CS238/AA228 Final Project Report: Decision-Making in the Context of Unconventional Oil Well Drilling Operations

Dante I. Orta Alemán and Clothilde M.A. Venereau
Stanford University - Energy and Resources Engineering Department

We study decision-making in the context of unconventional oil well drilling, under uncertainty in oil prices and levels of production generated at each well. Given a set of potential well locations, we use a Markov Decision Process formulation to find an optimal policy to sequentially drill wells, with the aim of maximizing profits from oil production to repay initial investment in the project. We use the value iteration algorithm on a discretized state space to solve for the optimal policy. The obtained optimal policy specifies the action to take at every individual state such that a reward function including costs and revenues from oil production and repayment of initial investment is maximized. We find that whenever the action of drilling was still feasible, the optimal policy suggested to drill a new well to increase oil production and therefore revenues. Whenever the maximum possible number of wells drilled was reached, the policy recommended the action of not drilling in the majority of cases. These results are aligned with our prior expectations of how to make decisions in drilling operations.

I. Introduction

We aim to study the profitability of an unconventional/tight oil company in its drilling operations. Unconventional oil production through fracking differs from conventional oil projects due to a variety of factors. An important difference is the very high decline rates in petroleum production over time associated with tight oil resources due to fundamental geological differences. Oil companies aim to make as high a profit as possible by maximizing oil production. The profits recovered from oil production must be balanced against expenditures which include operating costs, well drilling costs and acreage costs. Costs in oil exploration and production are substantial, and for this reason tight oil companies will finance operations using a balance of both debt and equity (termed leverage). Higher leverage is required in the early stages of production to build out the requisite infrastructure and technology. This debt can be repayed over time if the company is capable of sustaining production with a high enough oil price by drilling oil producing wells. The goal of this analysis is to come up with an optimal policy to guide the drilling operations of an unconventional oil company in a scenario of oil price uncertainty and uncertainty in the decline rates of producing wells. We simplify the financing structure for the oil production operations to only focus on generating sufficient revenues to pay back with a return the initial investment in the project.

II. Literature Review

The problem we are aiming to solve fits within the general Field Development Planning (FDP) optimization problem which has been studied using a variety of techniques. FDP optimization problems have been approached in the literature traditionally using either (i) black-box optimization techniques based on objective function evaluations through an oil reservoir simulator, (ii) stochastic global search methods (genetic algorithms, particle swarm optimization) or (iii) local optimization methods (pattern search or gradient based techniques) to find the optimal well drilling locations in a given field [1, 2]. In FDP, one of the most important decision to make is well placement. Well placement problems typically involve optimizing over parameters corresponding to positions and orientations of injector and producer wells [1]. The objective function usually corresponds to maximizing the total amount of oil recovered or the Net Present Value (NPV) of that oil extracted. The state of the art optimization method used for the well placement is mixed-integer nonlinear programming, which can be impractical in real reservoir applications due to the many objective function evaluations using a reservoir simulator over a set of geostatistical realizations of a reservoir's petrophysical properties [e.g., 3–6]. Rodriguez-Torrado et al. [6] first took into account the sequential nature of well placement and formulated the well placement optimization problem as a Partially Observable Markov Decision Process (POMDP). In their model, the information collected after drilling each well helps to decide where to drill subsequent wells by providing new constraints on spatial correlations in the geostatistical model representing the reservoir's physical properties. Making use of this property, Rodriguez-Torrado et al. (2017) [6] propose a dynamic programming approach to finding an optimal well placement policy.

However, most of the literature on FDP has been done for conventional petroleum reservoirs. Unconventional or 'tight' reservoirs are usually developed without the use of a reservoir model. On the contrary, the wells are placed in a equally spaced fashion and there is a poor understanding of the geological conditions that make an unconventional well successful [7]. The metric used for assessing the NPV of a tight oil field is an estimate of the Estimated Ultimate Recovery (EUR), that is the cumulative oil extracted from the wells during the life of the reservoir. Traditionally, the EUR is estimated using decline curve analysis and some research has been done using Markov Chain Monte Carlo methods to account for uncertainties in such analysis [8]. These factors make it infeasible to apply state of the art optimization methods presented above for conventional oil reservoir to tight oil cases. The problem of Field Development Planning in unconventional reservoirs has received increased attention in recent months as several companies filed for bankruptcy as fields were developed using flawed EUR estimates which led to unprofitable operations [9].

