AN INFORMATION-THEORETIC APPROACH TO SENSOR DEPLOYMENT FOR HYDROCARBON PRODUCTION SURVEILLANCE

Ashutosh Tewari*, Kuang-Hung Liu, Stijn de Waele and Dimitri Papageorgiou
Corporate Strategic Research, ExxonMobil Research & Engineering Co.
Annandale, NJ 08801

Abstract Overview

Production surveillance is the task of monitoring oil and gas production from every well in a hydrocarbon field. Accurate surveillance is a basic necessity for several reasons that include improved hydrocarbon accounting, improved resource management, reduced operational cost, and ultimately optimal hydrocarbon production. A key challenge in this task, for large fields with many wells, is the measurement of multi-phase fluid flow using a limited number of noisy sensors of varying characteristics. Current surveillance practices are based on fixed utilization schedules of these flow sensors, which rarely change over time. Such a *passive* mode of sensing is completely agnostic to surveillance performance, thus, it often fails to achieve a desired accuracy. Here we propose an approach of *active surveillance*, underpinned by the concept of *value of information* based sensing. Borrowing some well-known concepts from optimal experiment design, reinforcement learning and artificial neural networks, we demonstrate that a practical active surveillance strategy can be devised that not only can improve surveillance performance significantly, but also reduce usage of flow sensors.

Keywords

Active Sensing, Production Surveillance, Markov Decision Process, Deep Reinforcement Learning.

Introduction

In many industrial applications involving complex systems, there is a need to monitor/probe a physical domain to gather adequate data necessary for informed decision making. However, it is usually unclear how to optimally deploy sensors to fulfill such data requirements. The absence of a systematic sensor deployment strategy often leads to the collection of huge volumes of data, albeit with little information therein. In this era of ever-cheaper sensor technology and large sensor networks, the need for smarter sensor deployment strategies is stronger than ever. Consider the example of production surveillance in a large onshore hydrocarbon field, where the goal is to accurately monitor the production rates of thousands of wells individually. This is challenging for many reasons, such as disparate, multiphase flow-meters of varying fidelities and costs, well-comingling, evolving surface and subsurface conditions, time varying sensor noise and the sheer number of wells to be monitored continuously (Tewari et al., 2018). The current practice of production surveillance in the oil and gas industry is primarily based on predetermined schedules for sensor deployment. As a result, it may be possible to monitor production at a coarser level (a well group), but it is extremely difficult to obtain accurate production estimates for individual wells.

* To whom all correspondence should be addressed

In this work, we propose an *active sensing* strategy for production surveillance, where the sensors are no longer deployed based on a pre-determined schedule. Instead, at every time step (e.g., a day) an optimal deployment schedule is identified on the fly. Mathematically, the problem can be cast as a *Partially Observed Markov Decision Process* (POMDP), which is an extensively studied research area in the optimization and control communities. In a nutshell, a POMDP breaks down the optimal sensor deployment in two distinct steps: *state estimation* and *sensor optimization*. The former updates our current knowledge of the system state (in this case, the production rates of individual wells) given new measurements, and explicitly characterizes the uncertainties in state estimates. The latter determines an optimal plan of sensor deployment for the next time step by optimizing a cost function with two competing objectives; maximizing the information gain (uncertainty reduction) and minimizing the sensing cost.

Theory

Consider a setting where we are interested in inferring rates $x_t \in \mathbb{R}^n$ at time t from n wells using a subset of a total of m possible measurements $y_t = [y_t^1, y_t^2 \cdots y_t^m]^T$ of varying fidelities and costs. Assume that the measurements are a linear function of the rates with an additive Gaussian noise of known variance, i.e. $y_t^i | x_t \sim \mathcal{N}(h_i^T x_t, \sigma_i), \forall\, i \in \{1,2,\cdots,m\}$. Let $\mathcal{M}_t$ denote a set of indices of measurements selected at time t. Thus, the subset $y_t(\mathcal{M}_t)$ has a multivariate normal distribution, $y_t(\mathcal{M}_t)|x_t \sim \mathcal{N}(H(\mathcal{M}_t)x_t, R(\mathcal{M}_t))$, conditioned on the rates x_t. The matrix $H \in \mathbb{R}^{m\times n}$ contains all measurement functions: $H = [h_1, h_2 \cdots h_m]^T$, and the matrix $H(\mathcal{M}_t) \in \mathbb{R}^{|\mathcal{M}_t|\times n}$ is obtained by selecting the rows of H corresponding to the elements of $\mathcal{M}_t$. Similarly, the diagonal matrix $R(\mathcal{M}_t)$ encodes the noise variance of the selected measurements. Given a Gaussian distribution of rates at previous time step, $x_{t-1} \sim \mathcal{N}(\mu_{t-1}, \Sigma_{t-1})$, a linear dynamics on x specified by the matrix $A \in \mathbb{R}^{n\times n}$ and a process noise covariance $Q \in \mathbb{R}^{n\times n}$, a standard result from Bayesian inference yields a Gaussian posterior distribution of rates i.e. $x_t | y(\mathcal{M}_t) \sim \mathcal{N}(\mu_t, \Sigma_t)$, where covariance Σ_t is given by equation (1),

