Title

An Ensemble Optimization Framework for Coupled Design of Hydropower Contracts and Real-Time Reservoir Operating Rules

Journal
Water Resources Research, 54(10)

ISSN
0043-1397

Authors
Sahu, Reetik Kumar
McLaughlin, Dennis B

Publication Date
2018-10-01

DOI
10.1029/2018wr022753

An ensemble optimization framework for coupled design of hydropower contracts and real-time reservoir operating rules

Enter authors here: Reetik Kumar Sahu[1], and Dennis B. McLaughlin[1]

[1]Department of Civil and Environmental Engineering, Massachusetts Institute of Technology, Cambridge, Massachusetts, USA

Revised 1st October 2018

Corresponding author: Dennis B. McLaughlin (dennism@mit.edu)

Key Points:

- Hydropower reservoir revenue can be optimized by simultaneously adjusting contract specifications and the release operating rule.
- Predictive operating rules based on stochastic models of the reservoir and its inflows can perform better than standard operating rules.
- Stochastic performance assessments provide convenient measures of revenue uncertainty and facilitate quantitative performance comparisons of alternative operating strategies.

Abstract

Revenues from hydropower generation often depend on the operator's ability to provide firm power in the presence of uncertain inflows. The primary options available for optimizing revenue are negotiation of a firm power contract before operations begin and adjustment of the reservoir release during operations. Contract and release strategy optimization are closely coupled and most appropriately analyzed with stochastic real-time control methods. Here we use an ensemble-based approach to stochastic optimization that provides a convenient way to construct non-parametric revenue probability distributions to explore the implications of uncertainty. The firm power contract is a simplified bilateral fixed price agreement that partially insulates operator and buyer from price fluctuations. The release control laws and firm energy target are jointly optimized to maximize the operator's expected revenue. Revenue probability distributions and related spill performance statistics indicate that predictive operating strategies such as stochastic dynamic programming and model predictive control can give significantly better performance than standard deterministic operating rules. The performance obtained from batch optimization with perfect inflow information establishes a convenient upper bound on potential revenue and provides a baseline for assessing the significance of differences between real-time operating strategies. Sensitivity analysis indicates that the benefits of predictive operational strategies are greatest for reservoirs with medium non-dimensional residence times and less important for reservoirs with large residence times. Overall, probabilistic analysis of the coupled hydropower contract-operations problem provides a realistic way to assess revenue and risk for reservoirs that must provide firm power when inflows are uncertain.

1 Introduction

Hydroelectricity contributes 71% of global renewable electrical energy and 16% of total global electricity demand. Much of this energy is sold to institutional and industrial buyers under power purchase agreements that specify the price to be paid per unit energy for a firm amount of energy to be delivered over a designated time period. Firm power output is particularly important for industrial clients with predictable, possibly constant, energy demand. It is difficult for a hydropower facility to precisely track a firm power target since reservoir inflow and consequently power output are variable and uncertain. A hydropower operator must decide how much water to release at a given time without knowing with certainty how this will affect future power output. When selecting a contracted firm energy target, the operator must trade off the risk of having to purchase make-up power during low inflow periods when the target cannot be met vs. the risk of forgoing future income by releasing excess water during high inflow periods. Both of these situations can reduce revenue if the energy target is not properly chosen. Management decisions are further complicated when power may be obtained from multiple sources that have different characteristics and contractual arrangements. The various costs and demand functions encountered in these arrangements change continuously. Financial and economic uncertainties adds to the physical uncertainty contributed by fluctuations in reservoir inflows. All of these factors combine to make it difficult for a hydropower operator to determine the best way to manage a reservoir in real time or to identify the best strategy to pursue when negotiating a power purchase agreement.

In this paper we consider a simplified version of the hydropower firm power generation problem that enables us to focus on two particular factors that impact revenue performance: i) the operating strategy used to determine releases and ii) the firm power target. Our objective is to use insights from our probabilistic problem formulation to derive a release strategy and a firm power target that together maximize expected operator revenue when reservoir inflows are uncertain. We consider several different operating strategies that bracket the likely range of revenues that can be achieved with variable inflow and a given firm power target. These strategies include options that examine multiple replicates of future inflows to better anticipate the long-term effects of current releases. The power agreement we adopt specifies fixed prices over a long-term contract period and accounts for the decreasing marginal value of power. This type of agreement is becoming popular for renewable energy purchases by corporate entities with well-defined demands (Baker & A.McKenzie, 2015). Our use of a long-term fixed price agreement acknowledges the increasing desire among hydropower operators and buyers to reduce the risk associated with dependence on short-term market prices (Barroso et. al., 2006; Boneville Power Administration, 2013; World Energy Council, 2016).

Many optimization studies have addressed aspects of our problem formulation, including release operating strategies and power purchase agreements. Yeh (1985), Wurbs (1993), Labadie (2004), and Rani & Moreira (2010) provide comprehensive literature reviews on reservoir operating strategies. . In practice, most hydropower reservoirs are managed with deterministic operating rules that fall under the umbrella of a Standard Operating Policy (SOP) with hedging (Neelakantan & Pundarikanthan, 1999; Tu, Hsu, & Yeh, 2003; You & Cai, 2008). These rules typically relate the current reservoir release to the current reservoir storage and do not attempt to predict or adjust for future inflow variations.

Reservoir operators may be able to extract more benefit than can be achieved with Standard Operating Policies if they use decision rules that rely on probabilistic models of future inflows. The gold standard of this approach is Stochastic Dynamic Programming (SDP) (Bertsekas, 1995). The SDP method has been applied to the control of a single reservoir and to networks of multiple reservoirs (Zhao et al., 2014; Stedinger et al., 2013; Castelletti et al., 2007; Cervellera et al., 2006; Hall et al., 1968; Hooper et al., 1991; Yakowitz, 1982; Yeh, 1985). The popularity of SDP lies in its flexibility to accept a variety of objectives, constraints (equality and/or inequality), and random inflow models. Its main limitation is its computational complexity, which grows very quickly with the number of state and control variables used to describe the reservoir system (the so-called “curse of dimensionality’). This reflects the fact that SDP derives a general control law that specifies the optimal release at any time as a function of current state (Bertsekas, 1995). This control law depends on a reservoir inflow sequence that is specified for the entire operating period. In the stochastic version of dynamic programming the revenue to be maximized is averaged over an ensemble of many randomly sampled inflow realizations, using a version of Monte Carlo simulation, for each possible value of the current state. Several approximate SDP techniques have been developed to deal with the method’s computational demands. These take advantage of distinctive structural features applicable to reservoir operations problems (Bar-Shalom & Tse, 1974; Bertsekas, 1995; Bertsekas & Castañon, 1999).

Model Predictive Control (MPC) (García et al., 1989; Mayne et al., 2000; Rawlings, 2000) method is a limited look ahead real-time optimization technique that can use either deterministic or probabilistic inflow models. The stochastic version of MPC (SMPC) adopted here assesses the expected revenue for a given current release by averaging over an ensemble of random inflow replicates, following an approach similar to the one used in an ensemble SDP algorithm. An SMPC algorithm is generally less computationally intensive than SDP, and is able to readily handle complex constraints. This relative efficiency of SMPC reflects the fact that it optimizes the current release only for a particular (observed) current state rather than for all possible values of the state. SMPC plays an important role in process control where efficiency requires operating the system near specified bounds on the state and the control. Examples of relevant SMPC applications include process control (Arnold & Andersson, 2011), reservoir operations (Barjas Blanco et al., 2010; Linke, 2010; Tu et al., 2003), irrigation (Negenborn et al. 2009; van Overloop at al., 2008), and supply chain management (Perea-López et al., 2003; Qin & Badgwell, 2003). Here we consider the performance obtained with all three of the reservoir operating rules mentioned above (SOP, SDP, and SMPC) when they are coupled with a contract optimization procedure.

There is also an extensive literature on power purchase agreements and contracts. In particular, (Baker & A.McKenzie, 2015; ACORE, 2016; Shrestha et al. , 2005) discuss fixed price bilateral agreements that share important features with the one adopted in this paper. Many other types of contracts are available to manage operator and buyer risk. Examples relevant to hydropower applications are discussed in the financial risk literature (Catalão, Pousinho, & Contreras, 2012; Foster, Kern, & Characklis, 2015; Mo, Gjelsvik, & Grundt, 2001; Stickler et al., 2013) These include index methods that provide insurance to protect the operator from uncertainty. They do not generally consider the use of advanced operating strategies such as SDP and SMPC.

This paper synthesizes topics addressed in the literature cited above by examining connections between real-time operations and firm power contracts, with a focus on the effects of reservoir inflow variability. When developing a strategy for managing inflow uncertainty, it is best to determine the operating rule and contracted firm power target together, since they affect one another. The maximum revenue attainable by adjusting the operating rule depends on the firm power target and the maximum revenue attainable by adjusting the power target depends on the operating rule. We believe the synergy between operating strategy and firm power target is most easily examined with a bilateral fixed price contract. Such a contract partially insulates the operator and buyer from price fluctuations and is attractive when buyer demand is predictable and both parties seek to reduce risk from dependence on spot market prices.

