ and $\mathbf{C}$ are given, the new population $(\mathbf{x}_i)_{1\leq i\leq\lambda}$ is sampled according to:

$$\mathbf{x}_i = \mathbf{m} + \sigma\,\aleph_i(0,\mathbf{C}) \ , \qquad \forall\, i = 1,...,\lambda \ , \tag{3}$$

where $(\aleph_i(0,\mathbf{C}))_{1\leq i\leq\lambda}$ are λ independent multivariate normal distributions with zero mean vector and covariance matrix $\mathbf{C}$. A more detailed overview of CMA-ES can be found in Bouzarkouna et al. (2010).
We refer to Hansen & Kern (2004) for the equation updates of the parameters of the multivariate normal distribution $\mathbf{m}$, σ and $\mathbf{C}$. We note that updates rely only on the ranking of the μ (in general $=\lambda / 2$) best points of the population.

CMA-ES with meta-models: lmm-CMA
In order to reduce the number of reservoir simulations, CMA-ES is coupled to local meta-models. A meta-model is an approximate model built based on the true objective function evaluations and used during the optimization process to save evaluations on the expensive objective function.

To build an approximate model of the objective function f: $\mathrm{IR}^n \to \mathrm{IR}$, denoted $\hat{f}$, we use a locally weighted regression. During the optimization, a training set is built by storing, after every evaluation on the true objective function, points together with their objective function values. Assuming that the training set contains a sufficient number m of elements, for each point denoted now by $\mathbf{q} \in \mathrm{IR}^n$, we select the k nearest points to $\mathbf{q}$ in the training set $(\mathbf{x}_i)_{1\leq i\leq k}$ using the Mahalanobis distance d with respect to the current covariance matrix $\mathbf{C}$ defined for a given point $\mathbf{z} \in \mathrm{IR}^n$ as $d(\mathbf{z},\mathbf{q}) = \sqrt{(\mathbf{z}-\mathbf{q})^T\mathbf{C}^{-1}(\mathbf{z}-\mathbf{q})}$. Locally weighted regression builds an approximate objective function using (true) evaluations $(y_i)_{1\leq i\leq k}$ stored in the training set and corresponding to the k selected nearest points to $\mathbf{q}$. Kern et al. (2006) suggested the use of a full quadratic meta-model. Hence, using a vector $\beta \in IR^{\frac{n(n+3)}{2}+1}$, we define $\hat{f}$ as $\hat{f}(\mathbf{z},\beta) = \beta^T(z_1^2,...,z_n^2,z_1z_2,...,z_{n-1}z_n,z_1,...,z_n,1)\ \forall \mathbf{z} = (z_1,...,z_n)\in IR^n$. The vector of parameters β is found by minimizing the following criteria with respect to β :

$$A(\mathbf{q}) = \sum_{i=1}^{k}[(\hat{f}(\mathbf{x}_i,\beta) - y_i)^2 K(\frac{d(\mathbf{x}_i,\mathbf{q})}{h})] \ . \tag{4}$$

The kernel weighting function $K(.)$ is defined by $K(\zeta)=(1-\zeta^2)^2$, and h is the bandwidth defined by the distance of the k^{th} nearest neighbor data point to $\mathbf{q}$. In order to build the full quadratic meta-model, k must be greater or equal to $\frac{n(n+3)}{2}+1$.

Le us consider now an approximate model $\hat{f}$ of the objective function f. In order to incorporate the approximate model, we use the approximate ranking procedure (Runarsson, 2004). This procedure aims at deciding at every iteration which points must be evaluated on the true objective function, i.e., using a reservoir simulation, and which one can be approximated with the meta-model. In the following, we recall the variant used in Bouzarkouna et al. (2010), which will be used in this paper. For a given iteration, let us denote the population of CMA-ES to be evaluated by $(\mathbf{x}_i)_{i=1,...,\lambda}$, where each $(\mathbf{x}_i)$ represents a well configuration. The following procedure is then performed:

