STOCHASTIC OPTIMIZATION OF A HYDROELECTRIC RESERVOIR USING PIECEWISE POLYNOMIAL APPROXIMATIONS

BERNARD F. LAMOND

Département Opérations et systèmes de décision,
Faculté des sciences de l'administration,
Université Laval,
Québec (Qc) Canada G1K 7P4
e-mail: Bernard.Lamond@fsa.ulaval.ca

ABSTRACT

We propose a method for optimizing a single hydro-electric reservoir using a piecewise polynomial approximation of the future value functions. Unlike previous methods based on splines, we avoid discretizing the inflow distribution. Instead, we carry out the expectation step of dynamic programming using an exact, easy-to-evaluate formula for the integral of a piecewise polynomial function. We then apply our method to solving a model which assumes a piecewise linear reward function of the energy produced, and takes into account the turbine head effects.

Keywords: Markov decision processes, stochastic dynamic programming, piecewise polynomial approximation, optimal reservoir management, energy planning.

RÉSUMÉ

Nous proposons une méthode pour optimiser un réservoir hydro-électrique unique en utilisant une approximation polynomiale par morceau des fonctions de valeur future. Contrairement aux méthodes précédentes basées sur les splines, nous évitons la discrétisation de la distribution des apports. À la place, nous effectuons l'étape du calcul de l'espérance en programmation dynamique en utilisant une formule exacte et facile à évaluer pour l'intégrale de la fonction polynomiale par morceau. Nous appliquons ensuite notre méthode à la résolution d'un modèle qui assume une fonction de récompense linéaire par morceau de l'énergie produite, et tenant compte de l'effet de hauteur de chute des turbines.

Mots-clés : Processus de décision markoviens, programmation dynamique stochastique, approximation polynomiale par morceau, gestion optimale des réservoirs, planification énergétique.

1. INTRODUCTION

We consider a discrete time, finite horizon, stochastic optimization model for a hydro-electric system with a single reservoir. We assume the rewards from electricity production are given by a concave, piecewise linear function of the energy produced. The energy produced may depend on the storage as well as the discharge, thus allowing turbine head effects to be taken into account. See Lamond and Boukhtouta (1996), Lamond and Boukhtouta (2001) and the references therein for surveys of models and methods for stochastic reservoir optimization.

As in the usual stochastic dynamic programming (DP) approach (see, e.g., Puterman (1994)), we define the following random variables for $t = 1, \ldots, T$

$$S_t = \text{state at beginning of period } t$$
$$A_t = \text{action taken in period } t$$
$$R_t = \text{immediate reward received in period } t.$$

Let also S_{T+1} be the terminal state. Realizations of the state and action variables are denoted s_t and a_t, respectively, with state space $\mathcal{S}$ and conditional action set $\mathcal{A}_t(s_t)$ when $S_t = s_t \in \mathcal{S}$. We define the immediate reward functions

Recd. Oct. 2001, Rev. Oct. 2002, Acc. Feb. 2003

$$r_t(s_t, a_t) = E[R_t \mid S_t = s_t, A_t = a_t]$$

for $t = 1, \ldots, T$ and the terminal reward function $V_{T+1}(s_{T+1})$. Then the functions $V_t(s_t)$ of optimal, expected cumulative rewards from period t onward are given recursively by Bellman's principle of optimality, for $t = T, T-1, \ldots, 1$:

$$V_t(s_t) = \max_{a_t \in \mathcal{A}_t(s_t)} r_t(s_t, a_t) + \beta W_{t+1}(s_t, a_t) \tag{1}$$

where β is a discount factor and

$$W_{t+1}(s_t, a_t) = E[V_{t+1}(S_{t+1}) \mid S_t = s_t, A_t = a_t]. \tag{2}$$

A simple solution method is to discretize the state and action variables s_t and a_t, and to assume discrete transition probabilities, so that equations (1, 2) can be solved using discrete dynamic programming. A very fine discretization is usually required in order to obtain a satisfactory accuracy, however, resulting in a large amount of computation (Kitanidis and Foufoula-Georgiou (1987)). This approach is, in fact, impractical for solving multidimensional models of systems with many reservoirs. Nonetheless, some authors have proposed aggregation methods for multireservoir systems that solve a sequence of problems with two reservoirs (Turgeon (1980), Turgeon (1981)) or three reservoirs (Archibald et al. (1997)). In this context, the small models ought to be solved quickly yet accurately, because they have to be solved repetitively. Here, we propose an approach for solving a one dimensional problem which is faster and more accurate than discrete DP.

One such approach is to use continuous state and action variables and to approximate the functions $V_t(s_t)$ using, for example, polynomial or piecewise polynomial functions (Chen et al. (1999), Foufoula-Georgiou and Kitanidis (1988), Johnson et al. (1993), Philbrick and Kitanidis (2001)). When the function is suitably smooth, spline approximations are used to preserve continuity of the first derivative at the grid points. Then eq. (1) is solved by a continuous, nonlinear programming method, and the expectation in eq. (2) is evaluated using a Gaussian quadrature rule for numerical integration, which corresponds to a coarse discretization of the distribution of the random variables (for example, expectations are computed using only two or three quadrature points in Chen et al. (1999), Foufoula-Georgiou and Kitanidis (1988), Johnson et al. (1993), and Philbrick and Kitanidis (2001)).

Here we also use continuous state and action variables and a piecewise polynomial approximation for $V_t(s_t)$. However, due to the piecewise linear revenues, our function $V_t(s_t)$ is not as smooth as in Chen et al. (1999), Foufoula-Georgiou and Kitanidis (1988), Johnson et al. (1993), and Philbrick and Kitanidis (2001), hence their (coarse) Gaussian quadrature scheme is not suitable for computing the expectation in our case. Therefore, we derive an exact formula for efficiently evaluating the expectation function $W_{t+1}(s_t, a_t)$ of eq. (2) using a continuous distribution instead of a coarse discretization. Our approach extends the work of Drouin et al. (1996) and Lamond and Lang (1996), where a method for efficiently computing the expectation of a piecewise linear function is given for a discrete distribution over a fine grid. Here, we extend this result in two ways: we allow (*i*) piecewise polynomial functions of higher degree and (*ii*) continuous distributions.

Moreover, we extend the results of Gessford and Karlin (1958) and Lamond et al. (1995) on the structure of optimal solutions under piecewise linear rewards. In particular, we show that the optimal decision rules are always determined by two critical points, say $\hat{a}_{1t}$ and $\hat{a}_{2t}$, even in the presence of significant turbine head effect. Previous results (Drouin et al. (1996), Gessford and Karlin (1958), Lamond and Lang (1996), Lamond et al. (1995)) assumed constant turbine efficiency, independent of storage, which is considered unrealistic in practice for many hydroelectric systems (at the exception of run-of-the-river plants). We exploit this result to derive a piecewise polynomial approximation of the function $V_t(s_t)$ in which the points where the function is not smooth are taken as "not-a-knot" nodes (see, e.g., De Boor (1978)), and we present numerical results showing that this approach is much more adequate than the ordinary spline functions when the revenues are piecewise linear, as we assume in our model.

Some authors have also addressed the problem of approximating the value function of a multireservoir system by a piecewise linear function, using linear programming and Benders decomposition: Pereira and Pinto (1991) assume the reward is proportional to the energy produced and Archibald et al. (1999) assume the reward is a concave, piecewise linear function of energy. Moreover, some extensions of the work of Gessford and Karlin (1958) to multireservoir systems were also examined in Lang (1994) and Archibald et al. (2001). Unlike the present paper, these studies all assume a constant head on the turbines.

The paper is organized as follows. In §2 we derive a formula for computing the exact expectation of a piecewise polynomial function. In §3, we give a model of a single reservoir with stochastic inflows and piecewise linear revenues of hydro-electric production. In §4 we derive the structure of optimal solutions. In §5 we summarize our DP algorithm using the structure of optimal solutions and the expectation formula. Then, a numerical illustration is given in §6, with concluding remarks in §7.

