Environmental Modelling & Software 99 (2018) 39–51

Contents lists available at ScienceDirect

Environmental Modelling & Software

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

Policy tree optimization for threshold-based water resources management over multiple timescales

Jonathan D. Herman [a,*], Matteo Giuliani [b]

[a] Department of Civil & Environmental Engineering, University of California, Davis, CA, USA
[b] Department of Electronics, Information, and Bioengineering, Politecnico di Milano, Milan, Italy

ARTICLE INFO

Article history:
Received 12 June 2017
Received in revised form
14 September 2017
Accepted 29 September 2017

ABSTRACT

Water resources systems face irreducible uncertainty in supply and demand, requiring policies to respond to changing conditions on multiple timescales. For both short-term operation and long-term adaptation, thresholds or "decision triggers", where a policy links observed indicators to actions, have featured prominently in recent studies. There remains a need for a general method to conceptualize threshold-based policies in an easily interpretable structure, and a corresponding search algorithm to design them. Here we propose a conceptual and computational framework where policies are formulated as binary trees, using a simulation-optimization approach. Folsom Reservoir, California serves as an illustrative case study, where policies define the thresholds triggering flood control and conservation actions. Candidate operating rules are generated across an ensemble of climate scenarios, incorporating indicator variables describing longer-term climate shifts to investigate opportunities for adaptation. Policy tree optimization and corresponding open-source software provide a generalizable, interpretable approach to policy design under uncertainty.

1. Introduction

Control policies for water and environmental systems must be designed at timescales for which projections of exogenous variables remain highly uncertain (Pahl-Wostl, 2007). This challenge is primarily driven by nonstationary climate, but is amplified by uncertainties in hydrologic and human systems (Hallegatte, 2009; Brown et al., 2015). The resulting "cascade of uncertainty" (Wilby and Dessai, 2010) inhibits traditional planning methods and has prompted the development of non-probabilistic approaches to vulnerability assessment and policy design (Dessai and Hulme, 2004; Kwakkel et al., 2016). Much of the work in this area has focused on long-term adaptation measures, such as infrastructure sequencing (Mortazavi-Naeini et al., 2014; Beh et al., 2015; Zeff et al., 2016). However, changes in short-term management policies, such as reservoir operations, urban water conservation, and agricultural practices, have also been explored as alternative or complementary adaptation measures (Giuliani et al., 2016c; Culley et al., 2016). The role of operating policies in adaptation is increasingly recognized, as the performance of a system is often driven by the accumulation of short-term decisions. Moreover, the extreme events for which policies are designed are often the most difficult to project with certainty, requiring robust policies able to withstand unanticipated conditions (Walker et al., 2001; Lempert, 2002; Herman et al., 2015; Giuliani and Castelletti, 2016).

A common theme among studies of both long-term adaptation and short-term management policies has been the use of thresholds or "decision triggers", where a system changes course in response to observations. The long-term adaptation problem is perhaps best framed by Haasnoot et al. (2013) as a choice from a set of candidate pathways. In the water resources field, this most commonly refers to climate adaptation, the "response to observed or expected changes in climatic stimuli" (Adger et al., 2005), implying decadal timescales or greater. By contrast, in the socio-environmental systems field, adaptation may refer to any action taken to manage a system, including responses to short-timescale stimuli (Walker et al., 2004; Folke et al., 2010). Most reservoir operating policies adapt to short-term observations via feedback control loops, which explicitly condition management decisions on the system state (Castelletti et al., 2008); embedding long-term adaptation mechanisms in operating policies remains an open issue. Regardless of timescale, designing decision thresholds relies

* Corresponding author.
E-mail addresses: jdherman@ucdavis.edu (J.D. Herman), matteo.giuliani@polimi.it (M. Giuliani).

https://doi.org/10.1016/j.envsoft.2017.09.016
1364-8152/© 2017 Elsevier Ltd. All rights reserved.

on the concept of tipping points, which define the conditions under which a current or planned policy will cease to meet its objectives (Kwadijk et al., 2010). These tipping points may then be translated to a scenario-dependent ending date for each policy (Haasnoot et al., 2012). In a nonstationary climate, this adaptive mechanism allows the management policy to respond to new information as it becomes available while remaining robust to unexpected change (Hall et al., 2014). Computational methods to generate adaptive pathways have included evolutionary algorithms, with mixed-integer formulations to explore the combinatorial problem of sequencing actions on a fixed time interval (Kwakkel et al., 2015).

For shorter operational timescales, policies with decisions triggered by observations are fundamental in water resources (e.g., Young, 1967). Reservoir operating rules often involve discrete thresholds leading to a release decision (Shih and ReVelle, 1995), either for water supply or flood control. Similarly, in urban water management studies, decision triggers have been used to implement conservation measures or activate alternative supplies (Mortazavi et al., 2012; Zeff et al., 2014; Mortazavi-Naeini et al., 2015; Borgomeo et al., 2016). These studies in particular have used simulation-optimization approaches where the structure of the decision rule is fixed, and threshold values are searched as continuous decision variables. The design of short-term policies can play an important role in long-term adaptation, which might involve redefining policies more suitable to projected climate scenarios (Giuliani et al., 2016a) or the inclusion of longer-timescale information in the policy itself.

