



Planning structural inspection and maintenance policies via dynamic programming and Markov processes. Part II: POMDP implementation



K.G. Papakonstantinou*, M. Shinozuka

Department of Civil and Environmental Engineering, University of California Irvine, Irvine, USA

ARTICLE INFO

Available online 24 April 2014

Keywords:

Partially Observable Markov Decision Processes
Optimal stochastic control
Belief space
Uncertain observations
Structural life-cycle cost
Infrastructure management

ABSTRACT

The overall objective of this two part study is to highlight the advanced attributes, capabilities and use of stochastic control techniques, and especially Partially Observable Markov Decision Processes (POMDPs), that can address the conundrum of planning optimum inspection/monitoring and maintenance policies based on stochastic models and uncertain structural data in real time. In this second part of the study a distinct, advanced, infinite horizon POMDP formulation with 332 states is cast and solved, related to a corroding reinforced concrete structure and its minimum life-cycle cost. The formation and solution of the problem modernize and extend relevant approaches and motivate the use of POMDP methods in challenging practical applications. Apart from uncertain observations the presented framework can also support uncertain action outcomes, non-periodic inspections and choice availability of inspection/monitoring types and intervals, as well as maintenance actions and action times. It is thus no surprise that the estimated optimum policy consists of a complex combination of a variety of actions, which cannot be achieved by any other method. To be able to solve the problem we resort to a point-based value iteration solver and we evaluate its performance and solution quality for this type of applications. Simpler approximate solvers based on MDPs are also used and compared and the important notions of observation gathering actions and the value of information are briefly discussed.

© 2014 Elsevier Ltd. All rights reserved.

1. Introduction

This paper complements a companion paper [1], that provided a thorough theoretical background on the use of Markov Decision Processes (MDPs) for infrastructure management. As shown in our Part I companion paper the Partially Observable case of Markov Decision Processes (POMDPs) provides an excellent choice for decision making and asset management under uncertainty, with firm mathematical foundations and superior attributes. POMDPs are capable of describing a huge number of realistic situations and can support a diverse range of formulations and objective functions, including condition-based, reliability and/or risk-based problems [1]. The optimum life-cycle policies are provided based on stochastic control, probabilistic models, uncertain structural data and Bayesian principles. This Part II paper emphasizes more on the details of POMDP models and solvers for optimum structural inspection and maintenance policies. A detailed application example for a corroding reinforced concrete structure is provided and among others the cost-benefit of information is naturally

incorporated in the formulation. Most quantities and notions already defined in Part I [1] are not defined again in detail in this work and any assumptions made in the companion paper are valid in this paper as well. For example, in this paper we will again only refer to rewards, since cost can be simply perceived as negative rewards.

Markov Decision Processes have a long, successful history of implementation in risk management and minimum life-cycle costing of civil engineering structures [2]. Perhaps the strongest 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, like bridges, transportation networks, pavements, etc., [3–5]. In United States, PONTIS, the predominant management system for bridges and other infrastructures, uses MDPs as its core optimization tool, [6–8]. PONTIS is currently a registered trademark of AASHTO and it is licensed and used by the majority of U.S. state transportation departments and other organizations in the U.S. and other countries.

Although 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. POMDPs are a much more general tool that inherit all the valuable attributes of MDPs and

* Corresponding author. Tel.: +1 949 228 8986.

E-mail address: kpapakon@uci.edu (K.G. Papakonstantinou).

add more. However, POMDPs comprise a newer scientific field, open to extensive scientific research currently and not as mature as the MDP one. These reasons, 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, in comparison to other approaches in this area, [1]. In Madanat and Ben-Akiva [9] a POMDP problem with 8 states and a finite horizon of 10 years is solved and in Smilowitz and Madanat [10] a problem of just 3 states, concerning a network of highway pavements, is presented. Both of these works use a fixed, regular grid and the nearest neighbor interpolation–extrapolation rule (enduring all the disadvantages of this method, discussed in the Part I paper [1]) and convert the problems into fully observable MDPs, which are then solved by dynamic and linear programming, respectively. In Ellis et al. [11] and Jiang et al. [12] 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 and the authors solve the problems with an exact algorithm, taking advantage of the small number of states and the finite horizon formulation, which in this occasion is computationally beneficial, in comparison to an infinite horizon case. Faddoul et al. [13] studied an inspection and maintenance problem, regarding a reinforced concrete highway bridge deck, and sought optimum policies based on the nearest neighbor method and a 5 state POMDP with a horizon length of 20 years. Use of POMDPs in works pertinent to the discussion herein can be further found in [14–17]. The maximum state space size used in these cases is 9, [14]. It is apparent that in all these works the POMDP formulation of the problems hindered the researchers from describing the system in a more refined way, with larger state spaces. In MDP formulations, where solutions can be much more easily found, state space sizes in the order of hundreds or thousands are commonly encountered and can be even considered small. An example for instance is the work by Robelin and Madanat [18] where a bridge deck management problem, formed as a MDP with 840 states, is solved. The state space in this case consists of the reliability index of the deck and history dependent parameters.