2. PIECEWISE POLYNOMIAL EXPECTATION

To illustrate the difficulty with Gaussian quadrature rules when the function to integrate is not smooth enough, we evaluate the expectation of the function $v(X) = |X|$ under a standard normal distribution. We see easily that

$$I = \int_{-\infty}^{\infty} \frac{1}{\sqrt{2\pi}} e^{-x^2/2} |x| \, dx = \sqrt{2/\pi} \approx 0.80.$$

Using Gaussian quadrature with 3 nodes as in eq. (E5) of Foufoula-Georgiou and Kitanidis (1988), we obtain

$$Q_3 = \frac{2}{3} v(0) + \frac{1}{6} \left[v\left(-\sqrt{3}\right) + v\left(\sqrt{3}\right) \right] = 1/\sqrt{3} \approx 0.58.$$

The absolute error is $= I - Q_3 = 0.22$ and thus the relative error is 28%, which is hardly negligible. See Lamond and Bachar (1998) for a numerical study of the effect of discretization on the optimal value function $V_t(s_t)$ under an infinite planning horizon. We now present a method to integrate exactly a piecewise polynomial function (which is not necessarily smooth at the break points). The following notation is used in the sequel:

$$x \wedge y = \min(x, y)$$
$$x \vee y = \max(x, y)$$
$$(x)^+ = \max(x, 0).$$

Consider a continuous, nonnegative random variable X with density function $f(x)$. For $k = 0, 1, 2, \ldots$ and $x \geq 0$, define the cumulative functions

$$F_k(x) = \int_0^x F_{k-1}(y) \, dy, \tag{3}$$

where $F_{-1}(x) \equiv f(x)$. It is convenient to define $F_k(x) = 0$ for all $x < 0$. We will suppose these functions can be computed accurately and efficiently (this can be done, for instance, if some representation has already been constructed, analytically or by a numerical method, and can be used, say, by calling a library subroutine). An example is given in §6.

Suppose $v(x)$ is a piecewise polynomial function of degree n over its domain $[0, U]$, and let the numbers $0 = b_0 < b_1 < b_2 < \cdots < b_m < b_{m+1} = U$ be the break points. Let also the ith polynomial be

$$p_i(x) = \sum_{k=0}^{n} c_i^{(k)} x^k, \tag{4}$$

with its ℓth derivative

$$p_i^{(\ell)}(x) = \sum_{k=\ell}^{n} \frac{k!}{k-\ell)!} c_i^{(k)} x^{k-\ell}, \tag{5}$$

for $\ell = 0, \ldots, n$. Here $c_i^{(k)}$ is the coefficient of x^k in the ith polynomial $p_i(x)$. Of course, $p_i^{(0)}(x) \equiv p_i(x)$ and $p_i^{(\ell)}(x) \equiv 0$ for $\ell > n$. Then we can write

$$v(x) = \sum_{i=1}^{m+1} g_i(x), \tag{6}$$

where

$$g_i(x) = \begin{cases} p_i(x) & \text{if } b_{i-1} \leq x \leq b_i, \\ 0 & \text{else.} \end{cases} \tag{7}$$

The integration method is based on the following results.

Lemma 2.1
For $i = 1, \ldots, m+1$ and $a, x \geq 0$,

$$\int_0^x p_i(a+y) f(y)\, dy = \sum_{\ell=0}^{n} (-1)^\ell p_i^{(\ell)}(a+x) F_\ell(x). \tag{8}$$

Proof
For $\ell = 0, 1, \ldots, n+1$, define the functions

$$H_i^{(\ell)}(a,x) = \int_0^x p_i^{(\ell)}(a+y) F_{\ell-1}(y)\, dy.$$

We want to find $H_i^{(0)}(a,x)$. Now for $\ell = n+1$, we have trivially $H_i^{(n+1)}(a,x) = 0$. For $\ell \leq n$, integration by parts yields directly

$$H_i^{(\ell)}(a,x) = p_i^{(\ell)}(a+x) F_\ell(x) - H_i^{(\ell+1)}(a,x). \tag{9}$$

Successive application of eq. (9) for $\ell = n, n-1, \ldots, 0$ leads to eq. (8). □

Lemma 2.2
For $i = 1, 2, \ldots, m+1$ and $0 \leq a \leq U$, define

$$w_i(a) = \int_0^{U-a} g_i(a+y) f(y)\, dy. \tag{10}$$

Then

$$w_i(a) = \sum_{\ell=0}^{n} (-1)^\ell [p_i^{(\ell)}(b_i) F_\ell(b_i - a) - p_i^{(\ell)}(b_{i-1}) F_\ell(b_{i-1} - a)]. \tag{11}$$

Proof
From eq. (10), a simple change of variable gives

$$w_i(a) = \int_a^U g_i(t) f(t-a)\, dt.$$

We claim this is equivalent to

$$w_i(a) = \int_{a \vee b_{i-1}}^{a \vee b_i} p_i(t) f(t-a)\, dt. \tag{12}$$

Figure 1: Expectation of a polynomial piece

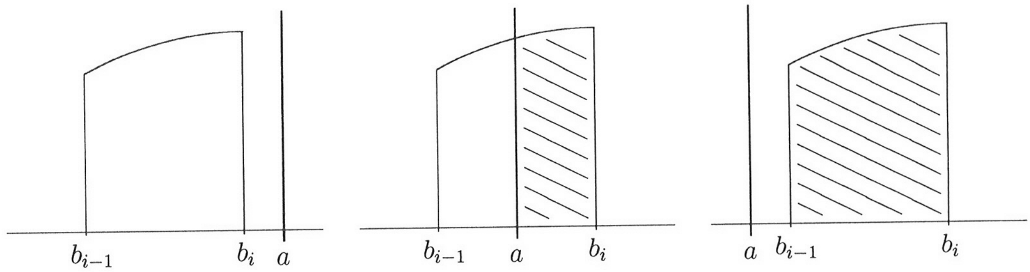

$(i)\ a > b_i$ $(ii)\ b_{i-1} < a < b_i$ $(iii)\ a < b_{i-1}$

Indeed, according to eq. (7), $g_i(t) = 0$ for $t < b_{i-1}$ and $t > b_i$. There are three cases to consider: (*i*) $a > b_i$, (*ii*) $b_{i-1} < a < b_i$ and (*iii*) $a < b_{i-1}$ (see Figure 1). In case (*i*), eq. (12) gives $w_i(a) = 0$, which is the correct value. In cases (*ii*) and (*iii*), eq. (12) gives the upper limit b_i and the lower limits a and b_{i-1} respectively, as required. Changing back to the original variable, we get

$$w_i(a) = \int_0^{(b_i - a)^+} p_i(a+y) f(y)\, dy - \int_0^{(b_{i-1} - a)^+} p_i(a+y) f(y)\, dy.$$

The result follows from Lemma 2.1 after we drop the positive parts in eq. (8) because $F_\ell(b_i - a) = 0$ for $a > b_i$, giving eq. (11). □

Theorem 2.3
For $0 \leq a \leq U$, *and with* $v(x)$ *given in eq. (6), define the function of expected values*

$$w(a) = E[v((a+X) \wedge U)]. \tag{13}$$

Then

$$w(a) = [1 - F_0(U - a)]v(U) + \sum_{i=j(a)}^{m+1} w_i(a), \tag{14}$$

where $j(a) = \min\{i : b_i \geq a\}$.