III. Methodology

A. Problem Statement

Following a similar approach as Torrado et al. (2017) [6], we consider the well placement optimization problem for an unconventional oil field as sequential in nature and we formulate it as a Markov Decision Process (MDP). We provide a detailed description below of our simplified approach to this problem.

It is assumed that the company owns a gridded field, with each grid block having the capacity to hold one well. We assume 3 potential well locations. A decline curve approximation for oil production of the field O_t at time t can be modeled as:

$$O_t = \frac{Q_i}{(1 + bDt)^{1/b}} \tag{1}$$

Where Q_i is the initial production rate while b and D are set decline rate constants specific to the field. Our company will start with an amount I_{init} of cash invested. At each decision step, the company needs to choose to drill or not drill a well at one of the three potential locations. Drilling a well will generate revenues from the oil produced and sold at a fluctuating oil price. Drilling a well will incur a cost $C_{drilling}$ associated with the drilling operation as well as operational costs $OPEX$ that are proportional to the production levels reached. Our company generates revenues by selling the ultimate oil recovered over time at an oil price P_{O_t}.

We will treat the problem as having a finite horizon with three decision epochs. At each decision step t, we need to decide where to drill given our three potential locations to maximize our revenues over the problem horizon.

The components of the MDP (state space, action space, transition and reward models) for our problem are presented in sections III.A.1-4 below.

1. State space:

The state space is defined by three variables: Well ID, estimated ultimate recovery (EUR), and oil price P_{O_t}:

$$\mathcal{S} = \{\text{Well ID}, EUR, P_{O_t}\}$$

The EUR corresponds to the amount of oil expected to be recovered from a well by the end of its producing life. We assume a constant producing life of 5 years for all wells, which is a typical duration of production operations for unconventional oil wells. For each well, the EUR can be calculated as:

$$EUR_{well} = \sum_{t=1}^{5} O_{t,well} \tag{2}$$

with O_t defined in equation (1). Using the EUR definition for calculation of our revenues allows us to account for the fact that oil production decreases with time according to the decline curve approximation described by equation (1).

We discretize our state space as follows.

- The Well ID represents the three potential well locations that are characterized by a binary state variable: $\{\text{drilled, not drilled}\} = \{0, 1\}$.

- We approximate the oil production as the ultimate oil recovery EUR which can take on three different values corresponding to a low, medium and high recovery according to the corresponding Q_i (equations (1), (2)). The production can only be non zero when the well has been drilled. We model the initial well production level as $Q_i \sim \mathcal{N}(\mu_w, \sigma_w)$ (parameters for probability distributions defined in section III.B. below).
- We take into account uncertainty in oil prices by having a stochastic price at each decision epoch. We model the oil price P_{O_t} as a Markov Chain specified by $P_{o_t} \sim \mathcal{N}(P_{o_t-1}, \sigma_o)$.

Examples of states in our state space can be represented as the following matrices where the first column of each matrix represents the Well ID, the second column the EUR and the third the price of oil P_{o_t}. Each row of the state matrices represents a well location.

$$\begin{bmatrix} 0 & 0 & p_{O_0} \\ 0 & 0 & p_{O_0} \\ 0 & 0 & p_{O_0} \end{bmatrix}, \begin{bmatrix} 1 & EUR_1 & p_{O_1} \\ 0 & 0 & p_{O_1} \\ 0 & 0 & p_{O_1} \end{bmatrix}, \begin{bmatrix} 1 & EUR_1 & p_{O_1} \\ 1 & EUR_2 & p_{O_2} \\ 0 & 0 & p_{O_2} \end{bmatrix}, \begin{bmatrix} 1 & EUR_1 & p_{O_1} \\ 1 & EUR_2 & p_{O_2} \\ 1 & EUR_3 & p_{O_3} \end{bmatrix}$$

2. Action space:

The action space $\mathcal{A}$ contains four actions corresponding to drilling a well (corresponding to binary value $\{1\}$) at one of the three locations, or not drilling (corresponding to binary value $\{0\}$). We assume that we cannot drill more than one well per decision step.

$$\mathcal{A} = \left\{ \begin{bmatrix} 0 \\ 0 \\ 1 \end{bmatrix}, \begin{bmatrix} 0 \\ 1 \\ 0 \end{bmatrix}, \begin{bmatrix} 1 \\ 0 \\ 1 \end{bmatrix}, \begin{bmatrix} 0 \\ 0 \\ 0 \end{bmatrix} \right\}$$