1. approximate $\hat{f}(\mathbf{x}_i)$ for all well configurations of the current population.
2. determine the ensemble of the μ best well configurations and the best well configuration.
3. evaluate the best well configuration with the true objective function and add its objective function value to the training set.
4. for n_{ic} from 1 to $(\lambda - 1)$, we:
   a. approximate $\hat{f}(\mathbf{x}_i)$ for all well configurations of the current population.

b. determine the ensemble of the μ best well configurations and the best well configuration.
c. if less than one fourth of the population is evaluated, the meta-model is accepted if it succeeds in keeping both **the best well configuration** and **the ensemble of the (λ / 2) best well configurations** unchanged.
d. if more than one fourth of the population is evaluated, the meta-model is accepted if it succeeds in keeping **the best well configuration** unchanged.
e. if the meta-model is accepted, we break. If not, we evaluate the best unevaluated well configuration with the true objective function, add its objective function value to the training set, and loop to step 4, until reaching the acceptance criterion of the meta-model.

CMA-ES with partially separated meta-models: p-sep lmm-CMA

It is clear that the number of reservoir simulations performed in each iteration depends on the quality of the used approximate model. The more accurate the approximate model is, the easier it is accepted by the approximate ranking procedure and hence the less reservoir simulations are performed. The contribution of this paper is to define a new way to build the approximate model exploiting a specific structure of the objective function. As shown in the previous section, a meta-model can be built using a locally weighted regression to approximate the whole objective function. Although this approach is shown to lead to important savings in the number of reservoir simulations (Bouzarkouna et al., 2010), expensive real field applications remain demanding in terms of reducing the number of reservoir simulations.

The objective function describing the well placement decision, e.g., the NPV is modeled in this work as a partially separable function. A function f is said to be partially separable if it can be written as a sum of simple functions called *element functions*, where each element function depends only on a few variables called *element variables*.
Given the Eqs. 1 and 2, we note that the field production of a phase p in a period n denoted above by $Q_{n,p}$ is equal to the sum of the productions corresponding to the phase p in the period n for each well w denoted by $(Q_{n,p})_w$, i.e., $Q_{n,p} = \sum_{w=1}^{N_w} (Q_{n,p})_w$, where N_w is the number of wells. Hence, the NPV of the field can be written as a sum of simpler functions such the NPVs for each well, or the NPVs for a number of sets of wells. In what follows, we denote the objective function to optimize by $f : \mathrm{IR}^n \to \mathrm{IR}$, and we suppose that it can be written as a sum of nf element function $(f_i)_{i=1..nf}$, where each f_i depends on n_i parameters. For notation purposes, we define for every f_i a simple explicit function $\Phi^i : \mathrm{IR}^n \to \mathrm{IR}^{n_i}$, such as f_i depends only on $\Phi^i(\mathbf{x})$ instead of the whole $\mathbf{x}$.

$$f(\mathbf{x}) = \sum_{i=1}^{nf} f_i(\Phi^i(\mathbf{x})) \ . \tag{5}$$

We suppose that each element function f_i depends on a number n_i $(< n)$ of parameters. In general, for some wells or set of wells, we can discard some non-influencing parameters or with a negligible influence.
In this notation, we introduce the function Φ^i in order to take into consideration not only element functions depending on some parameters of the objective function but also element functions depending on a linear or a non-linear combination of these parameters.