Proof
By definition of expectation, we have

$$\begin{aligned} w(a) &= \int_0^{\infty} v((a+y) \wedge U) f(y)\, dy \\ &= v(U) \int_{U-a}^{\infty} f(y)\, dy + \int_0^{U-a} v(a+y) f(y)\, dy. \end{aligned}$$

This gives eq. (14) by definition of $v(x)$ in eq. (6) and $w_i(a)$ in eq. (10), and noting that eq. (11) implies $w_i(a) = 0$ for $i < j(a)$, because $F_\ell(b_i - a) = 0$ for $a > b_i$. See Figure 2. □

We now describe the integration procedure. For convenience, we will refer to it as the "PRIM algorithm" because it uses the primitives of the density function. It requires three major steps.

Figure 2: Expectation of a piecewise polynomial function

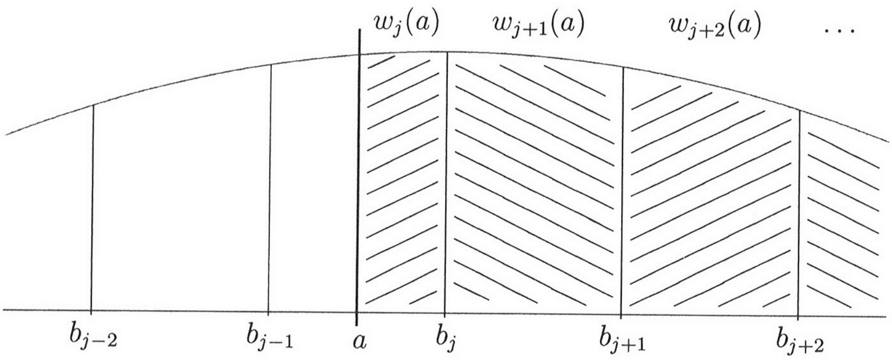

Algorithm PRIM for Piecewise Polynomial Integration

1. For a given probability density function $f(x)$, obtain the indefinite integrals analytically, if possible. Otherwise, tabulate the functions $F_\ell(x)$ in the interval $[0, U]$, for $\ell = 0, \ldots, n$, using a numerical integration method such as Simpson's rule, and construct, e.g., a cubic spline representation to interpolate $F_\ell(x)$ at arbitrary points $x \in [0, U]$.
2. For a given piecewise polynomial function $v(x)$ with break points b_i, use eq. (5) to evaluate the derivatives $p_i^{(\ell)}(b_i)$ and $p_i^{(\ell)}(b_{i-1})$, for $i = 1, \ldots, m+1$ and $\ell = 0, \ldots, n$.
3. To evaluate $w(a)$ at an arbitrary point a, use eq. (11) and eq. (14).

In this procedure, step 1 is a preparatory chore that may involve analytical integration, numerical analysis and computer programming. An example is given in §6 for the gamma distribution. At step 2, the derivatives $p_i^{(\ell)}(b_i)$ and $p_i^{(\ell)}(b_{i-1})$ are tabulated beforehand using $O(mn^2)$ arithmetic operations. At step 3, to evaluate $w(a)$ at a given point a, we need to evaluate each of the functions $F_0, \ldots, F_n$ at most $m+1$ times, and next $O(mn)$ arithmetic operations are required to compute $w_i(a)$ and $w(a)$. Usually, step 3 is repeated for several values of a.

3. STOCHASTIC MODEL FOR ONE RESERVOIR

Our notation is compatible with Drouin et al. (1996), Lamond and Lang (1996), and Lamond et al. (1995). Our model is basically the same, although our variables are continous as in Gessford and Karlin (1958). Our model differs in that it allows turbine head effects, as in Lamond and Bachar (1998). The system comprises one reservoir and one hydro-plant. We suppose the reservoir has a limited storage capacity of volume U, while the hydro-plant has unlimited production capacity. Decisions have to be made at the beginning of every T periods in the planning horizon. For $t = 1, \ldots, T$, let S_t denote the volume of water in storage in the reservoir at the beginning of period t and available for electricity production, and let Z_t denote the volume of water released through the turbines during period t. Then the drawdown volume, after turbine releases but before natural inflows are added, is $A_t = S_t - Z_t$. Let also D_t be the random variable giving the volume of water added to the contents of the reservoir during period t from uncontrolled natural inflows.

Moreover, we make the following three assumptions: (*i*) the volume D_t is not available for electricity production until period $t+1$, (*ii*) any volume of water in excess of the storage capacity is lost, without electricity production, and (*iii*) the random variables $D_1, \ldots, D_T$ are independent and have continuous distributions. With state and action variables S_t and A_t, respectively, the system dynamics are thus described by the following transition equation,

$$S_{t+1} = (A_t + D_t) \wedge U. \tag{15}$$

The state space is $\mathcal{S} = [0, U]$ and the set of allowable actions in period t, given $S_t = s$, is $\mathcal{A}_t(s) = [0, s]$.

The amount of electricity E_t produced during period t is given by the function $E(s, a)$ of the stored volumes $S_t = s$ and $A_t = a$. Here we neglect the effect of the random inflows D_t on the production of period t. Suppose the turbine efficiency is a function $\vartheta(x)$ of the stored volume x. Then

$$E(s, a) = \int_a^s \vartheta(x)\, dx = \Theta(s) - \Theta(a), \tag{16}$$

where $\Theta(x)$, the primitive of $\vartheta(x)$, gives the potential energy of a stored volume x. For example, in the special case when the turbine efficiency is a constant θ_0 independent of the stored volume, we have simply $E(s, a) = \theta_0(s - a)$, as in Drouin et al. (1996), Gessford and Karlin (1958), Lamond and Lang (1996), and Lamond et al. (1995). A slightly more realistic approach assumes, as in Lamond and Bachar (1998), the turbine efficiency is an affine function of storage

$$\vartheta(s) = \theta_0 + \theta_1 s,$$

where θ_0 and θ_1 are given constants. Then

$$\Theta(s) = \theta_0 s + \frac{\theta_1}{2} s^2, \tag{17}$$

so that $E(s, a)$ is a quadratic function of s and a. Polynomials of degree higher than two are often used in the literature as well. See, e.g., Boukhtouta and Lamond (2001), Mariño and Loaiciga (1985), and Soares and Carneiro (1990).

We assume the revenue from sales of electricity in period t is given by a piecewise linear function

$$\rho_t(E) = \begin{cases} h_{1t}E & \text{if } 0 \le E \le Q_t \\ h_{1t}Q_t + h_{2t}(E - Q_t) & \text{if } E \ge Q_t, \end{cases} \tag{18}$$

where E is the amount of hydro-electricity produced during the period. Under the assumption $h_{1t} > h_{2t}$, this corresponds to a (lucrative) primary market with regular price h_{1t} and limited demand Q_t and a secondary market for excess production with unlimited demand but at a discounted price h_{2t}. The reward R_t for period t is therefore determined by the function

$$r_t(s_t, a_t) = \rho_t\left(E(s_t, a_t)\right) = \rho_t\left(\Theta(s_t) - \Theta(a_t)\right). \tag{19}$$

As in the introduction, we assume there is a terminal reward function $V_{T+1}(s_{T+1})$ and we define the functions $V_t(s_t)$ recursively, using the dynamic programming equations (1, 2).

Overall, our modelling assumptions have the same limitations as in Drouin et al. (1996), Gessford and Karlin (1958), and Lamond et al. (1995), except that our model also allows the turbine head variations to be taken into account. As pointed out in Lamond et al. (1995), limited turbine capacity could also be included in the model by adding a third interval in eq. (18) and setting to zero the price of energy in excess of production capacity. We also argue that it is reasonable to neglect the impact of assumptions (*i*) and (*ii*) about the natural inflows if, for instance, the planning periods are short enough so that the mean inflow is small compared to the storage capacity of the reservoir. Moreover, the numerical study of Tejada-Guibert et al. (1995) suggests our independance assumption (*iii*) is justified as well in the context of planning hydropower production (we note here that our method of analysis could also be used when seasonal variations are taken into account).

