Journal of Hydrology 519 (2014) 1836–1847

Contents lists available at ScienceDirect

Journal of Hydrology

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

First-passage time criteria for the operation of reservoirs

Aiden J. Fisher [a,*], David A. Green [a], Andrew V. Metcalfe [a], Kunle Akande [b]

[a] School of Mathematical Sciences, University of Adelaide, SA 5005, Australia
[b] CH2MHILL, Burderop Park, Swindon, Wiltshire SN4 0QD, United Kingdom

ARTICLE INFO

Article history:
Received 30 June 2014
Received in revised form 19 September 2014
Accepted 21 September 2014
Available online 2 October 2014
This manuscript was handled by Geoff Syme, Editor-in-Chief

Keywords:
Multi-objective optimisation
Pareto-front
First-passage time criterion
Phase-type distribution

SUMMARY

A multi-objective optimisation for reservoir operation based on expected monetary value and expected first passage-time criterion is proposed. The computations are facilitated by the algorithms of matrix analytic methods. The formal structure, classifying states as levels and phases within levels, and associated algorithms of matrix analytic methods are introduced in the context of multi-reservoir systems. The algorithms underpin the feasibility of the computations for large systems and enable the calculation of the full distribution of first passage time. A new algorithm for computing results for a seasonal model, which reduces computing time by an order of magnitude for monthly time steps is presented. The methods are illustrated for a two reservoir system, with an option of pumping additional water from a transfer scheme, in the East of England. The Pareto front of Pareto optimal policies is shown.

1. Introduction

The optimal operation of reservoir systems, in terms of water releases, has been the subject of research for over fifty years, and given the stochastic input to reservoirs, the process is typically modelled as a Markov decision process (MDP). Howard (1960, pg. 128–137) implemented a stochastic dynamic programming (SDP) solution for MDPs. However, the need for algorithms to control increasingly large multi-reservoir systems providing hydro-power as well as water supply to different categories of users (e.g. Tejadaguibert et al., 1993; Huang et al., 2002; Archibald et al., 2006; Reddy and Nagesh Kumar, 2006) ensures that it remains an active area of research.

Stochastic dynamic programming typically focuses on expected monetary value (EMV) and does not include any measure of risk in the objective function (Mahootchi et al., 2010). To determine the optimal policy, the agency administering the resource (usually a government body) will typically assign value to the use of the water based upon various market factors and regulation, choosing the optimal policy based on EMV. The determined optimal policies for the management of water resources is sensitive to the relative value attributed to the needs of a potable water supply, agriculture, industry, the environment, as well as recreation and transport (Webby et al., 2009). The value of a water resource can be highly subjective. Water being used for industry is highly valuable, but using the same price for agriculture would be too prohibitive due to the large quantity needed to produce a relatively low value product. Nevertheless agriculture on a national scale is far too important to neglect. Young and McColl (2005) proposed that it is better to have a proportional allocation of the water remaining after essential human needs are met. That is, water resource management should be focused on maintaining a storage level to guard against future shortfalls and/or having a freeboard sufficient to protect against flooding, rather than just maximising EMV. In addition the policy maker may have other uses for water storage, such as, they may wish to maintain storage for recreational use, environmental flows, or to maintain a wetland environment. To meet these objectives based on water storage levels, rather than economic considerations, in this paper we propose an alternative criterion to EMV, based on expected first-passage time to a set of states.

The expected first-passage time is a property of the Markov chain modelling used in MDPs and as such eliminates the need for utility rewards typical of the EMV optimisation. We propose a new algorithm to find maximal or minimal expected first-passage time criterion (FPTC) that exploits the Markov chain basis of the MDP. To the best of the authors' knowledge a first-passage time criterion has not previously been used in an MDP to determine an optimal operating policy of reservoir storages, although it has been proposed in a general reliability context by Jianyong and

* Corresponding author.
E-mail addresses: aiden.fisher@adelaide.edu.au (A.J. Fisher), david.green@adelaide.edu.au (D.A. Green), andrew.metcalfe@adelaide.edu.au (A.V. Metcalfe), Kunle.Akande@ch2m.com (K. Akande).

http://dx.doi.org/10.1016/j.jhydrol.2014.09.061
0022-1694/

Huang (2001). First-passage times are used in the Gould matrix method (Gould, 1961) to find the mean recurrence time of a reservoir going from full to empty (McMahon et al., 2007). Here we present a general framework that can also perform these calculations. We also give a new application for the phase-type distribution, which was first discussed in a hydrological setting by Fisher et al. (2010).

The optimal FPTC policy is contrasted with the optimal policy found using EMV and a Metropolis–Hastings type algorithm is proposed to explore for other policies that are Pareto optimal for both of these criteria. This strategy provides the decision maker with a range of policies from maximising EMV through to maintaining the storage and shows the reduction in EMV as more emphasis is placed on maintaining the storage. An application to a system of two storage reservoirs in Suffolk and Essex counties of South East England is described.

2. Mathematical methods

Discrete time Markov chains have been a popular choice for modelling reservoir systems since Moran's influential monograph (Moran, 1955). A discrete time Markov chain is a random process in which the probability of moving from one state to another in a discrete time step is dependent only on the present state and not on the history of the process before the present state.