The concept of decision triggers holds promise for control problems across multiple timescales. To generalize, a threshold-based policy combines a set of candidate actions with a set of trigger conditions, defined by indicator variables and their threshold values. Several important challenges have been recognized, including the design and selection of indicators in cases where multiple observed variables are expected to influence the decision (Groves et al., 2015; Haasnoot et al., 2015; Giuliani et al., 2015). Accounting for this challenge requires a method in which the structure of the rule can be searched in addition to the threshold values. There remains a need to generalize the design of threshold-based policies without prespecifying the threshold values, sequence, or timing of actions—in other words, to optimize the set of conditions under which certain actions are taken, rather than the actions themselves.

In this paper, we frame the design of management policies over multiple timescales as a simulation-optimization problem in which the policy is represented as a binary tree. The tree structure allows actions to be conditioned on multiple indicator variables characterizing current system conditions as well as longer-term climatological variables. This combination of short- and long-term information improves the ability of the resulting solutions to adapt across a range of uncertain futures. In addition, this approach facilitates the interpretation of the optimized policies by decision makers, because the logical rules triggering a specific action can be easily visualized. This study focuses on triggering reversible actions, such as flood control and hedging operations, rather than irreversible ones such as infrastructure investment. Although precedent exists for the latter, statistical concerns arise for actions only triggered once—a topic for future study. We demonstrate the policy tree optimization method using Folsom Reservoir, California as an illustrative case study. First, we test the algorithmic performance under historical conditions against solutions designed with traditional control methods. The result serves as the baseline policy without any long-term adaptation measures. We then use the method to develop policy trees capable of navigating a wide range of plausible climate change scenarios. The resulting policies suggest how operating rules and thresholds might be adapted in the future despite highly divergent projections of water supply risk due to nonstationary climate.

2. Methods

2.1. Problem statement

Managed water systems can be modeled with the state transition equation:

$$\mathbf{s}_{t+1} = f(\mathbf{s}_t, \mathbf{u}_t, \mathbf{q}_{t+1}), \qquad t \in [0, H] \tag{1}$$

where $\mathbf{s}$, $\mathbf{u}$, and $\mathbf{q}$ are vectors of state, decision, and disturbance variables, and H is the evaluation horizon. The naming convention reflects a standard reservoir management problem where these are storage, release decision, and inflow vectors, respectively, though the approach can be applied to other types of policy design problems with different control variables. The time subscript of each variable denotes the time instant at which it assumes a deterministic value: the reservoir storage is measured at time t and thus is denoted as s_t, while inflow in the interval $[t, t+1)$ is denoted as q_{t+1} because it can be known only at the end of the time interval. The goal is to choose the sequence of release decisions $\mathbf{u}_t$ for $t = 0, \ldots, H-1$ that minimizes a cost function $J(\mathbf{s}_0, \mathbf{u}_0^{H-1}, \mathbf{q}_1^H)$ for the initial condition $\mathbf{s}_0$ and a particular inflow series $\mathbf{q}_1^H$, subject to the state transition (Equation (1)).

A common approach to determine the optimal sequence of release decisions is to define an operating policy as a parametric function P_θ that provides the release decisions at each timestep t given the current system conditions $\mathbf{x}_t$ (i.e., state variables along with additional indicator variables). In the literature, a number of operating rules based on storage and inflow have been proposed (e.g., Oliveira and Loucks, 1997; Lund and Guzman, 1999). An alternative to these rules is represented by the use of nonlinear approximating networks (e.g., Raman and Chandramouli, 1996; Busoniu et al., 2011; Giuliani et al., 2014), which provide more flexibility to the operating policy.

Given the selected parameterization of the operating policy P_θ, the exploration of the parameter space $\theta \in \Theta$ leads to the policy that optimizes the cost function $J(\cdot)$:

$$\theta^* = \mathrm{argmin}_\theta J\left(\mathbf{s}_0, P_\theta, \mathbf{q}_1^H\right) \tag{2}$$

This direct policy search approach originated in the artificial intelligence field (Rosenstein and Barto, 2001) and has been adopted in water resources as Parameterization-Simulation-Optimization (Koutsoyiannis and Economou, 2003), recently extended for multi-objective operations problems (Giuliani et al., 2016b, 2017; Quinn et al., 2017). The choice of the parameterized function family may be connected to the choice of the policy input variables: policy search is both an information selection problem (Hejazi et al., 2008; Giuliani et al., 2015) as well as a real-valued optimization problem (e.g., Zatarain-Salazar et al., 2016).

2.2. Policy trees

In this study, we represent policies as binary trees, replacing the real-valued optimization problem in Equation (2) with a search for a set of hierarchical rules to select decisions at each timestep from a discrete set of actions, $a_i \in \mathscr{A}$ (Fig. 1). Each action can be thought of as a different operating mode (e.g., flood control, water supply hedging) selected according to the current system conditions $\mathbf{x}_t$. We would like to search for the optimal policy tree T^* such that the sequence of actions $a_t^* = T^*(\mathbf{x}_t)$ minimizes the cost function subject to the dynamics of the system, i.e.