We want to find an optimal policy $\pi = (\gamma_1, \ldots, \gamma_T)$ that maximizes $V_1(s_1)$ for all $s_1 \in [0, U]$. Each decision rule $\gamma_t(s_t)$ is a function such that action $a_t = \gamma_t(s_t)$ is selected in period t if the storage at beginning of period t is $S_t = s_t$. The optimal decision rule $\gamma_t(s_t)$ gives the action a_t

maximizing eq. (1) for every state s_t. Thus our optimization algorithm must construct the T functions $\gamma_1, \ldots, \gamma_T$. Now from eq. (15), we can write the expectation in eq. (2) as

$$W_{t+1}(s_t, a_t) = W_{t+1}(a_t) = E[V_{t+1}((a_t + D_t) \wedge U)]. \tag{20}$$

This is the same expectation as in eq. (13) with X replaced by D_t. Hence we can compute $W_{t+1}(a_t)$ using Theorem 2.3, provided we approximate $V_{t+1}(\cdot)$ by a piecewise polynomial function.

4. STRUCTURE OF OPTIMAL SOLUTIONS

The optimal solution $V_t(s_t)$ of the dynamic programming equations (1, 2) cannot be obtained in closed form, although its essential properties can be derived analytically. A closed form solution exists, however, for the optimal decision rule $\gamma_t(s_t)$. We now derive this closed form and use it to describe the structure of the optimal solution $V_t(s_t)$ as well as the form of an adequate representation of $V_t(s_t)$ by a piecewise polynomial approximation. In the special case when the turbine efficiency is constant, i.e., $\vartheta(s) = \theta_0$, our solution reduces to the results of Gessford and Karlin (1958), and Lamond et al. (1995). Hence we extend previous results to the more realistic case when turbine efficiency varies due to head effects.

It is convenient to make a change of variables and express the state and action variables in potential energy units instead of storage volume. Let $x_t = \Theta(s_t)$ and $y_t = \Theta(a_t)$ be the new state and action variables, respectively. Using the inverse transformation $\Theta^{-1}(\cdot)$ to convert from potential energy to volume, we also define the functions

$$v_t(x_t) = V_t(s_t) = V_t(\Theta^{-1}(x_t))$$

and

$$w_{t+1}(y_t) = W_{t+1}(a_t) = W_{t+1}(\Theta^{-1}(y_t)).$$

Then the Bellman equations (1–2) become respectively

$$v_t(x_t) = \max_{y_t \in [0, x_t]} r_t(\Theta^{-1}(x_t), \Theta^{-1}(y_t)) + \beta w_{t+1}(y_t)$$

$$w_{t+1}(y_t) = E[v_{t+1}(\Theta(\Theta^{-1}(y_t) + D_t) \wedge \Theta(U))],$$

where we have used eq. (20) and D_t is the random variable of natural inflows in period t. Now, for an arbitrary constant $d \geq 0$, define the function

$$z_d(y) = \Theta(\Theta^{-1}(y) + d).$$

Then $z_d(y) \wedge \Theta(U)$ is the potential energy in the reservoir after an inflow of d units of water is added to an initial storage of potential energy y. In the sequel, we make the further assumption that the head function $\vartheta(a)$ is differentiable and that the ratio $\vartheta'(a)/\vartheta(a)$ is a nonincreasing function of a for all $a \in [0, U]$. This condition does not seem to be very restrictive in practice because one can show easily that it holds, in particular, when the head function $\vartheta(a)$ is concave increasing. In Lamond (2001), the latter condition is shown to hold for a number of examples of reservoir shapes in which the reservoir walls are leaning toward the outside, as would be the case with hydroelectric dams built on natural sites.

Lemma 4.1
Suppose the ratio $\vartheta'(a)/\vartheta(a)$ is a nonincreasing function of a for all $a \in [0, U]$. Suppose also that $v_t(x_t)$ is concave nondecreasing. Then $w_{t+1}(y_t)$ is concave nondecreasing.

Proof
It suffices to show that $v_t(z_d(y_t) \wedge \Theta(U))$ is concave nondecreasing for any value $d \geq 0$ of the random variable D_t. This, in turn, is true if $z_d(y_t)$ is concave nondecreasing. Now $z_d(y_t) = \Theta(a_t + d)$ and $y_t = \Theta(a_t)$, so we have

$$z_d'(y_t) = \frac{dz_d}{dy_t} = \frac{\vartheta(a_t + d)}{\vartheta(a_t)} \geq 0,$$

showing that $z_d(y_t)$ is a nondecreasing function. Applying the chain rule, we get

$$z_d''(y_t) = \frac{d}{da_t} z_d'(y_t) \frac{da_t}{dy_t} = \frac{\vartheta(a_t + d)}{(\vartheta(a_t))^2} \left[\frac{\vartheta'(a_t + d)}{\vartheta(a_t + d)} - \frac{\vartheta'(a_t)}{\vartheta(a_t)} \right] \leq 0,$$

where the last inequality follows because $\vartheta'(a)/\vartheta(a)$ is nonincreasing, by hypothesis. □

Lemma 4.2
Suppose the ratio $\vartheta'(a)/\vartheta(a)$ is a nonincreasing function of a for all $a \in [0, U]$. Suppose also $v_{T+1}(x_{T+1})$ is concave nondecreasing, and $h_{1t} > h_{2t}$. Then $v_t(x_t)$ is concave nondecreasing, for $t = 1, \ldots, T$.

Proof
By induction on t. The result is true for $t = T + 1$ by hypothesis. Suppose now $v_{t+1}(x_{t+1})$ is concave nondecreasing. Then $w_{t+1}(y_t)$ is concave nondecreasing too, by Lemma (4.1). Now, from eq. (1) and (19), we have that

$$v_t(x_t) = \max_{0 \leq y_t \leq x_t} \rho_t(x_t - y_t) + \beta w_{t+1}(y_t).$$

But the function $\rho_t(x_t - y_t) + \beta w_{t+1}(y_t)$ is concave on its domain $\mathcal{C}$, where

$$\mathcal{C} = \{(x, y) : 0 \leq x \leq \Theta(U),\ 0 \leq y \leq x\}.$$

Concavity of $v_t(x_t)$ therefore is a well-known result. See, e.g., pp. 525–526 of Heyman and Sobel (1984). Finally, for $x' > x$, the action set $[0, x']$ contains the set $[0, x]$. This implies $v_t(x_t)$ is nondecreasing. □

This result is remarkable, because in the volumetric domain, the functions $V_t(s_t)$ are not concave (see Figure 3d). In the energy domain, however, the problem recovers the same essential properties that were exploited in Gessford and Karlin (1958), and Lamond et al. (1995), although our results are not restricted to the special case of constant turbine efficiency.

Theorem 4.3
Suppose the conditions of lemmas 4.1 and 4.2 are satisfied. Then for each $t = 1, \ldots, T$, there are two critical numbers $\hat{y}_{1t}$ and $\hat{y}_{2t}$, given by

$$\hat{y}_{it} = \arg \max_{0 \leq y \leq \Theta(U)} -h_{it} y + \beta w_{t+1}(y), \tag{21}$$

for $i = 1, 2$, such that $\hat{y}_{1t} \leq \hat{y}_{2t}$. Moreover, an optimal set of actions is given by the decision rule

$$y_t^*(x_t) = \begin{cases} x_t & \text{if } 0 \leq x_t \leq \hat{y}_{1t} \\ \hat{y}_{1t} & \text{if } \hat{y}_{1t} \leq x_t \leq \hat{y}_{1t} + Q_t \\ x_t - Q_t & \text{if } \hat{y}_{1t} + Q_t \leq x_t \leq \hat{y}_{2t} + Q_t \\ \hat{y}_{2t} & \text{if } \hat{y}_{2t} + Q_t \leq x_t \leq U. \end{cases} \tag{22}$$

Table 1: Interpretation of optimal decision rules

Zone	State (s_t)	Action (a_t)	Explanation
1	$[0, \hat{a}_{1t}]$	s_t	no energy production
2	$[\hat{a}_{1t}, \hat{s}_{1t}]$	$\hat{a}_{1t}$	discharge to $\hat{a}_{1t}$ ($E_t < Q_t$)
3	$[\hat{s}_{1t}, \hat{s}_{2t}]$	$\varphi_t(s_t)$	produce energy $E_t = Q_t$
4	$[\hat{s}_{2t}, U]$	$\hat{a}_{2t}$	discharge to $\hat{a}_{2t}$ ($E_t > Q_t$)

Proof

The inequality $\hat{y}_{1t} \leq \hat{y}_{2t}$ is obvious.