2.1. Markov decision processes

For the purposes of this paper, a discrete-time Markov decision process is defined by 4 element object (X, D, P, R) and the unit of time step, where;

- X a finite state space representing the physical state of the system,
- D a discrete set of decisions that affect the system,
- P a probability transition matrix dependent on the decisions,
- R a matrix of relative rewards (a real number) for arriving in all states $j \in X$ from $i \in X$ under each decision $d \in D$.

The choice of the decision $d \in D$ can have two effects, it can change probabilities in P and it can also change the rewards in R. The probabilities and rewards under decision, d will be denoted $P(d)$ and $R(d)$ respectively.

The total reward is maximised at each time step for all states in X, by choosing a decision $d \in D$ that will give the highest expected monetary value (EMV) based on the highest EMV at the next time step. This is expressed in the form of a Bellman equation. Denote the EMV of being in state $i \in X$ at time t as $v_t(i)$. The Bellman equation is used to determine the maximum EMV, $v^*_{t-1}(i)$, at time $t-1$ given the maximum value, $v^*_t(i)$, at time t and is of the form,

$$v^*_{t-1}(i) = \max_{d_{t-1}} \left\{ \sum_j p_{ij}(d_{t-1}) \left[r_{ij}(d_{t-1}) + v^*_t(j) \right] \right\}, \qquad (1)$$

where $p_{ij}(d_{t-1})$ and $r_{ij}(d_{t-1})$ are the probabilities and reward respectively for the transition from state $i \in X$ at time $t-1$ to state $j \in X$ at time t under decision $d_{t-1} \in D$ made at time $t-1$.

Bellman's equation can be used to provide a stochastic dynamic programming (SDP) solution for the optimisation of a system of reservoirs, using a maximum EMV criterion, by working backwards in time from some final time T. Consider the decision problem from state i at time $T-1$. The problem specification includes rewards for arriving in each state j in X under each possible decision d_{T-1} together with transition probabilities from state i to j under each decision. The decisions that maximise the EMV can be determined, and this maximum EMV is written as $v^*_{T-1}(i)$ where

$$v^*_{T-1}(i) = \max_{d_{T-1}} \sum p_{ij}(d_{T-1}) r_{ij}(d_{T-1}),$$

the $p_{ij}(d_{T-1})$ being the transition probabilities of going from state i to state j under the decision d_{T-1}, and $r_{ij}(d_{T-1})$ being the reward for arriving in state j under decision d_{T-1}.

Now consider the decision problem at time $T-2$. The total value of a move from a state i to state j is the sum of the reward for arriving in state j under decision d_{T-2} and the maximised EMV in state j at time $T-1$. The decisions that maximise the total EMV are determined and this maximum EMV is written as $v^*_{T-2}(i)$. By recursion the general statement, known as Bellman's equation, in Eq. (1) assigns no value to ending in state j other than the reward for arriving in it. This is immaterial as Bellman's equation is solved recursively back from T until a fixed point where

$$d_{T-1} = d_t,$$

for all i in X, and this stable decision set is independent of any values assigned for ending in particular states.

In most hydrological applications seasonality needs to be accounted for and Butcher (1971) suggested a method for solving MDPs which have a seasonal component to them. In the model proposed, a decision process is expanded to $(X, D, P_0, \ldots, P_{s-1}, R_0, \ldots, R_{s-1})$, where each P_k is the probability transition function at time $k \equiv t \pmod s$ for a transition over the time interval $(t-1, t]$ and R_k is the reward matrix at time $k \equiv t \pmod s$, where s is the number of seasons in the cycle.

As the matrices are now time dependent, Eq. (1) becomes,

$$v^*_{t-1}(i) = \max_{d_{t-1}} \left\{ \sum_j p_{k,ij}(d_{t-1}) \left[r_{k,ij}(d_{t-1}) + v^*_t(j) \right] \right\}, \quad \text{where} \quad k \equiv t \pmod s.$$

However the Markov chain is now periodic. As a result when solving backwards the process is stopped when

$$d_{(t-1)\times s+m} = d_{t\times s+m},$$

for all $i \in X$ and $m = 1, \ldots, s$.

A policy is a vector denoted $\boldsymbol{u}$, where the elements of $\boldsymbol{u}$, $u_i \in D$, are the decisions that would be made in states i. Let $P(\boldsymbol{u})$ be the probability transition matrix for a reservoir operating under a policy $\boldsymbol{u}$, where the probability of transition from state $i \in S$ to state $j \in S$, in one time step, is the element in the ith row and jth column, denoted $p_{ij}(u_i)$. The MDP $(X, D, P_0, \ldots, P_{s-1}, R_0, \ldots, R_{s-1})$ is then solved iteratively until the same season's optimal decision for all states is made for two successive years. This optimal policy will be denoted u^*, where $u^*(i,k) \in D$ will be the optimal decision for state $i \in X$ in season k.

In the case of a dynamic programming model for a multi-reservoir system, with sufficient discretisation for realistic operation and seasonal variation included, the number of states may be very large indeed. One potential way of dealing with this is the adaptation of established matrix analytic methods used primarily in telecommunications and queueing models. These have been introduced in hydrological settings by Fisher et al. (2010) for ephemeral stream flow, and Piantadosi et al. (2010) for a sequence of three dams. Here the matrix analytic structure is further exploited to take account of seasonality and develop succinct methods of calculating measures of optimality.

2.1.1. Matrix analytic methods

Matrix analytic methods in discrete time extend the scope of the simple Markov chain by allowing other than geometric times between events, while retaining analytical tractability. The princi-