3. Reward function:

The reward $\mathcal{R}(s, a)$ is modeled as the profit the company makes from the cumulative oil produced minus the required return on initial investment:

$$\mathcal{R}(s, a) = \sum_{wells} (P_{O_t, well} EUR_{well}) - \sum_{wells} (EUR_{well} OPEX + C_{drilling}) - (1 + ROI) I_{init} \tag{3}$$

The costs associated with drilling a well are the initial drilling costs $C_{drilling}$ plus the operating costs $OPEX$ which are proportional to the EUR for a given well. As a simplifying assumption, the revenue for a well is the EUR times the price of oil at the time of drilling. This is equivalent to having a contracting scheme with a fixed oil price for each well. I_{init} corresponds to the initial investment that was made into our company. ROI is the return on investment required by shareholders . Thus, $\mathcal{R}(s, a)$ is only positive once the initial investment and desired return have been paid back, and negative otherwise.

4. Transition model:

Our state space contains three variables: $\mathcal{S} = \{\text{Well ID}, EUR, P_{o_t}\}$. We describe the transitions involved for each of the state variables given different actions below.

Transitions for the Well ID state variable: The most straightforward part of our transition model is: if we take the action to drill at a particular location, the first variable in our state space "Well ID" representing whether a well is drilled or not will transition from 0 to 1 with a probability of 1. Once a well has been drilled at one of the three locations, we cannot perform the action of drilling there again. Once a well has been drilled it cannot be "undrilled".

Transitions for the EUR state variable: If we drill a well i.e. Well ID going from 0 to 1, the oil production EUR for that well has to be updated to become non zero. Drilling a well results in an initial production Q_i distributed as $\mathcal{N}(\mu_w, \sigma_w)$. We discretize Q_i with $n = 3$ equally spaced bins between $\{0.4Q_{init}, 1.2Q_{init}\}$ with Q_{init} a fixed estimated initial production flow rate (see section III.B. Model Parameterization). Given the specified distribution, the probabilities of each oil production level are: $\{P(Q_i = 11{,}964) = 0.15, P(Q_i = 17{,}946) = 0.47, P(Q_i = 23{,}928) = 0.32\}$ where Q_i is measured in barrels of oil. Drilling a well will also incur a cost $C_{drilling}$. We assume identical lifetimes of 5 years for the three wells to calculate the EUR from the corresponding Q_i using equations (1) and (2). The EUR takes into account decline in production over time in a similar way for all wells.

Transitions for P_{O_t} state variable: For the wells that have not been drilled, the oil price also updates at each decision epoch. As previously described $P_{o_t} \sim \mathcal{N}(P_{o_t-1}, \sigma_o)$. Once a well has been drilled, its production and oil price are fixed. Drilling new wells will not change the oil price of wells that have already been drilled. This idea can be illustrated by the following example. At $t = 1$, we start with well at location 1 drilled, with production EUR_1 with associated oil price P_{O_1}. We now take the action to drill a well at location 3 which will result in production EUR_2 and price P_{O_2}. At this second time step $t = 2$, the price of oil for well location 1 will remain P_{O_1}. The price of oil for the location 2 that has not been drilled yet becomes P_{O_2}.

$$s = \begin{bmatrix} 1 & EUR_1 & p_{O_1} \\ 0 & 0 & p_{O_1} \\ 0 & 0 & p_{O_1} \end{bmatrix}, a = \begin{bmatrix} 0 \\ 0 \\ 1 \end{bmatrix}, s' = \begin{bmatrix} 1 & EUR_1 & p_{O_1} \\ 0 & 0 & p_{O_2} \\ 1 & EUR_2 & p_{O_2} \end{bmatrix}$$

To simplify our transition model, we create a function that eliminates invalid states in our state space. Invalid states include for example states that have no drilled wells but a production larger than zero. This reduces our state space size by a factor of ~ 10.

B. Model Parameterization

Table 1. describes the fixed parameters used in our model.