$$\Sigma_t = \tilde{\Sigma}_{t-1} - \tilde{\Sigma}_{t-1} H(\mathcal{M}_t)^T \left(H(\mathcal{M}_t)\tilde{\Sigma}_{t-1}H(\mathcal{M}_t)^T + R(\mathcal{M}_t) \right)^{-1} H(\mathcal{M}_t)\tilde{\Sigma}_{t-1}\ , \tag{1}$$

with $\tilde{\Sigma}_{t-1} = A\Sigma_{t-1}A^T + Q$. For static systems (time independent) where the measurements are to be chosen only once, the problem of best subset selection is thoroughly studied in experiment design research. To this end, several optimality criteria (A-optimality, D-optimality, Q-Optimality, etc.) have been proposed to optimize some metric defined on the posterior covariance matrix Σ_t (Shah and Sinha, 1989). For example, the A-optimality criterion finds a subset that minimizes the trace of posterior covariance. When coupled with the known costs of measurements, these criteria help navigate the tradeoff between information gain and sensing cost by posing an optimization problem of the following form,

$$\underset{\phi(\mathcal{M}_t)}{\arg\min}\ \ \phi(\mathcal{M}_t) \quad \text{s.t.} \quad \phi(\mathcal{M}_t) = \mathrm{tr}(\Sigma_t) + c(\mathcal{M}_t). \tag{2}$$

The composite cost $\phi(\mathcal{M}_t)$ combines the A-optimality criterion, $\mathrm{tr}(\Sigma_t)$, and the cost, $c(\mathcal{M}_t)$, of choosing the measurement subset $\mathcal{M}_t$. In a dynamic setting, the optimization becomes trickier as it is often desirable to minimize a cost defined over a future time horizon, as shown in (3), via a discount factor $\gamma \in [0,1)$. The discount factor provides a knob to appropriately down-weight future costs.

$$\underset{\mathcal{M}_t, \mathcal{M}_{t+1}, \mathcal{M}_{t+2}, \cdots}{\arg\min} \quad [\phi(\mathcal{M}_t) + \gamma\phi(\mathcal{M}_{t+1}) + \gamma^2\phi(\mathcal{M}_{t+2}) + \cdots] \tag{3}$$

The optimization problem in (3) is combinatorial and considerably harder to solve than (2) even if the horizon is finite. A key insight that allows an efficient solution of (3) is that the discounted cost is completely specified by the time evolution of Σ_t, which admits deterministic, Markovian dynamics as given by equation (1). As a result the problem of optimal measurement subset selection can be modeled as a *Markov Decision Process* (MDP) (Puterman, 1994). With this insight, we appeal to methods from *Deep Reinforcement Learning*, which exploit the *Bellman optimality criterion* for MDPs (Lagoudakis, 2017) and the universal function approximation capability of deep neural networks (Cybenko, 1989) to yield an extremely tractable optimization problem as given in (4),

$$\underset{\mathcal{M}_t}{\arg\min}\ \ Q^*(\Sigma_t, \mathcal{M}_t) \tag{4}$$

where $Q^*(\Sigma_t, \mathcal{M}_t)$ is called an optimal state-action value function that is learned using a deep neural network. Refer to (Minh et al., 2013) for more details on this subject.

Experimental Results

Experiments are conducted in a simulated environment to evaluate different sensor deployment strategies. The need of simulation stems mainly because it provides access to true rates, which is necessary for an objective comparison. Similar information from an oil field is not only proprietary but also cost prohibitive to obtain. The simulations are governed by a dynamic production model available in public domain (Tewari et al., 2018). We choose a well configuration, shown in figure 1, of eight wells feeding into a multiphase flow-meter called *test-separator*. The wells are grouped into two groups of four for the purpose of flow measurement, a practice known as *well-comingling*. During this process one or more wells from a group can be shut to gain more information about individual well production. The measurement noise variance of the test-separator is assumed to be 4 and 15 for groups 1 and 2, respectively.