2 Formulation of the reservoir operations problem

We examine connections between the firm power contract and operating strategy by considering a single purpose hydropower reservoir that provides energy to a single buyer according to a long-term bilateral contract agreed upon before operations start. The contract has a price structure that accounts for both power deficits and surpluses. The buyer agrees to pay a specified unit price ($ $MWhr^{-1}$) for the firm power, negotiated at the start of the contract period and held fixed until the end of the period. Recognizing that it may not always be possible to meet a particular firm power value when reservoir inflows are variable, the contract stipulates that shortfalls be covered by the buyer, who purchases makeup power at the market price and passes on a fixed unit charge to the operator. If the market price is below the shortfall charge the buyer benefits. If it is above then the operator benefits. Similarly, the buyer agrees to purchase surplus power from the operator for a negotiated fixed price. If this price is above the market price the operator benefits. If it is lower, then the buyer benefits. This contract arrangement has the advantage of providing both operator and buyer with a predictable pricing structure so that the only major source of operational uncertainty is inflow variability. We quantify this uncertainty with revenue probability distributions that apply for several different operating strategies.

In a discrete time problem formulation, the firm power value can be interpreted as an equivalent firm energy generated over a constant time step. Our coupled contract-operational design optimization focuses on two types of decision variables: 1) the firm energy value negotiated with the consumer and 2) the reservoir releases at a set of regularly spaced decision times throughout the contract period. The releases may be determined from an operating rule that is derived as part of the optimization process. The optimum firm energy value and operating rule maximize the operator's expected present value revenue in the presence of uncertain inflows, subject to relevant physical constraints.

This coupled stochastic optimization problem can be solved with an iterative algorithm that starts with an initial value for the firm energy value and an initial operating rule based on this value. The algorithm evaluates the resulting revenue, adjusts the contract energy to increase revenue, derives a new operating rule, and again evaluates the revenue, continuing until the process converges (Figure 1).

It is helpful to describe the contract in a mathematical form suitable for the optimization. Suppose for a given firm contract energy E_c that the reservoir generates actual energy E_k over the unit time interval $[t_k, t_{k+1}]$. The revenue obtained over this interval is determined by the piecewise linear concave revenue function illustrated in Figure 2:

$$
\begin{aligned}
g(E_k, E_c) &= \alpha_1(E_k - E_c) + \alpha_c E_c \quad \text{if} \quad E_k \leq E_c \\
g(E_k, E_c) &= \alpha_2(E_k - E_c) + \alpha_c E_c \quad \text{if} \quad E_k > E_c
\end{aligned}
\tag{1}
$$

Where E_c is the firm contract energy to be generated in each time interval during the contract period and $\alpha_1 > \alpha_c > \alpha_2$ are coefficients that define the price per unit energy in (\$ MWhr^{-1}) that applies for different situations. The term $\alpha_c E_c$ is the revenue obtained if the contract is exactly satisfied i.e. $E_k = E_c$. If E_k is greater than E_c, there is a surplus and the operator sells the additional energy $E_k - E_c$ at a lower rate $\alpha_2 < \alpha_c$. If E_k is less than E_c, there is a shortfall and the operator must purchase makeup energy $E_C - E_k$ at a higher rate $\alpha_1 > \alpha_c$. As mentioned earlier, we assume that the prices α_c, α_1 and α_2 are fixed, but that E_C is a decision variable. In contracts based on spot rather than fixed surplus and makeup power prices, variability in $\alpha_1, \alpha_c,$ and α_2 could be an important contributor to revenue uncertainty. This extension can be incorporated in the ensemble approach outlined here once the alternative contractural arrangement is precisely defined.

We suppose that the contract period extends from times t_1 to t_K and is divided into K-1 intervals of fixed duration Δt. The firm energy E_c needs to be known at the beginning of the contract period when the contract is negotiated. By contrast, the reservoir release u_k is most appropriately determined in real-time over each discrete time interval in the contract period, as illustrated in Figure 3. Real-time operation is important since it takes into account unanticipated variations in inflow and storage. To examine the real-time aspect in more detail we need to characterize the dynamic behavior of the reservoir system, which is described by the system states, releases, and energy output. These variables can be related by a set of stochastic constraints. For the single reservoir hydropower problem considered here the state vector x_k is partitioned into a scalar reservoir storage S_k, observed at time t_k, and a vector of states ψ_k that collectively describe an inflow time series model that could be estimated from observed inflow data using system identification techniques (Ljung, 2001). The associated state equations are

$$x_{k+1} = f(x_k, u_k, \omega_k)\ ;\ x_0 \text{ specified} \tag{2}$$

$$x_k = \begin{bmatrix} S_k \\ \psi_k \end{bmatrix}$$

$$S_{k+1} = f_S(S_k, \psi_k, u_k, \omega_k) \tag{3}$$

$$\psi_{k+1} = f_\psi(\psi_k, \omega_k) \tag{4}$$

$$I_k = M(\psi_k) \tag{5}$$

Here I_k is the total reservoir inflow over the time interval $[t_k, t_{k+1}]$, observed at t_k. This inflow is related to the time series model state ψ_k through a specified function $M(.)$. The scalar u_k (the control variable) is the total reservoir release over $[t_k, t_{k+1}]$ (specified by the operator at t_k) and ω_k is a sequence of independent random disturbances that drives the time series model. The time series model is used to predict inflows for the predictive operating rules considered in our example. Specific options for this model and its associated functions and variables are discussed in Section 4.

The storage state equation is a mass balance expression that neglects evaporation and seepage but includes spills:

$$S_{k+1} = f_S(S_k, \psi_k, u_k, \omega_k) = S_k + \Delta t[I_{k+1}(\psi_k, \omega_k) - u_k] - Z_k\ ;\ S_0 \text{ specified} \tag{6}$$

where the expressions in (4) and (5) can be used to write I_{k+1} in terms of the state vector ψ_k and disturbance ω_k. This state equation is used by the predictive operating rules to forecast storage from a particular predicted inflow sequence. The reservoir spill Z_k over $[t_k, t_{k+1}]$ is given by an additional constraint:

$$Z_k = \max\{S_k + \Delta t[I_{k+1} - u_k] - S_{max}, 0\} \tag{7}$$

Where S_{max} (L^3) is the reservoir capacity. The energy E_k generated by releases over $[t_k, t_{k+1}]$ is:

$$E_k = \phi(u_k, h_k, h_{k+1}) = u_k \int_{t=k}^{k+1} H(f_S)dt \qquad E_{max} = \phi(\bar{I}, h_{max}, h_{max}) \tag{8}$$

The reservoir head h_k at time t_k is related to the storage by a specified function $H(.)$ that depends on the reservoir geometry:

$$h_k = H(S_k) \qquad h_{max} = H(S_{max}) \tag{9}$$

The controlled release is constrained to be no greater than the turbine capacity u_{max}

$$u_k \leq u_{max} \tag{10}$$

For purposes of this study, the reservoir capacity, the head-storage function, and the turbine capacity are all assumed to be given. Note that fluxes are defined over K time intervals indexed by $k = 0{:}K-1$ and states are defined at K+1 discrete times indexed by $k = 0{:}K-1$. The time series notation can be made more concise if the entire sequence of releases defined through any time h_k is represented by the vector $u_{0,1,\ldots,k-1} = u_{0:k-1}$. Similar notation is used for sequences of other variables.

The desired solution to the real-time operations problem maximizes the following expected present value objective, which measures performance over the contract period $[t_0, t_K]$ for a given sequence of releases$u_{0:K-1}$, a given initial state x_0 , and a given firm energy value E_c :

$$J(u_{0:K-1}, x_0, E_c) = \mathcal{E}_{\omega_{0:K-1}}\left[\sum_{k=0}^{K-1}(1+r)^{-k}[g[E_k(u_k, x_k, x_{k+1}), E_c] - \alpha_Z Z_k(u_k, x_k)] + g_K(x_K)\right] \tag{11}$$

Dependence on the vector of random inflow disturbances $\omega_{0:K-1}$is removed by the expectation operation $\mathcal{E}_{\omega_{0:K-1}}$. The first term in the objective function expression is the present values of the hydropower revenue. The second term $\alpha_Z Z_k$ penalizes reservoir spills that can cause downstream flooding. The final term $g_K(x_K)$ (the salvage value) assigns a prescribed benefit to reservoir storage at the final time. This prevents the control strategy from emptying the reservoir at the end of the contract period. Specification of the spill and salvage value terms is discussed in more detail in Section 4.

