



Planning structural inspection and maintenance policies via dynamic programming and Markov processes. Part I: Theory



K.G. Papakonstantinou*, M. Shinozuka

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

ARTICLE INFO

Available online 19 April 2014

Keywords:

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

ABSTRACT

To address effectively the urgent societal need for safe structures and infrastructure systems under limited resources, science-based management of assets is needed. 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. Markov Decision Processes are in general controlled stochastic processes that move away from conventional optimization approaches in order to achieve minimum life-cycle costs and advice the decision-makers to take optimum sequential decisions based on the actual results of inspections or the non-destructive testings they perform. In this first part of the study we exclusively describe, out of the vast and multipurpose stochastic control field, methods that are fitting for structural management, starting from simpler to sophisticated techniques and modern solvers. We present Markov Decision Processes (MDPs), semi-MDP and POMDP methods in an overview framework, we have related each of these to the others, and we have described POMDP solutions in many forms, including both the problematic grid-based approximations that are routinely used in structural maintenance problems, and the advanced point-based solvers capable of solving large scale, realistic problems. Our approach in this paper is helpful for understanding shortcomings of the currently used methods, related complications, possible solutions and the significance different solvers have not only on the solution but also on the modeling choices of the problem. In the second part of the study we utilize almost all presented topics and notions in a very broad, infinite horizon, minimum life-cycle cost structural management example and we focus on point-based solvers implementation and comparison with simpler techniques, among others.

© 2014 Elsevier Ltd. All rights reserved.

1. Introduction

In this paper the framework of planning and making decisions under uncertainty is analyzed, with a focus on deciding optimum maintenance and inspection actions and intervals for civil engineering structures based on the structural conditions in real time. The problem of making optimum sequential decisions has a huge history in a big variety of scientific fields, like operations research, management, econometrics, machine maintenance, control and game theory, artificial intelligence, robotics and many more. From this immense range of problems and methods we carefully chose to analyze techniques that can particularly address the engineering and mathematical problem of structural management, and we also present them in a manner that we think is most appropriate for the potential

interested readers, who are dealing with this particular problem and/or structural safety.

A large variety of different formulations can be found addressing the problem of maintenance and management of aging civil infrastructure. In an effort to very succinctly present the most prevalent methodologies we classify them in five different general categories. The first category includes methods that rely on simulation of different predefined policies and indicative works can be found by Englund and Sorensen [1] and Alipour et al. [2]. Based on the simulation results, the solution that provides the best performance among these scenarios is chosen, which could be the one with the minimum cost or cost/benefit ratio, etc. It is evident that a problem with this approach is that the chosen policy, although better than the provided alternatives, will hardly be the optimal among all the possible ones that can actually be implemented. In the second category we include methods that are usually associated with a pre-specified reliability or risk threshold and several different procedures have been suggested in the literature. In Deodatis et al. [3] and Ito et al. [4] the structure is

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

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

maintained whenever the simulation model is reaching the reliability threshold, while in Zhu and Frangopol [5] the same logic is followed with the exception that the maintenance actions to take at the designated times are suggested by an optimization procedure. Thoft-Christensen and Sorensen [6] and Mori and Ellingwood [7] pre-assume a given number of lifetime repairs, in order to avoid the discrete nature of this variable in their non-linear, gradient-based optimization process, and based on their modeling they identify optimum maintenance times so that the reliability remains above the specified threshold. Zhu and Frangopol [5] also followed this approach but used a genetic algorithm, which has significant computational cost however, in order to drop the assumption of pre-determined number of lifetime repairs and to be able to model the available maintenance actions in a more realistic manner. Overall, the available methods in this category provide very basic policies and the simultaneous use of optimization algorithms in a probabilistic domain, in this context, usually compels use of rudimentary models. Unfortunately, this last statement, concerning a probabilistic domain, is also valid when the problem is cast in a generic optimization formulation, which we characterize as another category although the work in [5] would also fit in. Formulations in this class usually work well with deterministic models, the available number of possible different actions is typically greater than before and a multi-objective framework is enabled. The problem is frequently solved by genetic algorithms and a Pareto front is sought. The choice of genetic algorithms, or other heuristic search methods, for solving the problem is not accidental since these methods can also tackle the discrete part of the problem, like the number of lifetime actions and the chosen action type in each maintenance period. Unavoidably, the computational cost is significant nonetheless and probabilistic formats are problematic with these techniques. Representative works can be seen in [8–10], among others.

