Probabilistic Engineering Mechanics 37 (2014) 93–108

Contents lists available at ScienceDirect

Probabilistic Engineering Mechanics

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

Optimum inspection and maintenance policies for corroded structures using partially observable Markov decision processes and stochastic, physically based models

K.G. Papakonstantinou *, M. Shinozuka

Department of Civil Engineering and Engineering Mechanics, Columbia University, New York, USA

ARTICLE INFO

Article history:
Received 15 May 2013
Received in revised form
15 May 2014
Accepted 9 June 2014
Available online 24 June 2014

Keywords:
Partially observable Markov decision processes
Stochastic control
Uncertain observations
Structural life-cycle cost
Inspection and maintenance policies
Spatial stochastic corrosion model

ABSTRACT

Stochastic control methods have a history of implementation in risk management and life-cycle cost procedures for civil engineering structures. The synergy of stochastic control methods and Bayesian principles can result in Partially Observable Markov Decision Processes (POMDPs) that allow consideration of uncertainty within the entire domain of the model and expand available policy options compared to other state-of-the art methods. The superior attributes of POMDPs enable optimum decisions which are based on the belief space or otherwise only on the best knowledge that a decision-maker can have at each time. In this work the effort is mostly based in modeling and solving the problem of finding optimal policies for the maintenance and management of aging structures through a POMDP framework with large state spaces that can adequately and sufficiently describe real-life problems. In order to form the POMDP framework, stochastic, physically based models can be used and their connection to the control process is explained in detail. Specific examples of a corroded existing structure are presented, based on non-stationary POMDPs, for both infinite and finite horizon cases with 332 and 14,009 states respectively. Results from both cases are compared and discussed and the capabilities of the method become apparent.

1. Introduction

A nation's civil infrastructure state has significant impact on a nation's safety, quality of living and economy. It is thus very important to prevent structural degradation and maintain each structure throughout a sustainable period, in which the structure will be able to satisfy specified design and other requirements.

Aging effect on structures can be decisive. Preventing structural deterioration however is not an easy task. Research efforts currently are not giving much attention to maintenance, repair and inspection of existing infrastructure and inevitably existing practices lack the sophistication that such an important problem should deserve. Concerning bridges for example, usually inspections are performed periodically every two years, regardless of the bridge's condition, and are based on mere visual observations, which are obviously uncertain and can even be completely inadequate on some occasions. However, as will be explained further in the coming sections, in current practice whenever these inspection data are being used as input in a management software they are considered certain and accurate. It is easily understood that these procedures are simplistic and many important questions are still generally unanswered. In order to achieve an acceptable safety level at the minimum possible cost the number of inspections and the inspection times should not be predetermined, a variety of inspection and repair techniques should be taken into account and the cost and uncertainty of information should be considered.

A large variety of different formulations can be found in the literature addressing the problem of finding optimum policies for the maintenance and management of aging civil infrastructure. Several methods and valuable references can be found in review articles [1–4]. The majority of works in this field, however, although generally different in used methods and employed assumptions share a common limitation. They ignore in their modeling observations from inspection or monitoring along the structural life that reveal essential information about the true condition state of the structure, either by completely disregarding this vital aspect of the problem or by simply modeling them indirectly, through adjusting the probability of performing maintenance, e.g. [5], or the maintenance effectiveness, [6] for example. These formulations, although mathematically interesting, are generally impractical for real-life implementation. It is not surprising, therefore, that similar approaches are currently avoided by managing agencies which normally prefer to use asset management programs

* Corresponding author. Tel.: +1 949 228 8986.
E-mail address: kp2570@columbia.edu (K.G. Papakonstantinou).

http://dx.doi.org/10.1016/j.probengmech.2014.06.002
0266-8920/

that make use of a form of discrete stochastic optimal control or otherwise Markov Decision Processes (MDPs).

Markov Decision Processes are frequently employed for maintenance planning of degrading civil engineering structures [2]. A strong indication of their success and capabilities is their use from different state agencies all over the world for asset management of a variety of infrastructures, [7–9]. In United States, PONTIS, a registered trademark of AASHTO and the predominant management system for bridges and other infrastructures which manages at least 750,000 structures today, uses MDPs as its core optimization tool, [10–12]. The successful reliance of infrastructure management projects on MDPs is mainly due to the facts that MDPs advice the decision-maker to make optimum sequential decisions based on the actual inspection observations along the years, the ease in which extensive amount of inspection data can be incorporated in the mathematical framework and finally the comprehensive studies on these methods over the years, by several researchers in a variety of scientific fields.