The objective given in (11) could be maximized simultaneously with respect to the variables $u_{0:K-1}$ and E_c, using the methods of mathematical programming, imposing the constraints identified above. Since the contract must be determined before operations begin, at t_0, a simultaneous optimization of $u_{0:K-1}$ and E_c would require the entire release history to also be derived at t_0, before there are any observations of the actual states (open-loop control). Better revenue can generally be obtained if the contract is determined at the initial time but the releases are determined in real-time, as observations of the states become available (closed-loop control). This is possible if the release at each time is derived directly from the observed state, as specified by a closed loop operating rule (or decision function) of the following form:

$$u_k = \mu_k(x_k) \quad k = 0{:}K-1 \tag{12}$$

If (12) is substituted into (11) the objective J can be written as a functional $J_{\mu_{0:K-1}}(x_0, E_c)$ that maps the K decision functions $\mu_{0:K-1}$ to the scalar revenue.

$$J_{\mu_{0:K-1}}(x_0, E_c) = \mathcal{E}_{\omega_{0:K-1}} \left\{ \sum_{k=0}^{K-1} (1+r)^{-k} [g[E_k(x_0, \mu_{0:k}, \omega_{0:k}), E_c] - \alpha_Z Z_k(x_0, \mu_{0:k}, \omega_{0:k})] + g_K\big((x_0, \mu_{0:K-1}, \omega_{0:K-1})\big) \right\} \tag{13}$$

This is the real-time optimal control form of the optimization objective given in (11).

Note that the revenue, spill, and salvage terms are all random by virtue of their dependence on the random disturbance vector $\omega_{0:K-1}$. In our ensemble implementation of the stochastic optimal control problem, many random samples (or replicates) of this vector are drawn from a population determined by the statistics of the inflow time series model. Each inflow replicate gives a corresponding sample for each of the three terms in the objective and for the objective as a whole. The objective replicates provide equally likely predictions of the system performance for a given firm energy and decision strategy. The expected objective value is estimated by the arithmetic average of these replicates.

A real-time formulation of the operational part of the coupled optimization problem makes it possible to more precisely describe the iterative procedure outlined in Figure 1. If the current iterates (for iterations $l = 1, \ldots, L$) of the decision function and firm energy value are $\mu^l_{0:K-1}$and E^l_c the new decision strategy $\mu^{l+1}_{0:K-1}$is obtained by maximizing with respect to all decision functions that satisfy constraints (2)-(10).

$$\mu^{l+1}_{0:K-1} = \arg \max_{\mu_{o:K-1}} J_{\mu_{0:K-1}}(x_0, E^l_c) \tag{14}$$

This real-time optimal control sub-problem can be solved with the SDP, SMPC, and PI methods described in more detail in Section 3. Then the new firm energy value is obtained by maximizing $J_{\mu^{l+1}_{0:K-1}}(x_0, E_c)$ with respect to the scalar E_c:

$$E^{l+1}_c = \arg \max_{E_c} J_{\mu^{l+1}_{0:K-1}}(x_0, E_c) \tag{15}$$

This scalar optimization sub-problem can be readily solved with a one-dimensional search procedure (e.g. the Newton Raphson method). The iteration can be initialized with a plausible firm energy value, such as the energy that could be generated with a constant inflow somewhat less than the observed mean.

We denote the converged decision function and firm power by $\mu^*_{0:K-1}$ and E^*_c . We are unaware of a convergence proof for this algorithm but it has always converged in less than 20 iterations in the many sensitivity analyses we have performed for all of the predictive operating rules considered in Section 4. The iterates are well-constrained by the inflows and by the physical limitations of the reservoir system and all discontinuities (e.g. the reservoir spill expression) are approximated by locally smooth functions. Our experience has been that these factors lead to quick and reliable convergence.

The random inflow disturbance replicates generated in the iteration outlined above are used to guide the search procedure. In a practical application, the resulting optimum decision function and firm energy are used to determine the actual release from the reservoir. The corresponding actual inflow disturbance sequence will generally be different from any of the replicates used in the iteration. It is useful to quantify how well the reservoir system might work in such a situation. Since we do not know the actual inflows in advance such a performance assessment should account for uncertainty by considering a range of possible actual inflow disturbances. The framework for this assessment can be formulated in terms of the actual objective $J^a_{\mu^*_{0:K-1}}$,which depends on the actual inflow disturbance vector $\omega^a_{0:K-1}$ and the actual initial state as follows:

$$J^a_{\mu^*_{0:K-1}}(\omega^a_{0:K-1}, x^a_0, E^*_c) = \sum_{k=0}^{K-1}(1-r)^{-k}[g[E_k(x^a_0,\mu^*_{0:k},\omega^a_{0:k}),E^*_c] - \alpha_Z Z_k(x^a_0,\mu^*_{0:k},\omega^a_{0:k})] + g_K(x^a_0,\mu^*_{0:K-1},\omega^a_{0:K-1}) \tag{16}$$

Here E_c and the decision functions $\mu^*_{0:K-1}$ from (14) have been identified from the optimization procedure and can be considered given. At the initial time, before the inflows are observed, $\omega^a_{0:K-1}$ can be viewed as a random sequence sampled from the same population as the $W_{0:K-1}$ sequence that appears in (13). If x^a_0 is also unknown at the initial time it can also be treated as a random variable with a specified distribution. We call a collection of $\omega^a_{0:K-1}$ and x^a_0 samples a "meta-ensemble" to distinguish it from the ensemble $\omega_{0:K-1}$ used in the iterative search procedure.

If (16) is evaluated for a meta- ensemble of $\omega^a_{0:K-1}$ and x^a_0 samples we can derive the probability distribution of the actual present value revenue before inflows are actually observed. This distribution can be used to compute various revenue statistics such as the mean, upper quantile, etc. The process is carried out for selected decision rules in Section 4.

3 Options for deriving the real-time decision strategy

The options for deriving the operating rule $\boldsymbol{\mu_k(x_k)}$ use different methods to relate the current release to the current state. This section reviews some of the most promising alternatives.

3.1 Stochastic dynamic programming

Stochastic dynamic programming (SDP) provides a comprehensive approach for deriving real-time operating rules before real-time operations begin, without simplifying assumptions. In the discrete time version used here this method divides the real-time control problem of (12) and (13) into a sequence of *K* nested sub-problems that are solved with a recursion (Bellman, 1956). Each subproblem optimizes a time-dependent objective (the benefit-to-go) from a particular time to the end of the contract period. The objective for subproblem *k*, which is associated with time interval $[t_{k-1}, t_k]$ (commonly called Stage *k*) is:

$$J_{SDP,k}(x_k,E_c) = \max_{\mu_k(x_k)}\left[\mathcal{E}_{\omega_k}\left\{g[E_k(x_k,\mu_k(x_k),\omega_k),E_c] - \alpha_Z Z_k(x_k,\mu_k(x_k),\omega_k) + (1+r)^{-1}J_{SDP,k+1}(x_{k+1},E_c)\right\}\right]$$

(17)

The problems are nested because sub-problem k depends on the solution of sub-problem k+1. The solution is computed with a backward recursion that moves stage by stage from the final to initial contract times. A decision function $\mu_k(x_k)$ is derived and stored for sub-problem k (for $k = K-1, \ldots, 0$), for a given E_c. The recursion is initialized at $k = K$:

$$J_{SDP,K}(x_K, E_c) = g_K(x_K) \tag{18}$$

Note that the objective $J_{SDP,0}(x_0, E_c)$ obtained at the end of the recursion is equal to optimal revenue objective $J^*_{SDP,0}(x_0, E_c)$defined in (14). Also, the state equation can be used to express the term $J_{SDP,k+1}(x_{k+1}, E_c)$appearing in (17) as a functional that depends on $x_k, \mu_k(x_k), \omega_k$ and E_c .When the recursion is complete, the decision functions for all intervals are available and can be used to compute releases from actual observations in a forward real-time sweep (for $k = K-1, \ldots, 0$).

The maximization over $\mu_k(x_k)$ of the expected revenue in (17) gives the optimal release Stage k for any given value of the current state x_k. In practice, the state vector is usually discretized into a finite number of grid points and the optimum release value u^*_k is found at each of these points by maximizing the argument of (17), with E_c fixed. The releases at the grid points are interpolated to give a decision function $\mu_k(x_k)$ that applies at any feasible value of the state (Cervellera & Muselli, 2007; Johnson et al., 1993). The expectation operation appearing in (16) and (17) is approximated by the mean over an ensemble of synthetically generated ω_k samples, as discussed in Section 2.

