ARTICLE IN PRESS

Energy Policy xxx (xxxx) xxx–xxx

Contents lists available at ScienceDirect

Energy Policy

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

Reliability-constrained hydropower valuation☆

Antony Ware

The University of Calgary, 2500 University Avenue SW, Calgary, Alberta, Canada T2N 1N4

ARTICLE INFO

Keywords:
Hydropower
Electricity markets
Renewable energy
Valuation

ABSTRACT

Maximizing the long-term value of hydropower generation requires management of uncertain reservoir inflows, potentially variable constraints on outflows, and exposure to possibly wildly varying power prices. We describe a stochastic dynamic programming approach to the quantification of reservoir reliability (for example, measures of the risk of over-topping the reservoir or failing to satisfy downstream flow requirements) and a related approach to determining the reservoir flow strategy that maximizes expected revenue, subject to defined target reliability levels.

1. Introduction

Hydropower generation has an important role to play as our energy systems transition towards increasing reliance on renewable sources of power. Apart from representing in themselves a major renewable energy source, the fact that hydropower facilities can be dispatched extremely rapidly (in contrast to nuclear, coal-fired and even natural gas plants) means that they can help to offset the problems of intermittency that are associated with wind and solar power. In addition, when pump capability is included, the facility operates as a high-capacity battery offering storage benefits to the system.

The flexibility offered by hydropower facilities has limits. There are various kinds of constraints on the rate at which water can be allowed to flow from the reservoir. These limits arise from consideration of both the downstream impact of flow rates and the current level of water in the reservoir and the *anticipated* impact of flowing (or not flowing) water now on the reservoir reliability, i.e. its expected *future* ability to satisfy flow constraints.

Three key uncertain quantities affect the long-term reliability and value of the reservoir. Firstly, the *amount of water* flowing into the reservoir varies - from year to year and from day to day. Second, the *revenues available* on the energy and ancillary markets (used by the system operator to ensure stability of the system over short time scales) vary from moment to moment. Third, the range of *allowable flow rates*, which will be affected by the variable state of the river downstream from the reservoir, as well as on other factors such as plant outages. None of these quantities can be predicted with certainty for any specific future date, particularly more than a few days in advance.

The primary goal of the model we present is to provide quantitative estimates of both the reliability of the system and the marginal value of the water in the system. In practice, this value is generated by revenue from energy markets, ancillary markets (used by the system operator to ensure stability of the system over short time scales) and, sometimes, capacity markets (although in this work we consider only the energy revenue). An important component of the value is the the *optionality* that having water in the reservoir system provides - the flexibility to choose when to flow the water in order to take advantage of high prices. In this paper, the problem of determining how best to exploit that flexibility to maximize both the expected reliability and the revenue in the face of the uncertainties inherent in the system is formulated as a problem of stochastic optimal control.

In the remainder of this section we review some of many applications of stochastic optimal control problems in the management of hydropower systems, and introduce the particular hydropower facility we will use as an example to illustrate the proposed modelling approach. In Section 2 we consider the initial problem of maximizing the reliability of the system. This results in a 'safety-first' strategy, and allows us to quantify the risk of failing to satisfy the operational constraints. In Section 3 we show how this can then be incorporated into a revenue-optimizing strategy, constrained by the reliability. The final section contains some concluding remarks.

1.1. Stochastic dynamic programming for hydropower systems

The use of stochastic dynamic programming in the context of hydropower systems has a long history (see Yakowitz, 1982 for an early review, and Wallace and Fleten, 2003 for a slightly more recent overview of applications to hydropower and other energy systems). The work of Pereira and Pinto (1985), Pereira and Pinto (1991) has had a significant influence on this field; their introduction of a stochastic dual

☆ *"This article is part of a Virtual Special Issue entitled 'Energy and Environment: Transition Models and New Policy Challenges in the Post Paris Agreement'."*

E-mail address: aware@ucalgary.ca.

https://doi.org/10.1016/j.enpol.2018.03.053
Received 31 July 2017; Received in revised form 23 March 2018; Accepted 24 March 2018

Please cite this article as: Ware, A., Energy Policy (2018), https://doi.org/10.1016/j.enpol.2018.03.053

dynamic programming approach means that optimization problems associated with complex hydropower systems can be tackled. The approximate encoding of piecewise linear constraints allows them to lessen the impact of the 'curse of dimensionality'. See Rebennack et al. (2010a, 2010b) for many more recent examples, and Powell (2007) for a comprehensive overview of approximate dynamic programming and its applications.

Näsäkkälä and Keppo (2008) consider a Monte Carlo approach to studying the medium- and long-term planning problem in the face of inflow and price uncertainty using a simple parametrization for the optimal production strategy, and assuming that the power is sold using futures contracts with various tenors. Carmona and Ludkovski (2010) also work in a Monte Carlo setting, making use of an optimal stopping formulation of the stochastic dynamic programme.

Zhao and Davison (2009a), Zhao and Davison (2009b) use stochastic optimal control to study the operation and valuation of pump storage systems, assuming time-varying but deterministic prices, and develop an approach for estimating the value of accurate hydrological forecasts. Conversely, Bayón et al. (2009) develop an approach for forecasting short-term prices, and study the impact of the quality of such forecasts on optimal hydropower operation. Chen and Forsyth (2008) also consider deterministic inflow rate but stochastic prices. They study a finite difference method for solving the associated Hamilton-Jacobi-Bellman partial integro-differential equation, and prove the convergence of the numerical scheme to the viscosity solution. They are able to incorporate various operational constraints, and conclude that failing to consider these 'may considerably overestimate the value of hydroelectric power plant cashflows'.