For $x_t \leq Q_t$, we have $\rho_t(x_t - y_t) = h_{1t}(x_t - y_t)$. There are two cases: if $x_t < \hat{y}_{1t}$, then the optimal action is $y_t^* = x_t$, else it is $y_t^* = \hat{y}_{1t}$.

For $x_t > Q_t$, the function $\rho_t(x_t - y_t)$ has two linear pieces. There are two main cases. If $x_t > \hat{y}_{2t} + Q_t$, then the optimal action is $y_t^* = \hat{y}_{2t}$, else there are three subcases. If $x_t < \hat{y}_{1t}$ then $y_t^* = x_t$. If $\hat{y}_{1t} \leq x_t \leq \hat{y}_{1t} + Q_t$ then $y_t^* = \hat{y}_{1t}$. Finally, if $x_t > \hat{y}_{1t} + Q_t$ then $y_t^* = x_t - Q_t$. In all cases, the optimal action is given by eq. (22). □

To transform back to volumetric units, it is convenient to define the function $\varphi_t(s_t)$ giving the action a_t such that $E(s_t, a_t) = Q_t$. This action thus corresponds to a hydroelectric production of exactly Q_t units of energy. When storage is insufficient, we define $\varphi_t(s_t) = 0$. Then

$$\varphi_t(s_t) = \begin{cases} \Theta^{-1}(\Theta(s_t) - Q_t) & \text{if } \Theta(s_t) \geq Q_t, \\ 0 & \text{else.} \end{cases} \tag{23}$$

For example, with constant turbine efficiency, we have

$$\varphi_t(s_t) = (s_t - Q_t/\theta_0)^+.$$

When the turbine efficiency is an affine function of storage, we have from eq. (17) and (16) that $a = \varphi_t(s_t)$ satisfies the quadratic equation

$$\frac{\theta_1}{2}a^2 + \theta_0 a - C_t = 0,$$

where

$$C_t = \theta_0 s_t + \frac{\theta_1}{2}s_t^2 - Q_t,$$

provided $C_t \geq 0$. Thus we have in this case

$$\varphi_t(s_t) = \frac{-\theta_0 + \sqrt{\theta_0^2 + 2\theta_1 C_t}}{\theta_1}.$$

The optimal decision rules can now be described.

Corollary 4.4

Suppose the conditions of Theorem 4.3 are satisfied. Then for each $t = 1, \ldots, T$, there are two critical numbers $\hat{a}_{1t}$ and $\hat{a}_{2t}$, given by

$$\hat{a}_{it} = \arg \max_{0 \leq a \leq U} -h_{it}\Theta(a) + \beta W_{t+1}(a), \tag{24}$$

Figure 3: Structure of optimal solutions

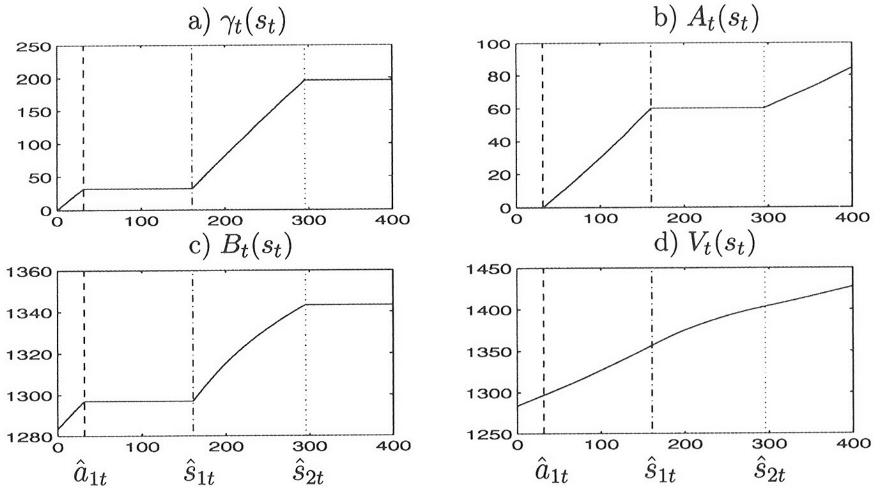

for $i = 1, 2$, such that $\hat{a}_{1t} \leq \hat{a}_{2t}$. Moreover, with

$$\hat{s}_{it} = \Theta^{-1}(\Theta(\hat{a}_{it}) + Q_t), \tag{25}$$

the optimal decision rules are given in Table 1.

This result has important consequences. First, the optimization of eq. (1) for all states $s_t \in [0, U]$ can be performed simply by doing a pair of line searches to find $\hat{a}_{1t}$ and $\hat{a}_{2t}$. Next, the optimal decision rules have a simple interpretation, with the state space in period t partitioned in four zones, as in Table 1. Finally, the special structure of $\gamma_t(s_t)$, plotted in Figure 3a, can be exploited to derive a suitable piecewise polynomial approximation of the function $V_t(s_t)$.

Indeed, replacing a_t by $\gamma_t(s_t)$ in eq. (1), and using (19, 20), we have

$$V_t(s_t) = A_t(s_t) + B_t(s_t),$$

where

$$A_t(s_t) = \rho_t \left(\Theta(s_t) - \Theta(\gamma_t(s_t))\right)$$
$$B_t(s_t) = \beta W_{t+1}(\gamma_t(s_t)).$$

Using Table 1, we can get a separate expression for these functions in each of the four zones. An illustrative example is plotted in Figure 3. One can see easily, as in Figure 3b, that $A_t(s_t)$ is constant in zones 1 and 3. Furthermore, if the potential energy function $\Theta(s_t)$ is a polynomial, then $A_t(s_t)$ is polynomial in zones 2 and 4. Similarly, $B_t(s_t)$ is constant in zones 2 and 4, as in Figure 3c, but it is non polynomial in zones 1 and 3.

The resulting expressions for $V_t(s_t)$ are given in Table 2, where the constants

$$K_{2t} = -h_{1t}\Theta(\hat{a}_{1t}) + \beta W_{t+1}(\hat{a}_{1t})$$
$$K_{4t} = h_{1t}Q_t - h_{2t}(\Theta(\hat{a}_{2t}) + Q_t) + \beta W_{t+1}(\hat{a}_{2t})$$

are used. If $\Theta(s_t)$ is a polynomial then $V_t(s_t)$ naturally has a polynomial representation in zones 2 and 4. In zones 1 and 3, however, a piecewise polynomial approximation using splines is suitable. In specifying a spline approximation over an interval, however, one has to decide what boundary conditions should be imposed at the end points. When the first derivative of the

Table 2: Expressions for the function $V_t(s_t)$

Zone	State (s_t)	$V_t(s_t)$	Polynomial representation
1	$[0, \hat{a}_{1t}]$	$\beta W_{t+1}(s_t)$	spline approximation
2	$[\hat{a}_{1t}, \hat{s}_{1t}]$	$h_{1t}\Theta(s_t) + K_{2t}$	exact polynomial
3	$[\hat{s}_{1t}, \hat{s}_{2t}]$	$\beta W_{t+1}(\varphi_t(s_t)) + h_{1t}Q_t$	spline approximation
4	$[\hat{s}_{2t}, U]$	$h_{2t}\Theta(s_t) + K_{4t}$	exact polynomial