Some distinctive aspects of the dynamic programming approach include: 1) the decision rules for all times are derived prior to the start of operations but each reservoir release is derived in real-time, after the current state is observed; 2) the decision function in our formulation depends on the energy contract; 3) the computational effort grows rapidly as the problem size increases. If N_{x_t},$N_{u_t}$and $N_{\omega_t}$are the number of discretized states, controls and inflow disturbances and the optimization horizon is K time steps, then the SDP algorithm requires $KN_{x_t}N_{u_t}N_{\omega_t}$ functional evaluations of the objective function; 4) performance is dependent on the accuracy of the predictive inflow and storage models (the stochastic state equations) 5) the algorithm implicitly accounts for the information provided by future measurements by relying on conditional probabilities that determine the likelihood of a transition from a particular observed state at t_k to another state at t_{k+1}. The computational demands of SDP tend to limit its application to problems with relatively small state vectors. In the hydropower operations context this implies that the problem needs to include only a few reservoirs and/or low dimensional inflow models.

3. 2 Stochastic model predictive control

Stochastic model predictive control (SMPC) derives the optimal release $\boldsymbol{u}^*_k$ at each decision time by maximizing expected revenue over a limited duration window extending into the future. The complete series of reservoir releases is computed by carrying out a new optimization at every time step rather than using a pre-computed decision rule. The objective for

Problem $\boldsymbol{k}$ originating at $\boldsymbol{t_k}$ is the present value revenue from $\boldsymbol{t_k}$ to $\boldsymbol{t_K}$, based on (13) and written directly in terms of releases rather than in terms of a decision function:

$$J_{SMPC,k}(u_{k:k+w-1}, x_k, E_c) = E_{\omega_{k:k+w-1}}\left\{\sum_{i=k}^{k+w-1}(1+r)^{-i}[g[E_i(x_k, u_{k:i}, \omega_{k:i}), E_c] - \alpha_Z Z_i(x_i, u_{k:i}, \omega_{k:i})] + g_{k+w}(x_k, u_{k:k+w-1}, \omega_{k:k+w-1})\right\} \quad (19)$$

The expectation operator is approximated by the mean over an ensemble of synthetically generated samples $\omega_{k:k+w-1}$. The optimization is carried out over a moving window of length $w \leq K-k$ time steps. This window spans the interval $[t_k, t_{k+w}]$.

An optimal release sequence over the current SMPC window is obtained by maximizing $J_{SMPC,k}(x_k, E_c)$ with respect to the releases:

$$u^*_{k:k+w-1} = \underset{u_{k:k+w-1}}{\arg\max}\, J_{SMPC,k}(u_{k:k+w-1}, x_k, E_c) \quad (20)$$

Although this optimization gives an entire sequence of optimal releases over the current time horizon, only the first release u_k is actually applied to the reservoir system (at t_k) since the remaining releases are recomputed at t_{k+1} when a new value of the state x_{k+1} is observed. This process is repeated for every decision time, until the moving window reaches the end of the contract period. The vector of current states $x_{0:K-1}$ and the associated vector of SMPC releases $u_{0:K-1}$ implicitly define a set of time-dependent decision functions $\mu_{0:K-1}$ through the relationship $u_k = \mu_k(x_k)$ for $k = 0: K-1$. For convenience, we refer to the SMPC decision as a function in the discussion below, even though SMPC does not explicitly derive such a function.

The distinctive aspects of model predictive control include: 1) releases are evaluated only for observed state values, not all possible values; 2) the decision function in our formulation depends on the contract energy 3) the decision function is defined implicitly and is available during operations only at the current time (not earlier), 4) future revenue is evaluated approximately, over a limited duration time horizon, 5) performance is dependent on the accuracy of the predictive inflow and storage models (the stochastic state equations) 6) SMPC is approximate, even in the limit as the time horizon becomes infinitely long, because it does not account for the impact of the future measurements, 7) computational effort is generally less than SDP, especially for large problems.

3.3 Standard operating policies

Both SDP and SMPC make an effort to predict the effect of future uncertain performance by averaging present value revenue over an ensemble of possible inflow disturbances. By contrast, deterministic (non-predictive) operating rules, such as the Standard Operating Policy (SOP) (Wurbs, 1993; Yeh, 1985) do not consider the possible impact of future inflows. These rules typically are heuristic and time-invariant (Figure 4). They do not optimize a particular objective and they are specified rather than derived functions of the system state. Non-predictive standard operating policies are easy to implement and convenient for multi-purpose reservoir operations but cannot generally be expected to perform as well in a single-purpose hydropower application as alternatives that utilize information about inflow variability and reservoir dynamics. They are considered here because they are widely used in practice and they provide benchmarks for assessing the potential performance improvement offered by predictive operating rules such as SDP and SMPC. Figure 4 shows two SOP variants. The simplest option, indicated by the black curve, releases all available water up to a nominal value equal to the mean inflow $u_{nom} = \bar{I}$ when the storage $S_{nom} - 0.1\Delta S_{max}$ .This nominal release is maintained until a nominal storage level $S_{nom} + 0.1\Delta S_{max}$ is reached. At that point additional water is released up to the maximum turbine capacity u_{max} . Beyond that, excess water must be spilled. The modified red curve hedges the release rule by smoothing abrupt transitions between low, nominal, and high storage conditions.

3. 4 Perfect information

Reservoir releases and revenues derived by assuming perfect knowledge of future inflows provide useful upper bounds on the performance that can be obtained for a particular actual inflow. In this case releases can be expressed in terms of a decision function but they need not be derived in real time. Instead, they can be computed by maximizing (11) with the assumption that the inflow disturbances $\boldsymbol{\omega_{0:i-1}} = \boldsymbol{\omega^a_{0:i-1}}$are not random but are known perfectly:

$$J_{PI}(u_{0:K-1}, \omega^a_{0:i-1}, x^a_0, E_c) = \sum_{i=0}^{K-1} (1+r)^{-i} \left[g[E_i(x^a_0, u_{0:i}, \omega^a_{0:i}), E_c] - \alpha_Z Z_i(x^a_0, u_{0:i}, \omega^a_{0:i}) \right] + g_K(x^a_0, u_{0:K-1}, \omega^a_{0:K-1}) \quad (21)$$

$$u^*_{0:K-1} = \underset{u_{0:K-1}}{\operatorname{argmax}} J_{PI}(u_{0:K-1}, \omega^a_{0:i-1}, x^a_0, E_c) \quad (22)$$

This problem can be solved with a standard non-linear programming algorithm since perfect information allows all releases to be computed at once, in batch rather than real-time mode. No reservoir operations method with imperfect information can do better than the perfect information case when presented with the same actual inflow.

4. Results and discussion

4.1 Setup of the example problem

The problem formulation and solution methods described above are tested here on a typical example using an ensemble of synthetically generated inflows. This Monte Carlo approach enables us to derive revenue probability distributions that quantify the risk associated with different contract selection/ real-time operations strategies. We suppose that the reservoir is designed primarily to generate hydropower, with operational objectives similar to those used in facilities such as Hoover Dam, USA; Tehri Dam, India; or and Itaipu Dam, Paraguay (Barros, et al., 2003; Fink, 2000) . Figure 5 shows the generic reservoir geometry and head-storage relation used in our example. The reservoir geometric information is provided in a tabular form that can be readily generalized to accommodate more complex head-storage functions.

The methods of this paper can be applied to any reservoir geometry as long as the storage vs. surface area and the head functions are provided. The standard operating policy used in the example is based on Figure 4 and uses a cubic function (red curve) to smooth transitions between the straight lines (black curve). The black lines are defined by the storage and release break points indicated in the figure. Note that these points depend on the value of the energy target E_c.

For the example we consider a single state random inflow model that gives sufficient variability to examine firm power shortages and surpluses as well as occasional spills. The normalized log of the inflow is a positive AR1 time series generated from a specified mean inflow, variance, and single lag correlation. The corresponding state equations are special cases of (3) and (4):

$$\begin{aligned} S_{k+1} &= f_S(S_k, \psi_k, u_k, \omega_k) \\ &= S_k + \Delta t[I_{k+1} - u_k] - Z_k \\ &= S_k + \Delta t\left[\bar{I}\exp\left(\rho_\psi \psi_k + \omega_k\right) - u_k\right] - Z_k; \quad S_0 \text{ specified} \\ \psi_{k+1} &= f_\psi(\psi_k, \omega_k) = \rho_\psi \psi_k + \omega_k \qquad \psi_0 \sim \mathcal{N}\left(\bar{\psi}, \sigma_\psi^2\right) \;\; \omega_k \sim \mathcal{N}(\bar{\omega}, \sigma_\omega^2) \end{aligned} \tag{23}$$

where ρ_ψ is the single lag correlation of $\psi_k$and the log normal inflow I_k is related to the unitless state ψ_k by:

$$I_k = M(\psi_k) = \bar{I}\exp(\psi_k) \tag{24}$$

The time-invariant mean and variance of ψ_k are computed from the specified mean and variance of I_k :

$$\bar{\psi} = -\frac{\sigma_\psi^2}{2}; \quad \sigma_\psi^2 = \ln\left(\frac{\sigma_I^2}{\bar{I}^2} + 1\right) \quad \bar{I}, \sigma_I^2 = \text{specified inflow mean and variance} \tag{25}$$

The time-invariant mean and variance of ω_k are obtained from:

$$\bar{\omega} = (1-\rho)\,\bar{\psi} \;\; \sigma_\omega^2 = (1-\rho^2)\sigma_\psi^2 \tag{26}$$

The AR1 model has the advantage, for testing purposes, of being having smaller correlation times than higher-order autoregressive models. The AR(1) model yields more variable inflows that are more difficult to predict than higher-model models. . Seasonality could be readily added if appropriate. In practice, the time series model should be estimated from historical inflow data and should be kept sufficiently low-dimensional to make an ensemble analysis of the predictive decision strategies computationally feasible.