In this paper, we rely on these previously presented works and efficacious approaches, as in [11] and [19], and significantly extend them towards large-scale modeling and solution of realistic problems. Utilizing the presented topics in the Part I companion paper, a distinct, advanced, demanding infinite horizon formulation is cast and solved in this work with non-stationary stochastic phenomena, connection to physically based stochastic models and a considerably larger state-space of 332 states. Choice availability of different monitoring and maintenance actions, uncertain observation and action outcomes and non-periodic structural visits are also incorporated in this work. With such an unprecedented formation and variety of unimpeded options the estimated optimum policy is a highly complex combination of a range of inspection/monitoring types and intervals, and maintenance actions and action times, which cannot be achieved by any other method. To be able to solve this challenging problem that could not have been solved by techniques in the aforementioned references, we resort to point-based solvers, as explained in our companion paper, [1]. Point-based methods have been mainly developed in the field of artificial intelligence for autonomous robot navigation which is a problem with inherently different characteristics from the structural management problem. Among others, uncertainty in usual robot navigation problems decreases with time, since the terrain is gradually explored, and not the opposite like in structural maintenance, and different observation actions are not typically sought during planning. Despite the differences, we demonstrate in this work that the point-based value iteration algorithm Perseus, [20], can perform successfully in

this type of applications, even for difficult problems with larger state-spaces than the ones currently described in the maintenance literature. Additional, recent attempts by the authors with larger models, with thousands of states, in a finite horizon formulation can be seen in [21], where the POMDP mapping to physically based stochastic models is also explained in greater detail. Apart from Perseus we also solve the problem with simple approximate solvers (MLS, QMDP, [1]) that are directly based on MDPs. Current structural management systems (like PONTIS) only rely on MDPs and hence they could straightforwardly utilize these methods. We demonstrate differences in performance and quality of solution between these methods and Perseus and based on this comparison we also shed some extra light to the important notions of observation gathering actions and the value of information, which of course implies that more accurate and precise inspection/monitoring techniques are self-evidently more expensive than cruder inspection/monitoring methods. Compendiously, the current paper provides in detail a generic framework that modernizes the way relevant problems are solved today, sets a step forward in large-scale modeling, showcases deployment of point-based POMDP methods and motivates their use in a wider variety of problems and practical applications.

The particular application used in this work, in order to demonstrate the suggested POMDP framework, its detailed implementation, solution and attributes, relates to a corroding reinforced concrete structure. Unfortunately, current non-destructive corrosion evaluation techniques are prone to measurement errors and have inherent deficiencies, which make it difficult to derive, certain, reliable engineering conclusions based on their output [22–24]. For this reason, a POMDP formulation of the problem is most appropriate. A spatial, stochastic, physically based model of steel corrosion in a wharf deck reinforced concrete slab is developed in Papakonstantinou and Shinozuka [25]. Based on this modeling an infinite horizon, non-stationary POMDP model with yearly time-steps and 332 discrete physical states is cast in this work and solved by asynchronous dynamic programming and Perseus [20]. The objective of the application is to identify an optimum life-cycle cost policy that can suggest, without any modeling restrictions, when and what type of inspection/monitoring and maintenance actions should be employed based on the conditions of the deteriorating structure in real time [1]. On the whole, 4 different maintenance actions are considered (including doing nothing and full replacement of the structure) and 3 different inspection/monitoring actions (including no inspection), resulting in a total number of 10 considered, different, combined (maintenance–inspection) action choices for the decision-maker. The uncertain observation outcomes are categorized in 4 different possible conditions and inform the decision-maker about the structural status according to the accuracy and precision of the chosen evaluation method. In general, the most important characteristic of all the actions for the optimum life-cycle cost policy is their relative to each other effectiveness and cost, which are explained in detail in the paper. In the remainder of this work, we thoroughly present the POMDP modeling of the problem, emphasizing mostly in the do-nothing action, which of course relates to the deterioration process. The rewards part of the modeling and the combination of maintenance–inspection actions are described as well. In Section 3 the point-based value iteration algorithm, Perseus, is analyzed, while in Section 4 specific implementation details and comprehensive results are provided and discussed.

2. POMDP modeling

POMDP is a 6-tuple (S, A, P, O, Po, R) where S, A and O finite set of states, actions and possible observations respectively P state

متن کامل مقاله

دریافت فوری ←

ISIArticles

مرجع مقالات تخصصی ایران

- ✓ امکان دانلود نسخه تمام متن مقالات انگلیسی
- ✓ امکان دانلود نسخه ترجمه شده مقالات
- ✓ پذیرش سفارش ترجمه تخصصی
- ✓ امکان جستجو در آرشیو جامعی از صدها موضوع و هزاران مقاله
- ✓ امکان دانلود رایگان ۲ صفحه اول هر مقاله
- ✓ امکان پرداخت اینترنتی با کلیه کارت های عضو شتاب
- ✓ دانلود فوری مقاله پس از پرداخت آنلاین
- ✓ پشتیبانی کامل خرید با بهره مندی از سیستم هوشمند رهگیری سفارشات