Table 3: Existence and boundary conditions for each zone

			Boundary conditions for splines	
Zone	State	Existence	Left	Right
1	$[0, \hat{a}_{1t}]$	$\hat{a}_{1t} > 0$	not-a-knot	not-a-knot if $\hat{a}_{1t} = U$ else $V_t'(\hat{a}_{1t}) = h_{1t}\vartheta(\hat{a}_{1t})$
2	$[\hat{a}_{1t}, \hat{s}_{1t}]$	$\hat{a}_{1t} < U$	N/A	N/A
3	$[\hat{s}_{1t}, \hat{s}_{2t}]$	$\hat{s}_{1t} < U$ and $\hat{a}_{2t} > 0$	not-a-knot if $\hat{a}_{1t} = 0$ else $V_t'(\hat{s}_{1t}) = h_{1t}\vartheta(\hat{s}_{1t})$	not-a-knot if $\hat{s}_{1t} = U$ else $V_t'(\hat{s}_{2t}) = h_{2t}\vartheta(\hat{s}_{2t})$
4	$[\hat{s}_{2t}, U]$	$\hat{s}_{2t} < U$	N/A	N/A

function is known, its value at an end point can be used as a boundary condition. Otherwise, the 'not-a-knot' condition is preferred. See De Boor (1978).

It is straightforward to see that the functions $V_t(s_t)$ are continuous. If the random variables D_t are continuous, then the functions $W_{t+1}(a_t)$ are differentiable everywhere. This is not always the case for $V_t(s_t)$. For example, if $V_{T+1}(s_{T+1}) = 0$ identically, then the optimal decision rule for period T is $\gamma_T(s_T) = 0$ for all $s_T \in [0, U]$. Then $V_T(s_T) = \rho_T(\Theta(s_T))$, which is not differentiable at $s_T = \Theta^{-1}(Q_T)$. We note that $\hat{a}_{1T} = \hat{a}_{2T} = 0$ in this case, and that the maximum was found at the left boundary of eq. (24). We also observe that the zones 1 and 3 do not exist, so that no spline approximations are required in this special case.

Lemma 4.5
The existence conditions for each zone are given in Table 3, along with the correct boundary conditions for the spline approximations.

Proof
The zone existence conditions are trivial. When both zones 1 and 2 exist, we have $0 < \hat{a}_{1t} < U$, in which case

$$-h_{1t}\Theta'(\hat{a}_{1t}) + \beta W_{t+1}'(\hat{a}_{1t}) = 0. \tag{26}$$

From Table 2, this implies continuity of $V_t'(s_t)$ at $s_t = \hat{a}_{1t}$. Hence the spline approximation in Zone 1 should be specified with a 'not a knot' condition on the left (at $s_t = 0$) and with $V_t'(\hat{a}_{1t}) = h_{1t}\vartheta(\hat{a}_{1t})$ on the right.

Similarly, existence of Zone 3 implies that of Zone 2. When $0 < \hat{a}_{1t} < U$, continuity of the derivative on the left of the interval, at $s_t = \hat{s}_{1t}$, is verified directly from Table 2, using the chain rule:

$$\beta W_{t+1}'(\hat{a}_{1t})\varphi_t'(\hat{s}_{1t}) = \frac{\beta W_{t+1}'(\hat{a}_{1t})\Theta'(\hat{s}_{1t})}{\Theta'(\hat{a}_{1t})} = h_{1t}\Theta'(\hat{s}_{1t}),$$

where the first equality follows from

$$\varphi_t'(\hat{s}_{1t}) = \Theta'(\hat{s}_{1t})/\Theta'(\hat{a}_{1t})$$

and the second equality follows from eq. (26). Further, when $0 < \hat{a}_{2t} < U$, we have

$$-h_{2t}\Theta'(\hat{a}_{2t}) + \beta W'_{t+1}(\hat{a}_{2t}) = 0. \tag{27}$$

Then continuity of the derivative on the right of the interval, at $s_t = \hat{s}_{2t}$, is verified in the same way, using eq. (27). Therefore, the spline approximation in Zone 3 should be specified with $V'_t(\hat{s}_{1t}) = h_{1t}\vartheta(\hat{s}_{1t})$ on the left, and $V'_t(\hat{s}_{2t}) = h_{2t}\vartheta(\hat{s}_{2t})$ on the right. □

5. DYNAMIC PROGRAMMING ALGORITHM

We now present a dynamic programming algorithm to solve eq. (1) for the model of §3, using the results of §§2–4. For $j = 1, 2, 3, 4$, let m_j be the number of subintervals in Zone j. If Zone j does not exist, set $m_j = 0$. For $j = 2, 4$, when Zone j exists, we have $m_j = 1$. For $j = 1, 3$, the parameters m_1 and m_3 need to be specified in the model data. In the notation of §2, the total number of intervals is then

$$m + 1 = m_1 + m_2 + m_3 + m_4,$$

and there are $m + 2$ break points, with $b_0 = 0$ and $b_{m+1} = U$. Let also ω_j denote the index of the first break point of Zone j. Then $\omega_1 = 0$ and $\omega_j = \omega_{j-1} + m_{j-1}$, for $j = 2, 3, 4, 5$, so that Zone j is the interval $[b_{\omega_j}, b_{\omega_{j+1}}]$.

Algorithm ZONE for Continuous Dynamic Programming

1. Set $t = T + 1$.
2. Piecewise polynomial approximation of terminal rewards $V_{T+1}(s_{T+1})$: choose a number m of grid points b_i, $i = 1, \ldots, m$, such that $0 = b_0 < b_1 < \ldots < b_m < b_{m+1} = U$, and obtain the coefficients $c_{ij}^{(k)}$ of the polynomials $p_{i,T+1}(x)$ in eq. (4), for $k = 0, \ldots, n$ and $i = 1, \ldots, m+1$.
3. Set $t = t - 1$. If $t = 0$ stop.
4. For $j = 1, 2$, perform a line search to find $\hat{a}_{jt}$ in eq. (24). At each iteration of the line search, evaluate the function $W_{t+1}(a_t)$ of eq. (20) using eq. (14) and (11).
5. Construct a piecewise polynomial approximation for $V_t(s_t)$ based on tables 2 and 3. If zones 1 and 3 exist, choose equally spaced break points for the spline approximations and evaluate $W_{t+1}(\cdot)$ at the interior break points using eq. (14) and (11).
6. Go to 3.

The computational complexity of the above DP algorithm is dominated by the number of evaluations of the expected value function $W_{t+1}(\cdot)$. There are T major iterations. Each line search, at Step 4, requires at most N function evaluations (say), and the spline constructions at Step 5 require $m_1 + m_3 - 2$ evaluations. Then the complete algorithm requires $T \times (2N + m_1 + m_3)$ evaluations of $W_{t+1}(\cdot)$. Detailed algorithms are given in Everitt (1987) for the line search, and in De Boor (1978) for the spline approximation. The number N of steps of the line search depends on the method used. For example, with the Golden Section method, the number of steps required to narrow down the search to an interval of width ϵ is $N = \ln(\epsilon/U)/\ln(1-g)$ where $g = (\sqrt{5} - 1)/2$ is the Golden Ratio.

6. NUMERICAL TESTS

We now illustrate the procedure with a numerical example. We implemented our DP algorithm in the C language on an IBM Risc-6000 computer, Model 3BT. We present numerical results for a continuous version of the model treated in Lamond and Bachar (1998), where 30 cases were considered for the natural inflows (i.e., 10 means × 3 coefficients of variation), under discretized normal and lognormal distributions. Here, we present only three cases, corresponding to cases

Table 4: Data for natural inflow distributions

Case	μ	σ/μ	r	λ
10	152	0.15	44	0.28947
11	152	0.25	16	0.10526
12	152	0.35	8	0.05263

10, 11 and 12 of Lamond and Bachar (1998), and we use the continuous gamma distribution, instead.