To illustrate the definition of Φ^i , let us consider an objective function denoted $g = g_1 + g_2$, where g_1 and g_2 are the element functions corresponding to g. We define g, g_1 and g_2 for a given $\mathbf{p} = (p_1, p_2, p_3, p_4) \in \mathrm{IR}^4$ as follows:

$$g : \mathrm{IR}^4 \to \mathrm{IR} \qquad\qquad g_1 : \mathrm{IR}^3 \to \mathrm{IR} \qquad\qquad g_2 : \mathrm{IR}^2 \to \mathrm{IR}$$
$$\mathbf{p} = (p_1, p_2, p_3, p_4) \mapsto g(\mathbf{p}) \qquad \tilde{\mathbf{p}} = (p_1, p_2, \sqrt{(p_3 - p_1)^2 + (p_4 - p_2)^2}) \mapsto g_1(\tilde{\mathbf{p}}) \qquad \overline{\mathbf{p}} = (p_3, p_4) \mapsto g_2(\overline{\mathbf{p}})$$

In this example, by defining Φ^1 and Φ^2 such as:

$$\Phi^1 : \mathrm{IR}^4 \to \mathrm{IR}^3 \qquad\qquad\qquad\qquad\qquad \Phi^2 : \mathrm{IR}^4 \to \mathrm{IR}^2$$
$$\mathbf{p} = (p_1, p_2, p_3, p_4) \mapsto (p_1, p_2, \sqrt{(p_3 - p_1)^2 + (p_4 - p_2)^2}) \qquad \mathbf{p} = (p_1, p_2, p_3, p_4) \mapsto (p_3, p_4)$$

we notice that $g(\mathbf{x}) = g_1(\Phi^1(\mathbf{x})) + g_2(\Phi^2(\mathbf{x}))$.

In this context, we propose to build an approximate model $\hat{f}_i$ for each element function f_i. The training set is built by storing for every evaluated point $\mathbf{x}$, both $\Phi^i(\mathbf{x})$ and their corresponding values on f_i, i.e., $f_i(\Phi^i(\mathbf{x}))$. Let us consider a point $\mathbf{q}$ to be

evaluated on $\hat{f}_i$. Assuming that the training set contains a sufficient number of elements, we select $k_i \in \text{IN}$ nearest points ($\Phi^i(\mathbf{x}_j)$, $j = 1...k_i$) to $\Phi^i(\mathbf{q})$ using the Mahalanobis distance d_i with respect to a matrix $\mathbf{C}_i$, defined for a given point $\mathbf{z} \in \text{IR}^n$ as

$$d_i(\mathbf{z},\mathbf{q}) = d_i(\Phi^i(\mathbf{z}),\Phi^i(\mathbf{q}))$$
$$= \sqrt{(\Phi^i(\mathbf{z}) - \Phi^i(\mathbf{q}))^T \mathbf{C}_i^{-1} (\Phi^i(\mathbf{z}) - \Phi^i(\mathbf{q}))} \quad , \tag{6}$$

where $\mathbf{C}_i$ is an n_i x n_i matrix adapted to the local shape of f_i landscape. The definition of $\mathbf{C}_i$ is borrowed from the covariance matrix update equation, defined by CMA-ES, which is by construction adapted to the local shape of the fitness function. Hence, we use the update equations but with a population and a corresponding objective function values defined using $\left(\Phi^i(\mathbf{X}_m), f_i\left(\Phi^i(\mathbf{X}_m)\right)\right)_{m=1...\lambda}$ instead of $(\mathbf{X}_m, f(\mathbf{X}_m))_{m=1...\lambda}$ where λ is the population size. We refer to Hansen & Kern (2004) for detailed update equations of the covariance matrix.

To build the meta-model, we use a full quadratic model. Using a vector $\boldsymbol{\beta}_i \in \text{IR}^{\frac{n_i(n_i+3)}{2}+1}$, $\hat{f}_i$ is defined for a given point $\mathbf{z} \in \text{IR}^n$, for which we denote $\Phi^i(\mathbf{z}) = (\tilde{u}_1,...,\tilde{u}_{n_i})$, as

$$\hat{f}_i(\mathbf{z},\boldsymbol{\beta}_i) = \hat{f}_i(\Phi^i(\mathbf{z}),\boldsymbol{\beta}_i) = \boldsymbol{\beta}_i^T \tilde{\mathbf{z}}_i \quad , \tag{7}$$