A sensitivity analysis of the results can be conveniently formulated in terms of a limited of non-dimensional variables and inputs that are formed from groups of dimensional variables introduced above, using the definitions given in Tables 1 and 2. These non-dimensional quantities are identified by primed subscripts. Note that there is no spill penalty ($\alpha_Z = 0$) in the nominal case. Also, the maximum possible sustainable energy $E_{max} = \Phi(\bar{I}, h_{max}, h_{max})$ appearing in Table 1 is achieved when the reservoir head is fixed at its maximum value $h_{max} = H(S_{max})$ and the reservoir inflow and turbine release are both fixed at $\bar{I}$. The actual energy generated over a given time step could exceed this value if the release exceeds the mean inflow. For the example the dimensional problem objective function given in (12) and the dimensional constraints given in (2) through (9) are converted to non-dimensional forms by applying the definitions in Tables 1 and 2, as described in Appendix A. All plots and sensitivity analysis results are expressed in terms of non-dimensional variables.

Table 1: Non-dimensional variables

Non-dimensional variable	Definition	Range or distribution
Storage	$S_k' = \frac{S_k}{S_{max}}$	0.0 – 1.0
Head	$h_k' = \frac{h_k}{h_{max}}$	0.0 – 1.0
Inflow	$I'_k = \frac{I_k}{\bar{I}}$	Log normal
Log inflow	$\psi_k' = \log(I_k')$	Normal AR1 timeseries
Release	$u_k' = \frac{u_k}{\bar{I}}$	Non-negative
Spill	$Z_k' = \frac{Z_k}{S_{max}}$	Non-negative
Current revenue	$g_k' = \frac{g_k}{\alpha_c E_{max}}$	Non-negative
Energy	$E_k' = \frac{E_k}{E_{max}}$ $E_c' = \frac{E_c}{E_{max}}$	Non-negative
Objective (discounted revenue ratio)	$R = \frac{J'_{\mu_{0:K-1}}}{\alpha_c E_{max}}$	Non-negative

Table 2: Non-dimensional inputs

Non-dimensional input	Definition	Value in example
Reservoir residence time	$\tau_{res} = \frac{S_{max}}{\bar{I}\Delta t}$	Nominal: τ_{res}^{low}=12 ; τ_{res}^{high}= 48
Maximum reservoir release	$u'_{max} = \frac{u_{max}}{\bar{I}}$	1.5
Contract Period	$K' = \frac{K}{\Delta t}$	100

MPC window length,	$w' = \frac{w}{\Delta t}$	12
Spill penalty coefficient	$\alpha'_Z = \frac{\alpha_Z S_{max}}{\alpha_c E_{max}}$	Nominal: $\alpha_Z^{low} = 0$; $\alpha_Z^{high} = 20$
Revenue coefficients	$\alpha'_1 = \frac{\alpha_1}{\alpha_c}$; $\alpha'_2 = \frac{\alpha_2}{\alpha_c}$;	$\alpha'_1 = 2$; $\alpha'_2 = 0.15$
Log inflow AR1 statistical parameters	ρ_ψ, σ_ψ^2	$\rho_\psi = 0.8$, $\sigma_\psi^2 = 0.18$
Number of replicates	N	50
Number of meta-replicates	N^a	200
Discount factor	r	4%

The following subsections examine the results obtained by simulating the reservoir operation with four different coupled contract selection / real-time operations strategies based on Stochastic Dynamic Programming (SDP), Stochastic model Predictive Control (SMPC), a Standard Operating Policy (SOP) and a Perfect Information Scenario (PIS). They also consider the effect of varying influential dimensionless inputs such as the non-dimensional residence time, spill coefficient, and log inflow statistics.

4.2 Hydropower revenue comparison

The overall performance of the four decision strategies described in Section 3 can be assessed in terms of a number of performance measures, such as the net present value of the hydropower revenue generated over the contract period, revenue volatility over time, spill magnitude and frequency, etc. In our ensemble analysis many of these performance measures are random variables by virtue of their dependence on random inflows. To illustrate the capabilities of an ensemble approach we compare probability distributions for the net present value of the four decision strategies introduced earlier. Similar comparisons can be made of other performance measures. It is convenient to compare revenue performance in terms of the dimensionless revenue ratio $\boldsymbol{R}$ defined in Table 1. We first consider performance for the nominal input values given in Table 2 and then for a few alternatives that use different values for some of these inputs.

The perfect information strategy is unique among those considered here since it relies on advance knowledge of the entire sequence of reservoir inflows. With perfect inflow information, it is possible to derive a different optimum $\boldsymbol{E_c}$ for each meta-replicate in the Monte Carlo simulation. By contrast, each of the other strategies work with a single $\boldsymbol{E_c}$ value that maximizes expected revenue over the entire inflow ensemble for that particular strategy.

Figure 6 compares the kernel density estimates of probability distribution of the revenue ratio for all four decision strategies for nominal inputs. The variation in revenue observed for the perfect information (PIS) case depends only on the intrinsic variability of the actual inflow, not on the algorithm's ability to predict this inflow (since it has access to perfect inflow information). If the inflow for a particular actual inflow meta-replicate is low for a prolonged period, revenue will be low, even though the inflow is known perfectly. The other three decision rules are affected both by the intrinsic variability of the actual inflow and by uncertainty in the inflow predictions used to make release decisions. That is why their distributions are shifted to the left, toward lower revenue. The SDP and SMPC strategies tend to be more sharply peaked near their modes but have relatively long tails at lower revenue values, reflecting the consequences of occasional poor predictions. The most visible property of the PIS is its greater probability of yielding high revenue (R >0.75).

Stochastic dynamic programming (SDP) is second among the alternatives in terms of mean revenue since it makes best use of the ensemble inflow predictions when optimizing the current release. The backward recursion stores release strategies that maximize the expected revenue for the remaining contract time from any value of the state. These strategies can be recovered as the actual state values become known. By contrast, stochastic model predictive control (SMPC) derives a current release that maximizes expected revenue only from the current state. The replicates used in this calculation may not reflect the actual evolution of the system at later times. Also, the SMPC maximization is limited to a window that can be significantly shorter than the remaining contract time. For these reasons, SMPC is somewhat less likely to give high revenues and more likely to give low revenues than SDP [Lee, 2011]. The non-predictive standard operating policy performs the worst among the four alternatives, generating the smallest mean revenue with the highest probability of low revenues. This reflects the method's inability to adjust releases when near-future inflows and storages are likely to be lower or higher than average, given current inflow and storage. By contrast, predictive methods such as SDP and SMPC adjust releases in anticipation of possible future conditions. Table 3 lists the average revenue ratio computed over all the inflow meta-replicates as well as the probability (in %) of achieving a low revenue ratio below 0.5 or high ratio above 0.75. These percentages complement information on the mean revenue by considering the probability of low or high revenue values when comparing decision strategies.

4.3 Sample time series

The Monte Carlo simulation conducted in our example provides individual replicates of relevant dynamic variables such as the inflow, storage, release, and energy output as well as the revenue probability distributions discussed above. Figure 7 compares these variables for four different decision strategies, all using the nominal inputs from Table 2. Each of these four cases maximizes one of the decision strategy objectives specified in Section 3 by selecting the best possible combination of contract firm energy and release history for a given actual inflow meta-replicate. The normalized values of E_c for this example (expressed as a fraction of E_{max}) are 0.61 for perfect information (PIS), 0.57 for stochastic dynamic programming (SDP), 0.51 for stochastic model predictive control (SMPC), and 0.48 for the standard operating policy (SOP). Comparing to Figure 6, the predictive strategies that generate higher firm power are also more likely to produce higher revenue.

The top panel of Figure 7 shows the non-dimensional reservoir inflow series together with four turbine release series computed in real time from the current storage and inflow values, one for each of the four operating rules. The middle panel shows the non-dimensional reservoir storage generated from these releases, with the maximum normalized storage given by 1.0. The challenge for the operating rule is to keep water levels high in order to maximize energy output while avoiding spills that may have adverse downstream consequences and that also reduce the quantity of water available for generating power.