All presented methods until now rely exclusively on simulation results and in essence do not take actual data into account in order to adjust or determine the performed actions, with the works in [3,4] being some sort of exception. While this may be sufficient for a variety of purposes, it is definitely incongruous for an applied, real world structural management policy. To address the issue a possible approach is suggested in the literature which is typically, but not utterly, associated with condition based thresholds. We classify these methods in a fourth category and a representative work can be seen in Castanier et al. [11]. The main idea behind these methods is to simulate deterioration based on a continuous state stochastic model, with Gamma processes being a favored candidate, and to set certain condition thresholds based on optimization, in between which a certain action takes place. Assuming perfect inspections, the related action is thus performed as soon as the structure exceeds a certain condition state during its lifetime. As probably understood already, the main weakness of this formulation is the usually unrealistic assumption about perfect observations. Due to this, although capabilities of the formulation are generally broad and versatile, including probabilistic outcome and duration of actions, the inspection part is lacking important attributes and analogous sophistication with other parts of the approach. A secondary concern with this approach can be also identified in the fact that the global optimum may be hard to find in non-convex spaces, although this is not a general limitation and is dependent on the specifics of the problem and the optimization algorithm used.

In the fifth category we include models that rely on stochastic control and optimum sequential decisions and these are the models of further interest in this paper. These approaches usually work in a discrete state space, and like the ones in the previously described category also take actual, real-time data into account in order to choose the best possible actions. In their most basic form

of Markov Decision Processes (MDPs) these models share the limitation of perfect observations, although they can generally provide more versatile, non-stationary policies, and taking advantage of their particular structure the search for the global optimum is typically unproblematic. Indicative of the successful implementation of MDPs in practical problems, Golabi et al. [12] and Thompson et al. [13] describe their use with fixed biannual inspection periods in PONTIS, the predominant bridge management system used in the United States. Most importantly however, as is also shown in detail in this paper, MDPs can be further extended considerably to a large variety of models and especially to Partially Observable Markov Decision Processes (POMDPs) that can take the notion of the cost of information into account and can even address the conundrum of planning optimum policies based on uncertain structural data and stochastic models. We believe that POMDP based models are adroit methods with superior attributes for the structural maintenance problem, in comparison to all other methods. They do not impose any unjustified constraints on the policy search space, such as periodic inspections, threshold performances, pre-determined number of lifetime repairs, etc., and can instead incorporate in their framework a diverse range of formulations, including condition-based, reliability and/or risk-based problems, periodic and aperiodic inspection intervals, perfect and imperfect inspections, deterministic and probabilistic choice and/or outcome of actions, perfect and partial repair, stationary and non-stationary environments, infinite and finite horizons, and many more. Representative works with a POMDP framework can be seen in Madanat and Ben-Akiva [14], Ellis et al. [15] and Corotis et al. [16], while further references about studies based on Markov Decision Processes are also given in the rest of this paper and in the second part of this work, [17].

To illustrate schematically a POMDP policy, with a minimum life-cycle cost objective, in a general, characteristic structural inspection and maintenance problem, Fig. 1 is provided. In this figure, the actual path of the deterioration process (continuous blue line) has been simulated based on one realization of a non-stationary Gamma process and is overall unknown to the decision-maker except when he decides to take an observation action. The gray area in Fig. 1 defines the mean ± 2 standard deviations uncertainty area which is given by the used stochastic model. This probabilistic outcome of the simulation model is the only base for maintenance planning for the decision-maker when actual observation data cannot be taken into account. Even with an accurate stochastic model, the fact that the actual deterioration process is never observed will usually result in non-optimum actions, for a certain structure, since the realized process can be, for example, in percentiles far away from the mean. Taking observation data into account the decision-maker can update his belief about the deterioration level of the structure according to his prior knowledge and the accuracy of observations. In Fig. 1 the belief updating is shown clearly based on the outcome of the first two different observation actions (marked with + in the figure). As seen, the first observation method is more accurate (probably at a higher cost), in comparison to the second, and directs more effectively to the true state of the system. Although rarely the case with structural inspection/monitoring methods, if a certain observation action can identify the state of the structure with certainty, the belief is then updated to this state with probability one. As is shown in detail in the rest of this paper, POMDPs plan their policy upon the belief state-space and this key feature enables them also to suggest times for inspection/monitoring and types of observation actions, without any restrictions, unlike any other method. Concerning maintenance actions, POMDPs can again optimally suggest the type and time of actions without any modeling limitations. Two different maintenance actions are shown as an example in Fig. 1, marked by the red rectangles. The length of the rectangles

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