where $\tilde{\mathbf{z}}_i = (\tilde{u}_1^2,...,\tilde{u}_{n_i}^2, \tilde{u}_1\tilde{u}_2,...,\tilde{u}_{n_i-1}\tilde{u}_{n_i}, \tilde{u}_1,...,\tilde{u}_{n_i},1)$. The full quadratic meta-model is built by minimizing with respect to $\boldsymbol{\beta}_i$ the following criterion:

$$B(\mathbf{q}) = \sum_{j=1}^{k_i}\left[\left(\hat{f}_i(\Phi^i(\mathbf{x}_j),\boldsymbol{\beta}_i) - f_i(\Phi^i(\mathbf{x}_j))\right)^2 K\left(\frac{d_i(\Phi^i(\mathbf{x}_j),\Phi^i(\mathbf{q}))}{h}\right)\right] \quad , \tag{8}$$

where $K(.)$ is the Kernel weighting function defined by $K(\zeta)=(1-\zeta^2)^2$, and h is the bandwidth defined by the distance d_i of the k_i^{th} nearest neighbor point to $\mathbf{q}$.

We note that k_i must be greater or equal to $\frac{n_i(n_i+3)}{2}+1$ to build the full quadratic meta-model. We recommend to choose

$$k_i = n_i(n_i+3)+2 \quad . \tag{9}$$

Hence, for each point, meta-models are built for all the element functions. The approximate function of the objective function f can be calculated as follows:

$$\hat{f}(\mathbf{x}) = \sum_{i=1}^{nf} \hat{f}_i(\Phi^i(\mathbf{x})) \quad . \tag{10}$$

To incorporate the built meta-model $\hat{f}$ into CMA-ES, we use the approximate ranking procedure as described for lmm-CMA with only one difference related to the acceptance criterion of the meta-model: for p-sep lmm-CMA, we use a less conservative criterion in which the meta-model is accepted if it succeeds in keeping only the best well configuration unchanged.

Application on the PUNQ-S3 reservoir case

This section shows an application of the proposed method (p-sep lmm-CMA) on the benchmark reservoir case PUNQ-S3. This application is compared to the CMA-ES optimizer and to the variant of CMA-ES with meta-models (lmm-CMA).

The PUNQ-S3 case has been taken from a reservoir engineering study on a real field. The model contains 19×28×5 grid blocks. The elevation of the field is shown in **Fig. 1**. An injection well denoted I1 is already drilled. **Fig. 2** represents the distribution of $\sum_{l=1}^{n\text{ layers}}(H_k \times \phi \times S_o)$, where H_k is the gross thickness of the layer l, S_o is the oil saturation and ϕ is the porosity. The location of I1 is also shown in **Fig. 2**, where I1 is an inclined well drilled in the layer 3.

We propose to drill 3 unilateral producers (denoted P1, P2 and P3) to maximize the NPV. The dimension of the problem is then equal to: 6×3 = 18. A producer limit bottomhole pressure is fixed to 150 bar, and an injector limit bottomhole pressure is fixed to 320 bar. A maximum length of 1000 m is imposed on the 3 producers to be drilled.

The population size λ is set, for all the methods used, to 60. The different optimizers are run with a stopping criterion corresponding to a maximum number of reservoir simulations equal to 1000. Other parameters of the optimization method were defined according to default settings (Hansen & Kern, 2004).

We perform 10 runs for each method to show the average performance measured by the number of reservoir simulations required to reach a given value of NPV.

For p-sep lmm-CMA, we should define the different element functions and their corresponding dependencies. In this application, a meta-model is built for each well (already drilled and to be drilled) approximating its NPV. We denote by NPV_1, NPV_2 and NPV_3, the NPVs corresponding respectively to P1, P2 and P3. The NPV corresponding to I1 is denoted NPV_4. Therefore, the objective function corresponding to the NPV of the field is equal to the sum of the different element functions corresponding to the NPV of each well, i.e., $(NPV_i)_{i=1...4}$.