Such variations in noise characteristics are quite common in the field. An additional mode of sensing, called *spin-cuts*, is assumed available, which involves manually taking multiphase fluid samples from a well and measuring water to oil ratio. Spin-cuts are accurate but expensive. The usage costs of test-separator and spin-cut are assumed 10 and 100 units respectively.

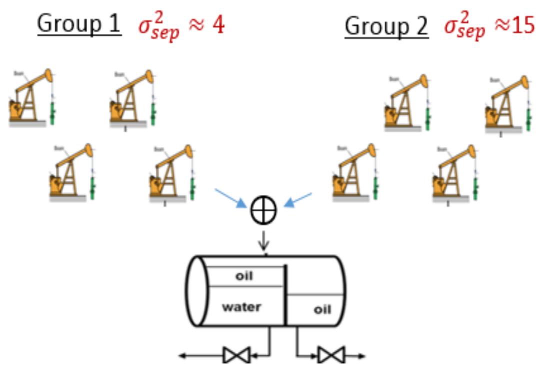

Figure 1. A well configuration used for comparative study of passive & active surveillance approaches

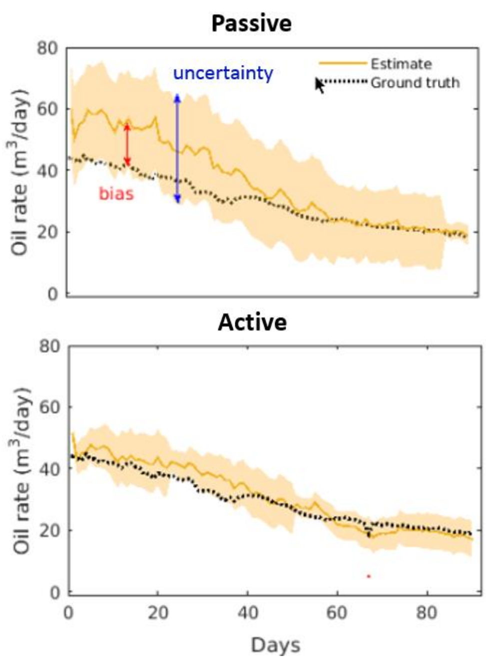

Figure 2. Production surveillance result of one of the eight wells using the two sensing approaches. The estimated rates using the passive sensing approach exhibit higher bias and uncertainty (with respect to the ground truth) compared to the active sensing approach. In addition, active sensing required 25 fewer well tests than its counterpart.

The passive sensing approach uses a fixed schedule of sensor deployment as follows. A well group is channeled through a test-separator every 3rd day, and a spin-cut is done on a well once every 3 months. Active-sensing, on the other hand, uses the proposed approach to deploy sensors as needed. Note that the decision of shutting a well (or wells) during a test-separator measurement is considered a type of sensor deployment. The estimated rates from one of the 8 wells is shown in figure 2 for both sensing strategies along with the true rates. Active sensing not only improved surveillance performance but also reduced the sensing cost by requiring 25 fewer test-separator measurements.

Conclusion

In this work we outline an information theoretic approach for production surveillance in oil fields, a task which is often hampered by inadequate sensing resources and continuously evolving subsurface conditions. The approach is grounded on the principal of information based planning that actively deploys sensors when and where the need is the most, while being cognizant of sensing cost. A proof of concept, in a simulated environment, clearly demonstrates advantages over the current practice of passive surveillance. Additionally, the proposed sensor planning approach is conceptually portable to other application areas of relevance to oil & gas industry such as seismic surveys, methane leak detection, loss-prevention systems, and seabed seep detection.

References

Cybenko, G. (1989). Approximations by superpositions of sigmoidal functions. *Mathematics of Control, Signals, and Systems*. 2(4), 303-314.

Lagoudakis M. G. (2017). Value Function Approximation. Encyclopedia of Machine Learning and Data Mining. *Springer*. Boston, MA.

Mnih, V., Kavukcuoglu, K., Silver, D., Graves, A., Antonoglou, I., Wierstra, D., and Riedmiller, M. (2013). Playing atari with deep reinforcement learning. *NIPS Deep Learning Workshop*.

Puterman, M. L. (1994). Markov decision processes: discrete stochastic dynamic programming, New York: Wiley.

Shah, K. R. and Sinha, B. K. (1989). Theory of Optimal Designs. *Lecture Notes in Statistics. 54*. Springer-Verlag. pp. 17.

Tewari, A., De Waele, S., Subrahmanya, N. (2018). Enhanced production surveillance using probabilistic dynamic models, *International Journal of Prognostic Health Management*. 9 (1), 1-12.