Nonetheless, even though MDPs provide a very strong and versatile mathematical framework for asset management, they also have some limitations which, at certain occasions, may be crucial for the quality of solutions they provide. A basic assumption of MDPs is the fact that inspections always reveal the true state of the system with certainty. While many problems in infrastructure management can perhaps support this feature, there are equally many, if not more, occasions where such an assumption is unrealistic. Another limitation, originating from the perfect inspections assumption, is that whenever an action is performed at a decision epoch, including the do-nothing action, necessarily an inspection has to precede it. Lastly, in the MDPs framework the notion of the cost of information is lost, since all inspections are assumed perfect. In reality, inspections of infrastructure facilities usually have a non-negligible cost and more accurate inspection techniques are self-evidently more expensive than cruder inspection methods. In addition to the previous limitations, MDPs are not very flexible on suggesting decision intervals either and hence PONTIS, for example, uses fixed, periodic, biannual inspection times in its modeling. However, unlike the previous unsurpassed limitations the latter one can be addressed with extreme care and refined modeling, [13]. To address the aforementioned limitations of MDPs a generalization of discrete stochastic optimal control based on Bayesian principles, which is usually called Partially Observable Markov Decision Processes (POMDPs), has been suggested. POMDPs constitute a much more general tool that inherits all the valuable attributes of MDPs and adds more. The enhanced attributes of POMDPs allow the combined solution of both inspection and maintenance optimization problems taking into account that the current state of the system at each time step may not be observed with certainty due to measurement errors. In a POMDP the system dynamics are determined by a MDP, but the decision-maker may not necessarily directly observe the underlying structural state with certainty. Instead, a probability distribution over the set of possible states is maintained, based on a set of observations and observation probabilities, and the underlying MDP.

POMDPs comprise a newer scientific field, not as mature as the MDP one. This reason, in addition to the fact that are much harder to be solved adequately for large, complex, realistic problems, has led until now to very few works addressing them, in the context of optimum inspection and maintenance of civil infrastructure systems. In Madanat and Ben-Akiva [14] a POMDP problem with 8 states is solved and in Smilowitz and Madanat [15] a problem of just 3 states, concerning a network of highway pavements, is presented. In Ellis et al. [16] and Jiang et al. [17] some finite horizon POMDP problems are analyzed, concerning structural degradation of bridge girders due to corrosion and fatigue. The maximum size of the state spaces in these two works is 13. Finally, Faddoul et al. [18] studied an inspection and maintenance problem, regarding a reinforced concrete highway bridge deck, and sought optimum policies for a 5-state POMDP with a horizon length of 20 years. In all these works the state space of the problems has been kept considerably small or in other words the system has not been modeled in a refined way that can suitably describe realistic problems. The reason for this is the significant difficulty in solving POMDPs of large domains with numerous states. In MDP formulations, where solutions can be much more easily found, state space sizes in the order of hundreds or thousands are commonly encountered.

In this work, a broader, scaled-up approach to the problem is achieved, based on a spatial stochastic, physically based model, asynchronous dynamic programming and Perseus, a point-based value iteration algorithm for POMDPs by Spaan and Vlassis [19]. In particular, modeling and solving the problem of finding optimal policies for the inspection and maintenance of aging structures through a POMDP framework with large state spaces is analyzed in detail and two example applications are presented and discussed, for both infinite and finite horizon cases, with 332 and 14,009 states respectively. In order to form the POMDP framework a spatial stochastic corrosion model is used and its connection to the control process is thoroughly explained. Algorithmic performance, potential limitations and possible future remedies are described as well. The suggested methodologies in this work result in very versatile and complex policies and the potential policy spectrum that the model can generally provide increases dramatically, compared to other methods. Without imposing any unjustified constraints on the policy search space (such as the usually used assumptions of periodicity, threshold performances, perfect inspections, predetermined number of interventions, deterministic environments and many more [13]), in order to make the optimization method capable of solving the problem, the POMDP model in this work accommodates aperiodic inspection intervals, imperfect observations, different inspection and maintenance choices, deterministic and probabilistic outcome of actions, perfect and partial repairs, non-stationary environments, infinite and finite horizons and suggested renewal times.

2. Discrete stochastic optimal control with full observability

Markov Decision Processes (MDPs) are controlled stochastic processes in which a decision-maker is uncertain about the exact effect of executing a certain action. A MDP assumes that at any decision time step t the environment is fully observed and is with certainty in a state $s \in S$, of a finite set of states, the agent takes an action $a \in A$, of a finite set of actions, and receives a reward (or cost) $R(s,a)$ as a result of this action, while the environment switches to a new state s' according to a known stochastic model, with transition probability $P(s'|s,a)$. In the rest of this work only rewards will be referred, since cost can be simply perceived as negative rewards. Thus, MDP is a 4-tuple (S, A, P, R) and the Markov property implies that the past of s is independent of its future, conditional on the present.

A decision problem within the MDP framework requires the decision-maker to find a sequence of actions that optimize some objective long-term reward function. Since the problem is stochastic, typically the objective function is additive and based on expectations. The most common function used is maximization of the total expected discounted rewards and this is the one considered in this work. One way to characterize a MDP policy is to consider its value function $V^{\pi} : S \rightarrow \mathbb{R}$, which represents the expected reward of some complete policy, π, which maps states to actions, $\pi : S \rightarrow A$. For every state s, V^{π} estimates the amount of discounted reward the decision-maker can gather when he starts