In the nominal case shown in Figure 7 the SDP decision strategy generally maintains higher storage than the other techniques, often approaching the reservoir capacity. This reflects SDP's somewhat better predictive capabilities and also the fact that spills are not explicitly penalized in the nominal case. SMPC behaves similarly but gives somewhat more erratic releases and energy production. Higher variability in energy together with a somewhat lower firm power value yield somewhat lower revenue for SMPC. The PIS is able to maintain much more stable release and energy production levels than any of the other methods. This reflects its ability to adjust releases in anticipation of future high or low inflow events, which are known perfectly. The advantage of perfect information also allows PIS to maintain a storage level that is generally lower than the other alternatives, even though the PIS average energy production and revenue are higher. The PIS result suggests the level of performance that SDP and SMPC could approach if they had access to very accurate inflow estimates.

4.4 Sensitivity analysis

All of the non-dimensional parameters listed in Table 2 effect the performance of the four different operational strategies considered here It is useful to examine in detail two key dimensionless inputs, the normalized spill penalty coefficient $\boldsymbol{\alpha'_z}$ and the residence time $\boldsymbol{\tau_{res}}$, and to briefly consider some of the others.

4.4.2 Sensitivity to spill penalty

A higher spill penalty tends to make the operational strategy more conservative, lowering the water level below the maximum to reduce the magnitude and frequency of spills. Table 3 includes a comparison of expected revenue and the probability of low and high revenues for a moderately high non-dimensional spill penalty value vs. the nominal case that does not penalize spills. Figure 8 shows revenue ratio probability distributions for the same two spill penalty options. Increasing the spill penalty consistently shifts the revenue probability density towards lower values (Figure 8). As the penalty coefficient increases spill occurrences decrease from 15% to 2.7% for dynamic programming, from 7% to 2.1% for model predictive control, and from 6% to 2.8% for the standard operating policy. Dynamic programming has the highest spill occurrence for the unpenalized case because its more complete description of uncertain future conditions benefits more from pushing the reservoir system to capacity in order to achieve maximum performance. Its more complete treatment of uncertainty also enables dynamic programming to significantly reduce spill occurrence when spills are penalized. By contrast, SOP gives a lower unpenalized spill occurrence but does not achieve as great a reduction when spills are penalized. Model predictive control falls somewhere in between.

The perfect information option shows a similar sensitivity to the spill penalty but gives a lower unpenalized spill occurrence than any of the alternatives. Perfect information makes the most difference during high inflow events that can cause spills since it enables the operating rule to draw down the reservoir before high flows occur. By reducing the amount of water lost to spills the perfect information option is able to generate more hydropower and greater revenue. It is possible to decrease spill occurrence somewhat further than indicated in Figure 8, by further increasing the spill penalty. But this effect is ultimately limited by the inflow statistics. Overall, perfect information and dynamic programming sacrifice revenue less than the other alternatives when spills are penalized.

Table 3: Comparison of the average revenue ratio R and probability of low R (< 0.5) and a high R (>0.75) between the four operational strategies

Technique	Low Spill Penalty, $\alpha_Z^{low} = 0$ (Nominal)			High Spill penalty, $\alpha_Z^{high} = 20$		
	$\mathcal{E}(R)$	P(R<.5)	P(R>.75)	$\mathcal{E}(R)$	P(R<.5)	P(R>.75)
Perfect information	0.69	3%	25%	0.62	18%	17%
Dynamic programming	0.64	5%	8%	0.62	10%	4%
Model predictive control	0.62	11%	6%	0.58	26%	2%

Standard operating policy	0.59	20%	5%	0.51	40%	0%

4.4.2 Sensitivity to residence time

The residence time ($\boldsymbol{\tau_{res}}$) provides a concise description of the combined effect of the reservoir capacity and the mean inflow. Increasing the residence time (low inflows/ large reservoirs) reduces sensitivity to inflow variability, generating higher revenue for extended periods. Figure 9 shows this behavior by plotting the revenue ratio probability distributions for two different residence time options: low $\boldsymbol{\tau_{res}^{low}}$ (nominal) vs. high $\boldsymbol{\tau_{res}^{high}}$. Increasing the residence time reduces the effects of inflow variability and shifts the revenue distributions towards higher values consistently across all four techniques. The revenue distribution also narrows, reducing the risk of lower revenues. Increasing the residence time increases the average revenue by 18% for dynamic programming, 20% for model predictive control, and 18% for the standard operating policy.

With a high residence time, high inflow events do not necessarily cause uncontrolled spills. They can be captured as storage, making it possible to temporarily allow releases greater than the mean inflow ($\boldsymbol{\bar{I}}$). This can yield revenue ratios (R) greater than 1 (see, for example, the perfect information case in Figure 9). For the nominal spill penalty coefficient spill occurrences decrease significantly with increasing residence time: 4.4% to 0.15% for perfect information, 15% to 3.8% for dynamic programming, 7% to 1.9% for model predictive control and 6% to 0.3% for standard operating policy. Although a large residence time reservoir is clearly desirable the potential for increased capacity is practically limited by site constraints and higher costs. When designing a new reservoir such considerations need to be included in the optimization process.

Table 4: Comparison of the average revenue ratio R and probability of low R (< 0.5) and a high R (>0.75) between the four operational strategies

Technique	Low residence time, $\tau_{res}^{low} = 12$ (Nominal)			High residence time , $\tau_{res}^{high} = 48$		
	$\mathcal{E}(R)$	P(R<.5)	P(R>.75)	$\mathcal{E}(R)$	P(R<.5)	P(R>.75)
Perfect information	0.69	3%	25%	0.82	0%	75%
Dynamic programming	0.64	5%	8%	0.78	1%	79%
Model predictive control	0.62	11%	6%	0.77	3%	78%
Standard operating policy	0.59	20%	5%	0.70	6%	38%

4.4.4 Sensitivity to other factors

The preceding sections show that the comparative performance between the three real-time operational strategies (SDP, SMPC and SOP) is sensitive to spill penalty and residence time. Performance also depends on other parameters such as reservoir geometry, inflow statistics, discount rate, and revenue function coefficients. For example, predictive operating strategies such as SDP and SMPC provide a greater performance benefit if the reservoir inflow series has a high serial correlation ρ_ψ and a low or moderate variance σ_ψ^2. In such cases it is easier to predict near-term inflows. On the other hand, if the correlation is close to 1 and the variance is high the possibility of extended periods of anomalous inflows leads to reduced benefit even for predictive algorithms. The revenue function parameters can also influence performance through their impact on both the contract value and real-time operations. For example, increasing α_1 increases the penalty of generating a shortfall, which leads to a more conservative contract that keeps reservoir storage near capacity and increases spill occurrence. The combined effect of many sensitivities determines the relative effectiveness of predictive vs. deterministic operating rules in any given situation.

5 Conclusions and Discussion

5.1 Summary of results

This paper describes a novel stochastic optimization approach that simultaneously selects a firm power target and a real-time release strategy for a hydropower reservoir. The probability distribution of operator revenue depends significantly on both of these design elements. Predictive techniques such as stochastic dynamic programming (SDP) and stochastic model predictive control (SMPC) give significantly better revenue (mean improvement > 10%) than a non-predictive standard operating policy (SOP) for the nominal conditions considered here. For other conditions the improvement may be either greater or less. Predictive techniques tend to work best in situations where reservoir inflow statistics favor the use of inflow and storage forecasts for optimizing revenue. Between the two predictive techniques, SDP generates higher revenue than SMPC but can be more computationally demanding, especially for multi-reservoir systems.

Sensitivity analysis indicates that a high spill penalty has a negative impact on revenue since it leads to strategies that operate the reservoir at a lower storage level. Reservoirs with a higher residence time generate higher revenues and result in less spill since the sensitivity to inflow variability decreases.

5.2 Generalization and extensions

The analysis described here makes certain simplifications that could be modified and generalized if appropriate. The emphasis is on single purpose hydropower reservoirs operated with the objective of maximizing revenue, subject to a penalty for excessive spills. Additional objectives could be incorporated; either through new terms in the objective function or through chance constraints that require specified measures of, for example, recreational, irrigation, or flood control benefits, to be exceeded with a certain probability. Tradeoffs among objectives could be examined with multi-objective visualization tools such as those described in (Woodruff, et al. , 2013). However, as more objectives are added and hydropower revenue is given lower priority the ability to optimize a firm power target and release strategy becomes more constrained. For this reason, the methods described here are most relevant for reservoirs that are primarily intended to generate hydropower.

The long term bilateral fixed price power purchase agreement used in our analysis insulates both operator and buyer from energy price fluctuations and is most appropriate when the buyer's demand is well defined and predictable. Such agreements are becoming more popular for bilateral corporate renewable energy transactions (Baker & A.McKenzie, 2015). However, it should be recognized that fixed price agreements may not be desirable or practical in all situations, especially where demand is uncertain and energy price fluctuations could have a significant effect on operator or buyer revenue. It would be reasonably straightforward to replace the fixed price agreement with alternatives with contract terms that vary with market prices. In such cases it is likely that the number of contract decision variables would increase beyond the single energy target value considered here.