Table 1 Fixed parameters used in the modelling

Parameter	Value	Description	Source
b	1.071	Oil production decline exponent	Calculated*
D	0.3898	Oil production decline term	Calculated*
I_{init}	$ 15,000,000	Initial investment into our company	Estimated
$OPEX$	$7/boe	Well drilling operating costs	[10]
Q_{init}	22,433 [bbl/month]	Baseline for initial oil production	Calculated*
μ_w	22,433 [bbl/month]	Mean Q_i distribution	Calculated*
σ_w	$0.1Q_{init}$	Standard deviation Q_i distribution	Estimated
σ_o	5	Standard deviation oil price update distribution	Estimated
$C_{drilling}$	$ 6,000,000	Cost of drilling a well	[11]
ROI	20%	Desired return on initial investment	Estimated

C. Algorithm Description

We choose to use value iteration to solve our problem given our discretized state space. Value iteration uses the Bellman equation to find the optimal value function as:

$$U^*(s) = \max_a (R(s,a) + \gamma \sum_{s'} T(s'|s,a) U^*(s')) \tag{4}$$

Where U^* is the optimal value function, $R(s,a)$ the reward given state s and action a, γ the discount factor and $T(s'|s,a)$ the transition model from state s and action a to state s'. The value iteration algorithm iteratively updates the estimate of U^* using the Bellman equation. Once U^* is known, we can extract the optimal policy $\pi^*(s)$ as:

$$\pi^*(s) \leftarrow \underset{a}{\text{argmax}} (R(s,a) + \gamma \sum_{s'} T(s'|s,a) U^*(s')) \tag{5}$$

We need to keep track of the Bellman residuals defined as $||U_k - U_{k-1}||$ as termination criteria for the loop in the value iteration algorithm and set a precision threshold δ. Algorithm 1 below describes the value iteration algorithm which we use.

*Calculated by averaging initial production from real data for 37 wells in the Delaware Basin.

We used the Python `mdptoolbox` to solve our MDP using value iteration with a discount of $\gamma = 0.9$ and $\delta = 0.01$.

Algorithm 1: Value iteration algorithm

```
Result: U_k
k ← 0 ;
U_0(s) ← 0 for all states s;
while ||U_k − U_{k−1}|| < δ do
    U_{k+1}(s) ← max_a[R(s,a) + γ Σ_{s'} T(s'|s,a)U*(s') for all states s;
    k ← k + 1;
end
```

IV. Results and Discussion

The outcome of the value iteration algorithm is a policy specifying the action to take at every individual state such that over a period of 5 years the reward $\mathcal{R}(s, a)$ (equation (3)) is maximized. For example, for a state s where only one well has already been drilled and has an associated EUR of 98,404 bbl (low productivity range well) and was contracted at a price of \$30/bbl, the action a suggested by the policy is to drill another well, in this case at the third potential well location.

$$s = \begin{bmatrix} 1 & 98,404 & \$30 \\ 0 & 0 & \$52 \\ 0 & 0 & \$52 \end{bmatrix}, a = \begin{bmatrix} 0 \\ 0 \\ 1 \end{bmatrix}$$

To better interpret the policy outcome of the value iteration algorithm, we studied the fraction of times the policy suggested to drill a new well, given the number of wells already in production and their productivity measured by the average EUR per well. Table 2 shows the results of this analysis. We group the EUR productivity in four bins: "low, medium low, medium high and high" and aggregate results over all states in our state space, grouping them by the number of wells drilled. The entries of the table can be interpreted as follows: a fraction close to 1 means that the majority of actions recommended by the policy involve the drilling of a new well; an fraction equal 0 means that the policy recommends to not drill. The "impossible cases" in Table 2 represent scenarios where no wells are drilled and productivity levels are high, which is infeasible.

From these results, we notice clearly that drilling a new well is the recommended action in most cases. If there is still an option to drill (i.e. the number of wells already drilled is 0, 1 or 2), the optimal policy suggests to drill a well (corresponding to the entries equal to 1 in Table 2). Drilling a new well enables more oil production and therefore more revenues from the sale of oil. However, when the number of wells drilled is already at its maximum (i.e. 3 wells), there is little incentive to decide to drill a well as the reward function would only penalize the state-action by the cost of drilling $C_{drilling}$ (equation (3)). This is clear when the EUR/well is at a high or medium high level and the average of recommended actions is 0, corresponding to the action to not drill. However, it is interesting to notice that at lower levels of EUR/well, the decision to drill a new well is still recommended by our policy in some cases (averages of 0.86 and 0.13 for the low and medium low levels) regardless of its lack of viability.

Table 2 Aggregated results from our policy obtained with value iteration. The entries in this table represent the fraction of times when the policy suggests to drill a well.