The goal of exploiting the partial separability of the objective function is to reduce the dimension of each element function, and therefore to allow a better accuracy of the meta-model. Some approximations should be made to define the dependencies of each element function.

For each well denoted by i, we define the following parameters:

- dp_i: the minimum distance between the well and the other producers.
- di_i: the minimum distance between the well and the injector.

In this application, we suppose that the NPVs of the wells to be drilled, i.e. $(NPV_i)_{i=1...3}$ depend on the parameters defining the location and trajectory of the considered well and on its corresponding dp_i and di_i.

The NPV of the well already drilled, i.e., NPV_4 depends only on 1 parameter: the dp_i related to the considered well.

Therefore, each $(NPV_i)_{i=1...3}$ depends only on 8 parameters. NPV_4 depends on 1 parameter.

The number of points used to build the partially separated meta-models, is chosen to be equal to 90 (according to Eq. 9), and the meta-model is used when the training set (storing the performed evaluations) contains at least 150 elements, i.e., before performing 150 simulations, all the points are evaluated with the true objective function and the partially separated meta-model is not used.

Fig. 3 shows the average performance of the proposed method, i.e., CMA-ES with partially separated meta-models (p-sep lmm-CMA). Results are reported together with those obtained using CMA-ES and CMA-ES with meta-models (lmm-CMA).

During the first iterations of the optimization, the performance of the 3 used algorithms is equivalent. For p-sep lmm-CMA, the meta-model is used if the training set contains at least 150 performed reservoir simulation results. Therefore, in the beginning of the optimization, the meta-model is not used which justifies the equivalent results for the three optimizers.

For lmm-CMA, building the meta-model requires more reservoir simulations compared to partially separated meta-models. Non-partially separated meta-models depend on 18 parameters. In the performed runs, the meta-model is built using 300 performed reservoir simulations ($k = 300$) and used when the training set contains at least 350 objective function evaluations. Hence, before reaching 350 simulations, lmm-CMA and CMA-ES are equivalent.

Except in the beginning of the optimization in which all the optimizers are equivalent, it is clear that CMA-ES with partially separated meta-models outperforms the other methods, when considering a restricted budget of 1000 reservoir simulations. The context of restricted budget of simulations is imposed to consider real applications in which the number of simulations is generally limited to several hundreds to several thousands, given the CPU time required by a simulation.

For a given number of reservoir simulations equal to 600, p-sep lmm-CMA is able to find a well configuration with an NPV equal to 1.26×10^{10} \$. However, CMA-ES reaches only an NPV equal to 1.17×10^{10} \$ and lmm-CMA offers only a maximum NPV equal to 1.21×10^{10} \$. As a conclusion, using a restricted budget of reservoir simulations, exploiting the partial separability allows reaching greater NPV values compared to CMA-ES and lmm-CMA.

To reach a value of NPV equal to 1.20×10^{10} \$, CMA-ES with partially separated meta-models requires 370 reservoir simulations. However, to reach the same value of NPV, using standard meta-models requires 510 reservoir simulations, and when using CMA-ES without meta-models, we need 930 reservoir simulations. Therefore, using partially separated meta-models saves 60% of the number of reservoir simulations compared to CMA-ES (without meta-models). The contribution of exploiting the partial separability is shown when comparing p-sep lmm-CMA with lmm-CMA. Exploiting the partial separability of the objective function saves 28% of the number of reservoir simulations compared to CMA-ES with standard meta-models approach.

Fig. 4 shows one of the obtained solution well configurations, with an NPV value equal to 1.38×10^{10} \$. Although, each of the performed runs proposes in general a different solution, the majority of the solution well configurations are located in the same regions.

Conclusions

This paper introduces a new approach to optimize well locations and trajectories with a reduced number of reservoir simulations. The proposed approach is based on a new variant of the stochastic optimizer CMA-ES combined with meta-models exploiting the problem structure, i.e., the partial separability of the objective function.

