European Journal of Operational Research 230 (2013) 688–702

Contents lists available at SciVerse ScienceDirect

European Journal of Operational Research

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

Innovative Applications of O.R.

Dynamic decision making for graphical models applied to oil exploration

Gabriele Martinelli [a,*], Jo Eidsvik [a], Ragnar Hauge [b]

[a] Dept. of Mathematical Sciences, Norwegian University of Science and Technology, Alfred Getz' vei 1, Trondheim, Norway
[b] Norwegian Computing Center, Gaustadalleen 23, Oslo, Norway

ARTICLE INFO

Article history:
Received 21 February 2012
Accepted 29 April 2013
Available online 11 May 2013

Keywords:
Bayesian Networks
Dynamic programming
Graphical model
Heuristics
Petroleum exploration

ABSTRACT

We present a framework for sequential decision making in problems described by graphical models. The setting is given by dependent discrete random variables with associated costs or revenues. In our examples, the dependent variables are the potential outcomes (oil, gas or dry) when drilling a petroleum well. The goal is to develop an optimal selection strategy of wells that incorporates a chosen utility function within an approximated dynamic programming scheme. We propose and compare different approximations, from naive and myopic heuristics to more complex look-ahead schemes, and we discuss their computational properties. We apply these strategies to oil exploration over multiple prospects modeled by a directed acyclic graph, and to a reservoir drilling decision problem modeled by a Markov random field. The results show that the suggested strategies clearly improve the naive or myopic constructions used in petroleum industry today. This is useful for decision makers planning petroleum exploration policies.

1. Introduction

This paper considers the problem of sequential decision making, where the outcome of one decision will influence the others. Our motivation and main applications are from oil and gas exploration, where a petroleum company must evaluate a set of potential drilling prospects. For each prospect, we may either drill or not. There is a cost of drilling, but revenues if the well discovers oil or gas. The prospects are statistically dependent, and drilling at one prospect gives information that is used to update the probability of success at other prospects. The goal is to find an optimal drilling sequence, including when to stop drilling and abandon the remaining prospects.

The optimization of the expected utility function is a trade-off between two factors: the direct reward from the exploitation, and the indirect gain of learning, or exploration, that helps us make informed future decisions. The balance between the two is controlled by a discounting factor. With no discounting, the problem becomes a maximization of the value of information, whereas a high discounting factor leads to a greedy search where only immediate gain counts.

In the oil industry prospects are typically evaluated one-by-one. The implicit working assumption is then independence between prospects, and a greedy search is optimal. As petroleum companies are now forced to look for smaller volumes, gains can be achieved by joint modeling of prospects. Recent work by VanWees et al. (2008) and Martinelli et al. (2011) use Bayesian Networks (BNs) to capture the geological dependencies between prospects, while Bhattacharjya et al. (2010) study the effect of various data acquisition schemes for reservoir units modeled by a Markov random field (MRF). Dependence means that we can update the probability model after exploring the most lucrative prospect. We can next go for the second best prospect, conditional on the outcome of the first, and so on. This line of thinking leads to a myopic (conditional greedy) approach, which uses the dependence in the model for forward learning about the prospects. As is common in sequential decision making, this forward selection approach can be improved by taking the expected value over all possible future drilling scenarios into account, which leads to the optimal solution given by a dynamic program (DP).

Our goal with the current paper is to compare various dynamic strategies for the large BN model in Martinelli et al. (2011) and on the MRF in Bhattacharjya et al. (2010). This challenge of constructing drilling strategies for dependent prospects has not been studied much, except certain special cases: Kokolis et al. (1999) describe a similar problem with a focus towards decision making under uncertainty and the technical risks connected to a project. Smith and Thompson (2008) analyze the consequences of dependent versus independent prospects, and give drilling guidelines that are optimal in special situations. In Bickel and Smith (2006) and Bickel et al. (2008) DP is used to compute the optimal sequences and profits from six dependent prospects. The big challenge which we address here is that related to the combinatorial increase in the number of scenarios. DP is not tractable when the number of prospects gets large.

* Corresponding author. Address: Schlumberger, SNTC, Aslakveien 14E, 0753 Oslo, Norway. Tel.: +47 45775460.

E-mail addresses: gmartinelli@slb.com (G. Martinelli), joeid@math.ntnu.no (J. Eidsvik), ragnar.hauge@nr.no (R. Hauge).

0377-2217/$ - see front matter © 2013 Elsevier B.V. All rights reserved.
http://dx.doi.org/10.1016/j.ejor.2013.04.057

A possible solution to large problems is offered by approximate DP methods, see Bertsekas and Tsitsiklis (1996) and Powell (2008). The optimization function is then replaced with a statistical model that captures the impact of decisions now on the future. For our graphical representation of dependent prospects it is not obvious how to find a statistical model that approximates the future value function. Instead we study look-ahead policies, where a DP is used for a finite future horizon and heuristics approximate the continuation value (CV). We also apply pruning of the decision tree, i.e. we ignore unlikely branches to reduce the combinatorial problem. The sequential decisions are thus made according to a rolling-horizon algorithm, where we pick one prospects at a time, update the probabilities, look-ahead using DP and select again. Similar strategies are discussed in Chapter 8 of Powell (2007). An application of such strategies to wind energy is presented in Zhou et al. (submitted for publication). Our methods are different because of the statistical modeling based on a BN and MRF. The operations research community have been interested in research ideas at this interface (Meisel and Mattfeld, 2010) in diverse applications, see e.g. Falzon (2006) for military operations. Moreover, the field of learning within BNs is quite active, see e.g. Dearden et al. (1999), Heckerman (1999) or Sucar et al. (2012), but there has been little focus on the sequential selection of nodes, which is our focus for the petroleum prospect selection problem.