As in Lamond and Bachar (1998), we take a discount factor $\beta = 0.95$. The reservoir capacity is $U = 400$ water units. The turbine efficiency function is $\vartheta(s) = 1 + 0.0025s$, corresponding to efficiencies of $\vartheta(0) = 1$ and $\vartheta(U) = 2$, respectively, at minimum and maximum storage. Thus our model represents a hydroelectric system with rather important head variations. We assume the primary demand is $Q_t = 160$ energy units, with the primary and secondary prices $h_{1t} = 0.375$ and $h_{2t} = 0.125$, respectively, in all periods. Then our model satisfies all the assumptions of §4 and our DP algorithm ZONE is applicable. We choose the number of periods $T = 200$ which, for all practical purposes, corresponds to an infinite horizon. Indeed, the largest one-period revenue is 122.5, corresponding to the conversion of a full reservoir into electricity. The difference between $V_1(s)$, with $T = 200$, and the infinite horizon value is bounded by

$$122.5 \times \frac{0.95^{200}}{1 - 0.95} < 0.086,$$

and computations show that the optimal infinite horizon value is greater than 1280 for all $s \in \mathcal{S}$.

As in Gessford and Karlin (1958), we assume the natural inflows follow a gamma distribution with integer shape parameter r. For simplicity, we suppose the parameters are the same in all periods, i.e.,

$$D_t \sim \text{Gamma}(r, \lambda),$$

for $t = 1, \ldots, T$. Then the expected inflow is $\mu = r/\lambda$ and the variance is $\sigma^2 = r/\lambda^2$, giving a coefficient of variation of $\sigma/\mu = 1/\sqrt{r}$. For large r, this approximates a normal distribution. Data for our three inflow cases are given in Table 4.

The gamma density function is, for $x \geq 0$:

$$f(x) = \frac{1}{(r-1)!}(\lambda x)^{r-1}\lambda e^{-\lambda x}. \tag{28}$$

Writing $F_{-1}(x) = f(x)$ as in §2, we obtain the primitives $F_n(x)$, $n = 0, 1, 2, \ldots$, analytically, as follows.

Lemma 6.1
Define $\gamma_{-1,1} = 1$ *and* $\gamma_{-1,k} = 0$, *for* $k = 2, 3, \ldots, r$. *Moreover, for* $n = 0, 1, 2, \ldots$, *define inductively the cumulative sums*

$$\gamma_{n,k} = \sum_{\ell=1}^{k} \gamma_{n-1,\ell}.$$

Then, for $n = 0, 1, 2, \ldots$, *and* $x \geq 0$,

$$F_n(x) = \frac{1}{\lambda^n}\left[\sum_{\ell=0}^{n}(-1)^{n-\ell}\frac{\gamma_{n-\ell,r}}{\ell!}(\lambda x)^\ell - (-1)^n e^{-\lambda x}\sum_{\ell=0}^{r-1}\frac{\gamma_{n,r-\ell}}{\ell!}(\lambda x)^\ell\right]. \tag{29}$$

Proof

By induction on n. The result is trivially true for $n = -1$. We also observe that for all $n \geq 0$,

$$F_n(0) = \frac{1}{\lambda^n}[(-1)^n \gamma_{n,r} - (-1)^n \gamma_{n,r}] = 0.$$

as required. Next, straightforward algebra shows that eq. (29) satisfies $F_n'(x) = F_{n-1}(x)$, and the result follows. □

Before using eq. (29) to integrate polynomials of degree n, one can compute all constant terms with about $n(n+1)/2$ additions and $n(n+1)/2 + nr$ divisions. Next, evaluation of the functions $F_\ell(x)$, for $\ell = 0, \ldots, n$, requires about $n(n+1)/2 + nr$ pairs of additions and multiplications, plus one evaluation of the exponential function. Hence each time Step 3 of algorithm PRIM is executed, i.e., at Step 5 of algorithm ZONE, the exponential function is executed $m+1$ times and there are $(m+1) \times [n(n+1)/2 + nr]$ pairs of additions and multiplications. Cubic splines are often used in practice, giving $n = 3$. But from Table 4, we have $r \in \{8, 16, 44\}$ in our test cases, giving respectively $30(m+1)$, $78(m+1)$ and $162(m+1)$ pairs of additions and multiplications every time the function $W_{t+1}(a_t)$ is evaluated.

The primary purpose of the numerical tests is to evaluate the performance of our continuous DP algorithm ZONE by comparison with the naive, discrete DP algorithm (often called *value iteration*) shown in the appendix. As in Lamond and Bachar (1998), we use a value iteration algorithm with integer state, action and inflow sets $\{0, 1, 2, \ldots, 400\}$. With $f(i)$ given in eq. (28), the discrete inflow distribution is, for $i = 0, 1, 2, \ldots$,

$$\alpha_i = P(D_t = i) = f(i) / \sum_{j=0}^{\infty} f(j).$$

There are thus three sources of discretization error in the value iteration algorithm. Our continuous DP algorithm ZONE, however, avoids inflow discretization because the integration procedure PRIM of §2 is exact. Nonetheless, there are two sources of truncation error present. One is due to the spline approximation with a finite number of break points in zones 1 and 3 (analogous to state discretization). The other is due to the finite number of iterations in the line search procedure (analogous to action discretization), which we implemented using the Golden Section method (see, e.g., Everitt (1987)).

Another important purpose of the tests is to demonstrate the much higher quality of solutions obtained by the ZONE algorithm, regarding the accuracy of $V_t'(s_t)$ and the concavity of $v_t(x_t)$. The comparison is made with value iteration but also with a continuous DP algorithm in which cubic splines are used naively to approximate $V_t(s_t)$ (without applying the "not-a-knot" condition at the non-smooth points). We denote the latter as the SPLINE algorithm. The SPLINE algorithm is similar to ZONE except that the piecewise polynomial approximation at Step 5 is constructed with $m+1$ equal intervals on $[0, U]$, and ignoring the four zones of §4. In a sense, value iteration is a simple, very coarse, brute force method. SPLINE uses the exact PRIM integration procedure but its piecewise polynomial approximation, although efficient, is not really adequate for our model, due to the discontinuities of the second derivative at certain points (see Figure 4). On the other hand, ZONE uses exact PRIM integration with efficient and accurate piecewise polynomial approximation.

To make a fair evaluation of the methods with respect to CPU time, all three algorithms were coded in C by the same programmer, and all three codes were executed on the same machine. Moreover, the precision parameters of ZONE and SPLINE (i.e., N for line search, m_1, m_2 and m for spline approximations) were calibrated to give the function $V_1(s_1)$ with the same accuracy as with value iteration. The line searches stopped when $\hat{a}_{it}$ was within 0.1 units of the true optimal solution (golden section search was used by both ZONE and SPLINE).