The ability of a particular release strategy to track a particular energy target depends significantly on the nature of reservoir inflow variability as well as the reservoir's physical properties. The example considered in this paper uses a log inflow that is a normally distributed AR(1) autoregressive time series with a specified mean, variance, and correlation time. This choice gives reasonable variability and persistence and provides the basis for the ensemble predictions used in the SDP and SMPC release strategies. In any given application, the actual inflows may vary in other ways that should be determined, as much as possible, from historical data. If the inflow model were changed the relative performance of the different decision strategies could also change. Both of the predictive release strategies, SDP and SMPC, are able to accommodate inflow models other than the AR(1) by increasing the dimensionality of the inflow state vector ψ_k in the problem formulation. It is important to note that the sensitivity of the results to the inflow model is mitigated somewhat by the real-time nature of the release decision rules. One of the primary goals of real-time control is to provide a mechanism that can use observations to compensate for model approximations and simplifications. The inflow model need not be perfect for the control strategy to improve performance over alternative methods.

The conceptual framework presented here provides a probabilistic perspective that quantifies both revenue and spill risk for a hydropower reservoir designed to meet a firm power target. This framework can be adapted to accommodate different reservoir shapes, inflow models, revenue functions, and contact structures. It can also be extended to multi-reservoir systems. The example considered here indicates that a stochastic approach that focuses on the probability distributions of inflow and revenue can provide useful insights and tangible benefits for both hydropower reservoir operations and contract negotiations.

Acknowledgments

The authors would like to thank Prof. Kenneth Strzepek and Prof. Saurabh Amin at MIT as well as three anonymous reviewers for their comments and help. Partial funding for this research was provided by the Abdul Latif Jameel World Water and Food Security Lab (J-WAFS) and the Environmental Solutions Initiative at MIT. Data associated with this study can be downloaded from [https://github.mit.edu/reetik/WRR_Hydropower.git].

Figure captions

Figure 1: Iterative search for optimum energy contract. The energy contract E_c proposed at each iteration requires a new decision rule (turbine release vs. storage) to maximize revenue.

Figure 2: Piecewise linear concave revenue function. The slope of the red-dotted and black lines shows how the unit revenue at the contract energy compares to the unit cost of makeup power at lower energy values and the unit revenue of surplus energy at higher energy values

Figure 3: Example representation of discrete reservoir variables defined over two consecutive time intervals. The bottom panel shows the piecewise linear storage state (S) over each interval. The top panel shows the piecewise constant turbine release (u) and inflow (I) over each interval, with the inflow measurement (I) observed at the end of the interval.

Figure 4: Schematic representation of two typical Standard Operating Policies, with the reservoir release expressed as a function of currently available storage. Deviations of the red (hedged) curve from the black (standard) curve indicate an effort to moderate abrupt transitions between low, nominal, and high storage conditions.

Figure 5: Reservoir geometry for the example problem. Left panel shows reservoir configuration and right panel plots the storage vs. head curve for the example

Figure 6: Probability density function of the revenue ratio for SDP, SMPC and PIS operational techniques

Figure 7: Example reservoir operations with the four techniques (SDP, SMPC, SOP and PIS) plotted for a particular inflow meta-replicate. Top panel: Reservoir inflow time series and turbine release; Middle-panel: Reservoir storage; Bottom panel: Energy generated. All quantities are non-dimensional.

Figure 8: Effect of spill penalty on the revenue density function. Mean revenue and spill frequency both decrease as the spill penalty is increased from nominal $\alpha_z^{low} = 0$ *to* $\alpha_z^{high} = 20$

Figure 9: Effect of residence time on the revenue density function. Increase in residence time τ_{res} *(nominal* $\tau_{res}^{low} = 12$, $\tau_{res}^{high} = 48$) *shifts the revenue distributions to higher revenue in every operational strategy*

Appendix A: Non-dimensional Problem Formulation

The following expressions give non-dimensional versions of the coupled contract-operational design problem objective and constraints. The non-dimensionalization is illustrated for the definitions given in Tables 1 and 2 and uses the AR1 log inflow model described in Section 4.1.

$$J'_{\mu'_{0:K-1}} = \mathcal{E}\left\{\sum_{k=0}^{K-1}(1+r)^{-k}[g'[E'_k, E'_c] - \alpha'_c Z'_k] + g'_K\right\} \quad \textit{objective} \tag{A-1}$$

$$g'(E'_k, E'_c) = \alpha_1(E'_k - E'_c) + E'_c \quad \text{if } E'_k \leq E'_c \tag{A-2}$$

$$g'(E'_k, E'_c) = \alpha_2(E'_k - E'_c) + E'_c \quad \text{if } E'_k > E'_c$$

$$\psi_k = \log(I'_k) \tag{A-3}$$

$$\psi_{k+1} = f_2(x_k, u_k, \omega_k) = \rho_\psi \psi_k + \omega_k \quad \textit{log inflow equation} \tag{A-4}$$

$$\begin{aligned} S'_{k+1} &= S'_k + \frac{1}{\tau_{res}}[I'_{k+1} - u'_k] - Z'_k \\ &= S'_k + \frac{1}{\tau_{res}}[\exp(\rho\psi_k + \omega_k) - u'_k] - Z'_k \end{aligned} \quad \text{Storage equation} \tag{A-5}$$

$$Z'_k = \max\left\{S'_k + \frac{1}{\tau}[I'_{k+1} - u'_k] - 1, 0\right\} \quad \textbf{spill equation} \tag{A-6}$$

$$E'_k = \phi'(u'_k, h'_k, h'_{k+1}) = \frac{1}{E_{max}}\phi(u_{max}u'_k, h_{max}h'_k, h_{max}h'_{k+1}) \quad \textbf{energy equation}$$
(A-7)

$$h'_k = H'(S'_k) = \frac{1}{h_{max}}H(S_{max}S'_k); \quad k = 0{:}K. \ \textit{head} - \textit{storage} \tag{A-8}$$

$$u'_k \leq u'_{max} \quad \textit{release upper bound} \tag{A-9}$$

Bibliography

Arnold, M., & Andersson, G. (2011). Model predictive control of energy storage including uncertain forecasts. *Power Systems Computation Conference (PSCC), Stockholm, Sweden*.

Baker, C. G. J., & A.McKenzie, K. (2015). The Rise of Corporate PPAs: A New Driver for Renewables. Retrieved September 5, 2018, from https://www.bakermckenzie.com/-/media/files/insight/publications/2015/12/the-rise-of-corporate-ppas/risecorporateppas.pdf?la=en

Bar-Shalom, Y., & Tse, E. (1974). Dual Effect, Certainty Equivalence, and Separation in Stochastic Control. *IEEE Transactions on Automatic Control*, *19*(5), 494–500. https://doi.org/10.1109/TAC.1974.1100635

Barjas Blanco, T., Willems, P., Chiang, P. K., Haverbeke, N., Berlamont, J., & De Moor, B.