In this paper, stochastic prices, inflows and downstream constraints are incorporated, and an approximate stochastic dynamic programming approach developed in order to quantify reservoir reliability and develop an operation strategy that maximizes expected revenue while maintaining a given level of reliability.

1.2. Bighorn

In this work, we use the Bighorn hydropower facility as the setting for the modelling and valuation approach we describe. This facility is owned and operated by TransAlta Corp., and is situated on the North Saskatchewan River in the foothills of the Rocky Mountains, west of Red Deer in southern Alberta, Canada. It feeds from Abraham Lake, which holds more than 1.1 million acre-feet of water in the operating range of the 120 MW hydropower facility. The water it releases passes through Rocky Mountain House and Edmonton on its way—eventually—to Hudson Bay. Inflow into the reservoir (i.e. Abraham Lake) is highly seasonal. In the winter, there is little glacial melt, and any upstream precipitation is in the form of snow and has little immediate impact on water levels in the reservoir. The highest inflows come during the spring runoff, when warmer temperatures lead to melting snow and ice, and more precipitation in the form of rainfall. This is evident in the 1985–2011 inflow data[1] shown in Fig. 1, where the top graph shows cumulative annual inflows, and the lower graph shows daily inflow amounts. These data should be viewed with a modicum of scepticism, however, since reliable measurements of inflow amounts over short time periods are particularly difficult to obtain. This is particularly evident during the winter months, where some days appear to show significant volume reduction, far in excess of what might be lost through processes such as sublimation.

There are some minimum and maximum flow constraints that must be satisfied. The minimum flow constraints arise from the requirement to provide sufficient water supply to the downstream users. Maximum flow constraints arise from the need to respect the carrying capacity of the downstream river system. Added to this, in Alberta, temperatures in the winter months can fall to a level low enough that ice forms on the river. When this happens, flow rates must be kept steady in order to allow a 'clean' ice cover to form. Once this has happened, flow rates can be steadily increased to 'push' the ice up and create a natural tunnel, which, once established, will allow variable flows to occur underneath the ice.

2. Reliability

In this paper, we seek to quantify the reliability of a hydropower system. Koutsoyiannis (2005) discusses various formulations of the reliability of such a system, and our quantification can be seen as an extension of one setting of his framework. We consider failure of the system to have occurred if the reservoir overtops its limits (so that water is spilled in an uncontrolled fashion), or if the reservoir becomes empty, and so is unable to satisfy minimum flow constraints. For a given operation strategy, we consider the probability that failure will be avoided over a time horizon T. We define the *reliability* to be the maximum possible probability over all possible operation strategies. Any strategy that results in this maximum reliability we refer to as a *safety-first* strategy.

2.1. Modelling inflows and river states

The rate of inflow $I(t)$ on day t (measured in units of years) is modelled as a log-normal random variable. Specifically, we have

$$\ln I(t) \sim N(\beta(t), \sigma^2(t)), \tag{1}$$

where the mean, $\beta(t)$, and the variance, $\sigma^2(t)$, are seasonally-varying functions. They are modelled using trigonometric functions with frequencies up to N and N_σ:

$$\beta(t) = \beta_0 + \beta_1 \cos 2\pi t + \beta_2 \sin 2\pi t + \cdots + \beta_{2N-1} \cos 2\pi N t + \beta_{2N} \sin 2\pi N t,$$

and

$$\sigma^2(t) = c_0 + c_1 \cos 2\pi t + c_2 \sin 2\pi t + \cdots + c_{2N_\sigma-1} \cos 2\pi N_\sigma t + c_{2N_\sigma} \sin 2\pi N_\sigma t.$$

The parameters β_i and c_i are determined from a positive subset of historical (daily) inflows $\hat{I}(t)$ by choosing them so as to maximize the log-likelihood

$$\sum_t -\ln(2\pi\sigma^2(t)) - \frac{(\ln \hat{I}(t) - \beta(t))^2}{2\sigma^2(t)}.$$

Fig. 2 shows the fitted inflows, along with percentiles generated by the functions $\beta(t)$ and $\sigma^2(t)$, with $N_\beta = 6$ and $N_\sigma = 4$.

As described in Section 1, the ice cover downstream from the reservoir has a significant impact on the rate at which water can be released. This varies in an unpredicable way from year to year, with ice typically forming somewhere between November and January, and staying anywhere from a few weeks to a few months before breaking up in the spring.

When ice is in the process of forming, it is important to keep the flow consistent so that a stable 'ice bridge' can form, under which water can continue to flow at a potentially higher rate.

We use a continuous-time discrete-space Markov chain to model the state of the ice. We identify three ice states: 1: 'clear', 2: 'forming' and 3: 'solid'. Under these states, flows are constrained to be within the limits shown in Table 1 (given in cubic feet per second (cf/s)).

The probability of changing from one ice state to another is determined by a transition rate λ_{ij}, so that between two times t and $t + h$ the probability of a change from state i to state j is $\lambda_{ij} h + o(h)$. We make each λ_{ij} time-dependent so as to capture the way in which the probability of being in each of the various ice states changes depending on the time of year. Fig. 3 show the transition rates and resulting

[1] Data available from the HYDAT database https://ec.gc.ca/rhc-wsc/default.asp?n=9018B5EC-1.