Note that when considering a set of independent prospects, the optimal sequential decisions are offered by Gittins indeces (Gittins, 1979), used for a petroleum example by Benkherouf and Bather (1988). In our model the correlation is much more complex, and the actions influence the model probabilities in a complicated manner. Branch and bound methods are non-heuristic in the sense that they produce lower and upper bounds for the values (Goel et al., 1979). In practice the gap between bounds can be wide, and in our context we will typically lack monotonicity when computing the best (discounted) sequence. See Brown and Smith (submitted for publication) for promising work in this direction, using the BN that we consider here as an example. See also Ryzhov et al. (2012) for continuous examples in this context, with statistical dependence.

We have no theoretical restrictions on the underlying statistical model for dependence. There is a practical requirement that conditional distributions can be computed fast, since many of these conditionals will be evaluated when designing a strategy. For comparing strategies, it is advantageous if we can easily simulate from the models. The BNs and MRFs we consider here are fast to update and easy to simulate from.

The paper develops as follows: In Section 2 we motivate by introducing the notation and statistical model for the oil and gas exploration examples. In Section 3 we present the DP algorithm for our problem. In Sections 4 and 5 we propose the various heuristic strategies, and the algorithms used to evaluate the properties of the sequential exploration strategies. In Section 6 we provide results for a small BN model and the BN case study of 25 prospects in the North Sea. In Section 7 we analyze a MRF for a oil reservoir represented on a 5×20 lattice.

2. Background, modeling and notation

We consider a set of N prospects. These N prospect nodes are a subset of the total M nodes in a graph. The remaining $M - N$ auxiliary nodes impose the specified dependency structure in the model, but are not observable. For every node $i = 1, \ldots, M$ we have a discrete random variable $x_i \in \{1, \ldots, k_i\}$. In the examples below we use $k_i = k$, and $k = 3$. The random vector of all variables is $\mathbf{x} = (x_1, \ldots, x_M)$, where the N first components correspond to the prospect variables. We model the probability distribution of $\mathbf{x}$ by a BN or a MRF. We will next motivate our problem description via our main case study.

The BN in our main example is defined via a directed graph, which means the joint probability model $p(\mathbf{x})$ is the product of conditional distributions $p(x_i|x_i^{\mathrm{pa}})$, for all nodes $i = 1, \ldots, M$, and x_i^{pa} denotes the set of outcomes at parent nodes of i. Fig. 1 shows the directed graph connecting parent nodes to nodes via edges. The graph contains $N = 25$ prospect nodes, while there are $M = 42$ nodes in total.

The graph in Fig. 1 is built from the causal large scale geological processes required to make sufficient amounts of oil and gas, see

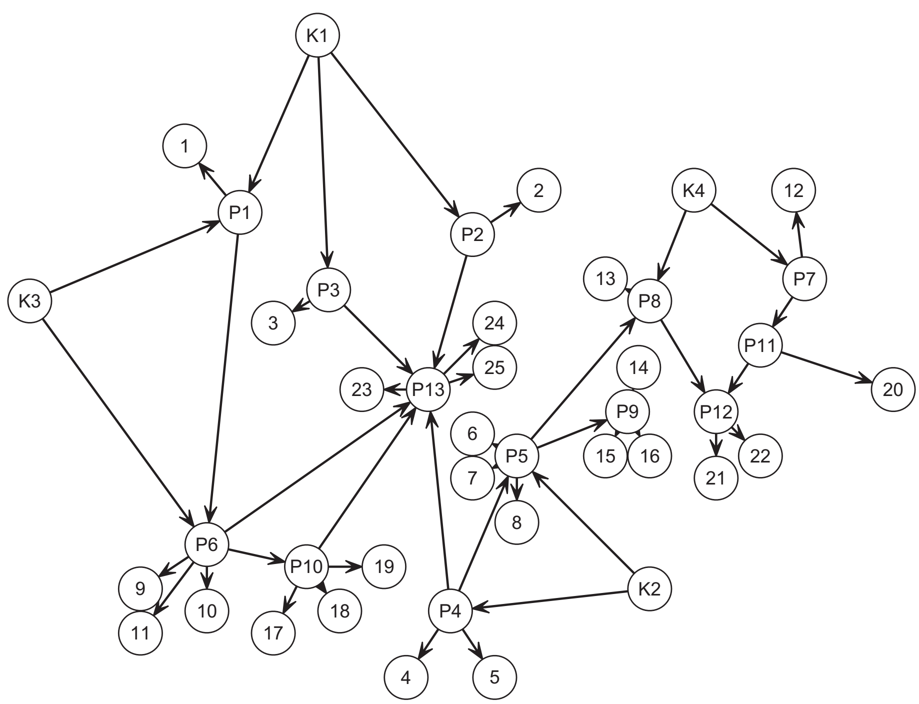

Fig. 1. Motivation network. In this case we have 25 drilling prospects, identified with the nodes from 1 to 25, where we can possibly drill. The BN was first presented in Martinelli et al. (2011).