(2010). Flood regulation using nonlinear model predictive control. *Control Engineering Practice*, *18*(10), 1147–1157. https://doi.org/10.1016/j.conengprac.2010.06.005
Barros, M. T. L., Tsai, F. T.-C., Yang, S., Lopes, J. E. G., & Yeh, W. W.-G. (2003). Optimization of Large-Scale Hydropower System Operations. *Journal of Water Resources Planning and Management*, *129*(June), 11. https://doi.org/10.1061/(ASCE)0733-9496(2003)129:3(178)
Barroso, L. a., Rosenblatt, J., Guimaraes, A., Bezerra, B., & Pereira, M. V. (2006). Auctions of contracts and energy call options to ensure supply adequacy in the second stage of the Brazilian power sector reform. In *2006 IEEE Power Engineering Society General Meeting* (p. 8 pp.). https://doi.org/10.1109/PES.2006.1708974
Bellman, R. (1956). Dynamic Programming and Lagrange Multipliers. *Proceedings of the National Academy of Sciences*, *42*(10), 767–769. https://doi.org/10.1073/pnas.42.10.767
Bertsekas, D. (1995). *Dynamic programming and optimal control*. *Control* (Vol. II). Retrieved from http://live.iugaza.edu.ps/NR/rdonlyres/Electrical-Engineering-and-Computer-Science/6-231Dynamic-Programming-and-Stochastic-ControlFall2002/5CFA7384-8ED4-4F99-99A2-B83B3B2290EF/0/DP_2ndEDITION_Corrections.pdf
Bertsekas, D. P., & Castañon, D. A. (1999). Rollout algorithms for stochastic scheduling problems. *Journal of Heuristics*, *5*(1), 89–108. https://doi.org/10.1023/A:1009634810396
Boneville Power Administration. (2013). *Marketing Hydropower: Fact Sheet*. Retrieved from https://www.bpa.gov/news/pubs/FactSheets/fs-201304-Marketing-Hydropower.pdf
Castelletti, A., de Rigo, D., Rizzoli, A. E., Soncini-Sessa, R., & Weber, E. (2007). Neuro-dynamic programming for designing water reservoir network management policies. *Control Engineering Practice*, *15*(8), 1031–1038. https://doi.org/10.1016/j.conengprac.2006.02.011
Catalão, J. P. S., Pousinho, H. M. I., & Contreras, J. (2012). Optimal hydro scheduling and offering strategies considering price uncertainty and risk management. *Energy*, *37*(1), 237–244. https://doi.org/10.1016/j.energy.2011.11.041
Cervellera, C., Chen, V. C. P., & Wen, A. (2006). Optimization of a large-scale water reservoir network by stochastic dynamic programming with efficient state space discretization. *European Journal of Operational Research*, *171*(3), 1139–1151. https://doi.org/10.1016/j.ejor.2005.01.022
Cervellera, C., & Muselli, M. (2007). Efficient sampling in approximate dynamic programming algorithms. *Computational Optimization and Applications*, *38*(3), 417–443. https://doi.org/10.1007/s10589-007-9054-8
Council, W. E. (2016). *World Energy Resources: Hydropower 2016*. *Energy Conservation*. Retrieved from https://www.worldenergy.org/wp-content/uploads/2017/03/WEResources_Hydropower_2016.pdf
Energy, A. C. on R. (2016). *Corporate renewable energy procurement survey insights*. Retrieved from https://acore.org/wp-content/uploads/2017/12/Corporate-Renewable-Energy-Procurement-Industry-Insights.pdf
Fink, A. K. (2000). Tehri hydro power complex on the Bhagirathi river in India. *Hydrotechnical Construction*, *34*(8–9), 479–484. https://doi.org/10.1023/A:1004187208788
Foster, B. T., Kern, J. D., & Characklis, G. W. (2015). Mitigating hydrologic financial risk in hydropower generation using index-based financial instruments. *Water Resources and Economics*, *10*, 45–67. https://doi.org/10.1016/j.wre.2015.04.001
García, C. E., Prett, D. M., & Morari, M. (1989). Model predictive control: Theory and practice-A survey. *Automatica*, *25*(3), 335–348. https://doi.org/10.1016/0005-1098(89)90002-2

Hall, W. A., Butcher, W. S., & Esogbue, A. (1968). Optimization of the Operation of a Multiple-Purpose Reservoir by Dynamic Programming. *Water Resources Research*, *4*(3), 471–477. https://doi.org/10.1029/WR004i003p00471

Hooper, E. R., Georgakakos, A. P., & Lettenniaier, D. P. (1991). Optimal Stochastic Operation of Salt River Project, Arizona. *Journal of Water Resources Planning and Management*, *117*(5), 566–587. https://doi.org/10.1061/(ASCE)0733-9496(1991)117:5(566)

Johnson, S. A., Stedinger, J. R., Shoemaker, C. A., Li, Y., & Tejada-Guibert, J. A. (1993). Numerical Solution of Continuous-State Dynamic Programs Using Linear and Spline Interpolation. *Operations Research*, *41*(3), 484–500. https://doi.org/10.1287/opre.41.3.484

Labadie, J. W. (2004). Optimal Operation of Multireservoir Systems: State-of-the-Art Review. *Journal of Water Resources Planning and Management*, *130*(2), 93–111. https://doi.org/10.1061/(ASCE)0733-9496(2004)130:2(93)

Linke, H. (2010). A model-predictive controller for optimal hydro-power utilization of river reservoirs. In *Proceedings of the IEEE International Conference on Control Applications* (pp. 1868–1873). IEEE. https://doi.org/10.1109/CCA.2010.5611081

Ljung, L. (2001). Estimating Linear Time-invariant Models of Nonlinear Time-varying Systems. *European Journal of Control*, *7*(2–3), 203–219. https://doi.org/10.3166/ejc.7.203-219

Mayne, D. Q., Rawlings, J. B., Rao, C. V., & Scokaert, P. O. M. (2000). Constrained model predictive control: Stability and optimality. *Automatica*, *36*(6), 789–814. https://doi.org/10.1016/S0005-1098(99)00214-9

Mo, B., Gjelsvik, A., & Grundt, A. (2001). Integrated risk management of hydro power scheduling and contract management. *IEEE Transactions on Power Systems*, *16*(2), 216–221. https://doi.org/10.1109/59.918289

Neelakantan, T. R., & Pundarikanthan, N. V. (1999). Hedging rule optimisation for water supply reservoirs system. *Water Resources Management*, *13*(6), 409–426. https://doi.org/10.1023/A:1008157316584

Negenborn, R. R., Overloop, P. J. Van, Keviczky, T., & De Schutter, B. (2009). Distributed model predictive control of irrigation canals. *Network and Heterogeneous Media*, *4*(2), 359–380. https://doi.org/10.3934/nhm.2009.4.359

Perea-López, E., Ydstie, B. E., & Grossmann, I. E. (2003). A model predictive control strategy for supply chain optimization. In *Computers and Chemical Engineering* (Vol. 27, pp. 1201–1218). https://doi.org/10.1016/S0098-1354(03)00047-4

Qin, S. ., & Badgwell, T. . (2003). A survey of industrial model predictive control technology. *Control Engineering Practice 11*, 733–764. Retrieved from http://www.sciencedirect.com/science/article/pii/S0967066102001867

Rani, D., & Moreira, M. M. (2010). Simulation-optimization modeling: A survey and potential application in reservoir systems operation. *Water Resources Management*, *24*(6), 1107–1138. https://doi.org/10.1007/s11269-009-9488-0

Rawlings, J. B. (2000). Tutorial Overview of Model Predictive Control. *IEEE Control Systems*, *20*(3), 38–52. https://doi.org/10.1109/37.845037

Shrestha, G. B., Pokharel, B. K., Lie, T. T., & Fleten, S. E. (2005). Medium term power planning with bilateral contracts. *IEEE Transactions on Power Systems*, *20*(2), 627–633. https://doi.org/10.1109/TPWRS.2005.846239

Stedinger, J. R., Faber, B. A., & Lamontagne, J. R. (2013). Developments in Stochastic Dynamic Programming for Reservoir Operation Optimization. *World Environmental and Water Resources Congress 2013*, 1266–1278. https://doi.org/10.1061/9780784412947.125

Stickler, C. M., Coe, M. T., Costa, M. H., Nepstad, D. C., McGrath, D. G., Dias, L. C. P., … Soares-Filho, B. S. (2013). Dependence of hydropower energy generation on forests in the Amazon Basin at local and regional scales. *Proceedings of the National Academy of Sciences*, *110*(23), 9601–9606. https://doi.org/10.1073/pnas.1215331110

Tu, M.-Y., Hsu, N.-S., & Yeh, W. W.-G. (2003). Optimization of Reservoir Management and Operation with Hedging Rules. *Journal of Water Resources Planning and Management*, *129*(2), 86–97. https://doi.org/10.1061/(ASCE)0733-9496(2003)129:2(86)

van Overloop, P. J., Weijs, S., & Dijkstra, S. (2008). Multiple Model Predictive Control on a drainage canal system. *Control Engineering Practice*, *16*(5), 531–540. https://doi.org/10.1016/j.conengprac.2007.06.002

Woodruff, M. J., Reed, P. M., & Simpson, T. W. (2013). Many objective visual analytics: Rethinking the design of complex engineered systems. *Structural and Multidisciplinary Optimization*, *48*(1), 201–219. https://doi.org/10.1007/s00158-013-0891-z

Wurbs, R. A. (1993). Reservoir-System Simulation and Optimization Models. *Journal of Water Resources Planning and Management*, *119*(4), 455–472. https://doi.org/10.1061/(ASCE)0733-9496(1993)119:4(455)

Yakowitz, S. (1983). Dynamic programming applications in water resources. *Water Resources Research*, *18*(4), 673–696. Retrieved from http://onlinelibrary.wiley.com/doi/10.1029/WR018i004p00673/full

Yeh, W. W. (1985). Reservoir Management and Operations Models: A State-of-the-Art Review. *Water Resources Research*, *21*(12), 1797–1818. https://doi.org/10.1029/WR021i012p01797

You, J. Y., & Cai, X. (2008). Hedging rule for reservoir operations: 2. A numerical model. *Water Resources Research*, *44*(1), n/a-n/a. https://doi.org/10.1029/2006WR005482

Zhao, T., Zhao, J., & Yang, D. (2014). Improved Dynamic Programming for Hydropower Reservoir Operation. *Journal of Water Resources Planning and Management*, *140*(3), 365–374. https://doi.org/10.1061/(ASCE)WR.1943-5452.0000343