Since the objective function can be split into separate element functions corresponding to each well or set of wells, we build a separate meta-model for each element function. The performance of the proposed approach is demonstrated on the synthetic benchmark reservoir case PUNQ-S3. Results show that the new approach leads to an important reduction of the number of reservoir simulations (around 60%) compared to the optimizer CMA-ES. The important savings in the number of reservoir simulations are justified by the reduced number of parameters required to build the meta-model of the element functions. Further studies are useful in order to define the best parameters for each element function and thus improve the accuracy of the partially separated meta-models.
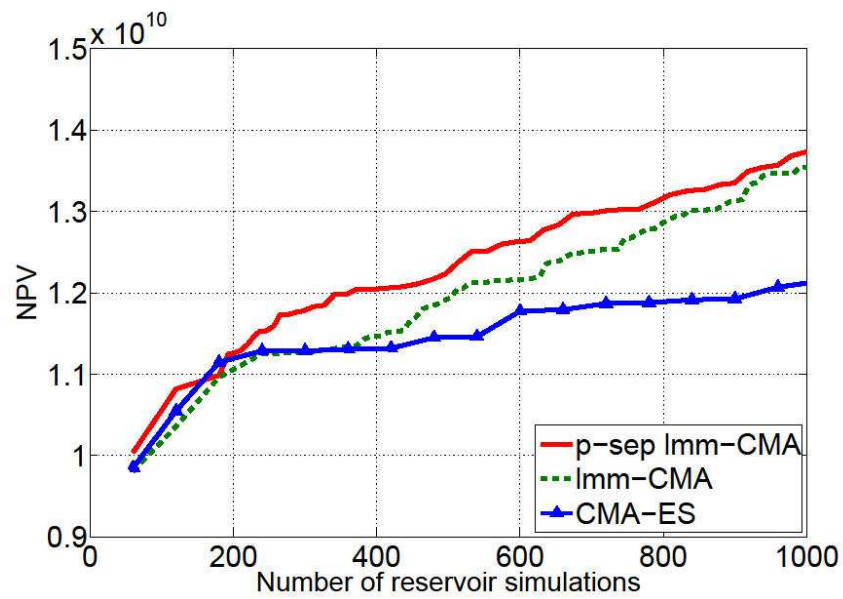
The proposed approach exploiting the partial separability of the objective function can also be combined with any other stochastic optimizer, in order to reduce the computational cost of the optimization.

References

Bouzarkouna, Z., Ding, D. & Auger, A. (2010). Using Evolution Strategy with Meta-models for Well Placement Optimization. In: Proceedings of the 12th European Conference on the Mathematics of Oil Recovery ECMOR XII.

Bukhamsin, A., Farshi, M. & Aziz, K. (2010). Optimization of multilateral well design and location in a real field using a continuous genetic algorithm. In: SPE/DS Saudi Arabia Section Technical Symposium and Exhibition, SPE 136944

Ding, D. (2008). Optimization of well placement using evolutionary algorithms. In: SPE EUROPEC/EAGE annual conference and exhibition, SPE 113525

Emerick, A., Silva, E., Messer, B., Almeida, L., Szwarcman, D., Pacheco, M. & Vellasco, M. (2009). Well placement optimization using a genetic algorithm with nonlinear constraints. In: SPE reservoir simulation symposium, SPE 118808

Forouzanfar, F., Li, G. & Reynolds, C. (2010). A two-stage well placement optimization method based on adjoint gradient. In: SPE annual technical conference and exhibition, SPE 135304

Guyaguler, B. & Horne, R. (2000). Optimization of well placement. Journal of Energy Resources Technology **122**(2), 64–70

Hansen, N. & Kern, S. (2004). Evaluating the CMA evolution strategy on multimodal test functions. In: X. Yao et al., editors, Parallel Problem Solving from Nature - PPSN VIII, *Lecture Notes in Computer Science*, vol. 3242, pp. 282–291. Springer Berlin / Heidelberg