	Level of EUR/well			
Number of wells drilled	**Low**	**Medium Low**	**Medium High**	**High**
0	1	Impossible case	Impossible case	Impossible case
1	1	1	1	1
2	1	1	1	1
3	0.86	0.13	0.0	0.0

Given our simplified model with a small number of wells that can be drilled (3 wells) and limited resolution of the price of oil and EUR variables (binned with $n = 3$ bins), it is difficult to reach interesting situations where the recommended action would be not to drill a well and wait. However, our implementation requires the transition probabilities matrix to be pre-computed, which is computationally expensive and scales with $O(n^2)$. Even when increasing the resolution of P_{O_t} and EUR to $n = 5$ bins each, the computation of the transition probability matrix takes more than 6 hours using a desktop computer. Using an alternative online methodology such as forward search, branch and bound or Monte Carlo Tree Search could help overcome these limitations in our current problem formulation and algorithm choice as the policy for the whole state space does not need to be computed.

V. Conclusion

The goal of this analysis was to obtain a policy to operate the drilling operations of an unconventional oil company given uncertainty in the oil prices and declining well production rates. We formulated the problem as an MDP. By solving the MDP with value iteration, we obtained a policy which determines the optimal drilling action to take given our state space which encompasses uncertain well productivity levels and oil prices. Whenever the action of drilling was still feasible, the policy suggested to drill a new well to increase oil production and therefore revenues. When the maximum number of wells drilled was reached, the policy recommended the action of not drilling in the cases of medium and high oil recoveries. These results are aligned with our prior expectations of what a policy should advise the decision-maker to do in this context. As future work, we are interested in extending the model to include the action of re-drilling a well at an existing location that has already been drilled. This will make our model more realistic because in real unconventional oil production, if the recovery at a specific location is not optimal, additional wells might be drilled at that location with the aim of increasing the productivity.

VI. Contributions

D. Orta Alemán and C. Venereau worked together for the entirety of the project. D. Orta Alemán focused more on writing the code implementand C. Venereau on the mathematical problem formulation and paper.

References

[1] Humphries, T.D., Haynes, R.D. and James, L.A. (2014), Simulataneous and sequential approaches to joint optimization of well placement and control, *Computational Geoscience*, 18:433-448, doi:10.1007/s10596-013-9375-x

[2] Shirangi, M.G. (2017), Advanced Techniques for Closed-Loop Reservoir Opimization Under Uncertainty, *PhD Thesis*, Palo Alto, CA, USA: Stanford University.

[3] Davidson, J.E. and Beckner, B.L. (2003), Integrated optimization for rate allocation in reservoir simulation, *SPE Reservoir Simulation Symposium*. 79701, Houston, TX, USA

[4] Isebor, O.J. (2013), Derivative-free optimization for generalized oil field development, *PhD Thesis*, Palo Alto, CA, USA: Stanford University.

[5] Wang, P. (2003), Development and applications of production optimization techniques for petroleum fields, *PhD Thesis*, Palo Alto, CA, USA: Stanford University.

[6] Rodriguez Torrado, R., Rios, J. and Tesauro, G. (2017), Optimal Sequential Drilling for Hydrocarbon Field Development Planning, *Proceedings of the Twenty-Ninth AAAI Conference on Innovative Applications (IAAI-17)*

[7] Balasubramanian, S., Chen, P., Bose, S., Alzahabi, A., and Thakur, G., Recent Advances in Enhanced Oil Recovery Technologies for Unconventional Oil Reservoirs. *Offshore Technology Conference. doi:10.4043/28973-MS*

[8] Purvis, D. C., and Kuzma, H., Evolution of Uncertainty Methods in Decline Curve Analysis. *Society of Petroleum Engineers. doi:10.2118/179980-MS*

[9] Elliott, R., Matthews, C. (2019), Oil and Gas Bankruptcies Grow as Investors Lose Appetite for Shale, *The Wall Street Journal*, August 2019.

[10] ENERCOM (2017), Energy Industry Data and Trends: Delaware Basin Study, [Online] Available at: https://www.oilandgas360.com/wp-content/uploads/2017/10/EnerCom-Energy-Industry-Data-and-Trends-September-Combined.pdf [Accessed: 5-Dec-2019]

[11] Energy Information Administration (2016), Trends in US Oil and Natural Gas Upstream Costs.