Kern, S., Hansen, N., Koumoutsakos, P. (2006). Local meta-models for optimization using evolution strategies. In: T. Runarsson et al., editors, Parallel Problem Solving from Nature – PPSN IX, *Lecture Notes in Computer Science*, vol. 4193, pp. 939–948. Springer Berlin / Heidelberg

Montes, G., Bartolome, P. & Udias, A. (2001). The use of genetic algorithms in well placement optimization. In: SPE Latin American and Caribbean Petroleum Engineering Conference, SPE 69439

Morales, A., Nasrabadi, H. & Zhu, D. (2010). A modified genetic algorithm for horizontal well placement optimization in gas condensate reservoirs. In: SPE annual technical conference and exhibition, SPE 135182

Runarsson, T. (2004). Constrained evolutionary optimization by approximate ranking and surrogate models. In: X. Yao et al., editors, Parallel Problem Solving from Nature - PPSN VIII, *Lecture Notes in Computer Science*, vol. 3242, pp. 401–410. Springer Berlin / Heidelberg

Sarma, P. & Chen, W. (2008). Efficient well placement optimization with gradient-based algorithms and adjoint models. In: SPE intelligent energy conference and exhibition, SPE 112257

Vlemmix, S., Joosten, G., Brouwer, D. & Jansen, J. (2009). Adjoint-based well trajectory optimization in a thin oil rim. In: SPE EUROPEC/EAGE annual conference and exhibition, SPE 121891

Yeten, B., Durlofsky, L. & Aziz, K. (2003). Optimization of nonconventional well type, location and trajectory. SPE Journal **8**, 200–210

Zandvliet, M., Handels, M., van Essen, G., Brouwer, D. & Jansen, J. (2008). Adjoint-based well-placement optimization under production constraints. SPE Journal **13**(4), 392–399

Fig. 1: PUNQ-S3 test case.

Fig. 2: Map of $\sum_{l=1}^{n\ layers}(H_k \times \phi \times S_o)$ with the location of the injector already drilled I1.

Fig. 3: The mean value of NPV (in US$) and its corresponding standard deviation for well placement optimization using CMA-ES with partially separated meta-models denoted p-sep lmm-CMA (——), CMA-ES with meta-models denoted lmm-CMA (- -) and CMA-ES (▲). 10 runs are performed for each method.
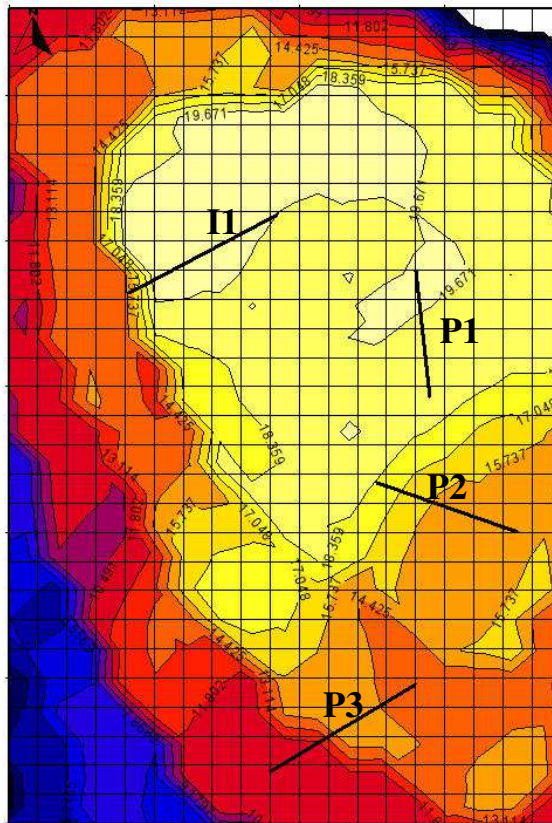
Fig. 4: Map of $\sum_{l=1}^{n\,layers}(H_k \times \phi \times S_o)$ with the location of the injector already drilled I1, and solution producers (P1, P2 and P3).