Table 5: Summary of numerical results

	Value Iteration		SPLINE			ZONE		
Case	Error	CPU	$m+1$	Error	CPU	$m_1 = m_2$	Error	CPU
10	0.12	164.73	14	0.12	224.94	4	0.13	115.23
11	0.15	163.25	13	0.13	100.53	4	0.17	47.23
12	0.20	164.88	12	0.18	67.67	3	0.15	33.14

Figure 4: Concavity of $v_1(x_1)$ and accuracy of $v_1'(x_1)$

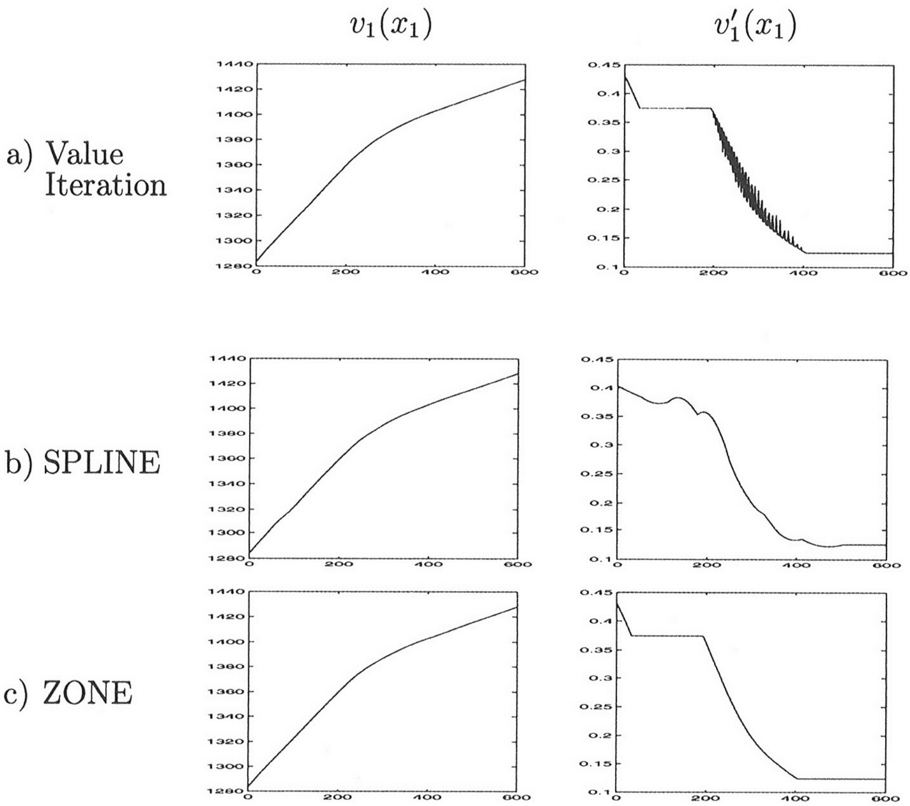

This was considered accurate enough since value iteration finds an integer solution, thus within 0.5 units of the true optimal solution. To calibrate the number of break points for the spline approximations, we estimated the true optimal function $V_1(s_1)$. This was accomplished through an initial run with $T = 200$ DP iterations, using an accuracy of 0.0001 for the line searches, and with $m_1 = m_2 = 29$ intervals for ZONE and $m+1 = 32$ intervals for SPLINE. Both methods gave very near solutions. The deviation between the solution obtained by value iteration and the true optimal value is less than 0.2 in all three inflow cases. With $V_t(s_t) \geq 1280$ for all $s_t \in S$, this gives a relative error of 0.016%.

Setting $T = 200$ iterations and a precision of 0.1 for the line searches, we successively solved SPLINE and ZONE with a decreasing number of intervals. The numerical results are summarized in Table 5. It gives the smallest number of intervals needed for a deviation of at most 0.2 from the optimal solution, along with the largest absolute deviation from the true optimal solution, and the CPU time in seconds. The efficiency of spline approximations is impressive, with 12 to 14 intervals for the SPLINE method, and only three or four intervals per zone for the ZONE method. The ZONE method has the smallest CPU time. Value iteration has the longest, except for Case 10, for which SPLINE is slower.

While all three methods give comparable accuracy of the function $V_1(s_1)$, the quality of the solutions varies considerably with respect to concavity of the function $v_1(x_1)$, in the energy domain, and accuracy of its first derivative $v_1'(x_1)$. As shown in Figure 4a, the derivative is very jumpy in Zone 3 for the function $v_1(x_1)$ obtained by value iteration, because of the effect of action discretization. Moreover, from Figure 4b we see that the derivative $v_1'(x_1)$ computed with the SPLINE algorithm is so rugged and irregular that $v_1'(x_1)$ is not monotone, therefore violating the concavity property of $v_1(x_1)$. This is due to the fact that the spline approximation assumes continuity of the second derivative everywhere, which is not true at $\hat{a}_{1t}$, $\hat{s}_{1t}$ and $\hat{s}_{2t}$. However, Figure 4c shows that the situation is entirely different with our ZONE algorithm, where $v_1'(x_1)$ is monotone nonincreasing and therefore $v_1(x_1)$ is concave, as required by Lemma 4.1.

7. CONCLUSION

In conclusion, we derived a new method (PRIM) for computing the expectation of a piecewise polynomial function and we extended previous results of Drouin et al. (1996), Gessford and Karlin (1958), Lamond and Lang (1996), and Lamond et al. (1995) on the structure of optimal solutions for a single reservoir and hydroplant with piecewise linear revenues. We used these theoretical results to obtain a continuous DP algorithm (ZONE) that is faster and more accurate than both discrete DP (value iteration) and a continuous DP method using splines on a fixed grid (SPLINE). While spline approximations on a fixed grid and a fixed quadrature rule are justified when the reward function is sufficiently smooth as in Foufoula-Georgiou and Kitanidis (1988), Johnson et al. (1993), and Philbrick and Kitanidis (2001), our work suggests they are not well-suited when the rewards are piecewise linear.

Finally, we note that further CPU time reduction could be achieved by approximating the primitive functions $F_n(x)$ (for example, using splines) rather than using exact evaluation of eq. (29). Another refinement would be to use a faster line search algorithm such as Newton's method.

8. ACKNOWLEDGMENTS

This research was supported in part by the Natural Sciences and Engineering Research Council of Canada, under Grant 0105560, and the *Fonds pour la Formation de Chercheurs et l'Aide à la Recherche du Québec*, under Grant 1570. The author would like to thank Michel Lantagne for his assistance with computer programming, implementation and numerical tests.

APPENDIX

We briefly review here the value iteration algorithm used in our numerical tests. We have the state set $\mathcal{S} = \{0, 1, 2, \ldots, U\}$ and the action sets $\mathcal{A}(s) = \{0, 1, 2, \ldots, s\}$, for $s \in \mathcal{S}$. The natural inflow probabilities are $\alpha_i = P(D_t = i) = f(i)/Q$ where $Q = \sum_{j=0}^{1000} f(j)$, with $f(x)$ given in eq. (28). Let also $\omega_i = 1 - \sum_{j=0}^{i-1} \alpha_j$. The immediate rewards $r_t(s, a)$ in period t are given by eq. (19), and the terminal rewards $V_{T+1}(s)$ are given data. The value iteration algorithm then proceeds as follows. See, e.g., Puterman (1994).

1. Let $t = T$.
2. For $a = 0, 1, 2, \ldots, U$, compute

$$W_{t+1}(a) = \sum_{j=0}^{U-a-1} \alpha_j V_{t+1}(a + j) + \omega_{U-a} V_{t+1}(U).$$

3. For $s = 0, 1, 2, \ldots, U$, compute

$$V_t(s) = \max_{a \in \mathcal{A}(s)} r_t(s, a) + \beta W_{t+1}(a).$$

4. Let $t = t - 1$. If $t = 0$ stop.
5. Go to 2.

The maximization at Step 3 is done by enumerating all feasible actions ($a = 0, 1, 2, \ldots, s$), without exploiting the special structure of optimal solutions.

REFERENCES

1 T.W. Archibald, K.I.M. McKinnon and L. C. Thomas (1997). An aggregate stochastic dynamic programming model of multireservoir systems. *Water Resources Res.* 33, p. 333–340.

2 T.W. Archibald, C.S. Buchanan, K.I.M. McKinnon and L. C. Thomas (1999). Nested Benders decomposition and dynamic programming for reservoir optimization. *J. Oper. Res. Soc.* 50, pp. 468–479.

3 T.W. Archibald, C.S. Buchanan, L.C. Thomas and K.I.M. McKinnon (2001). Controlling multi-reservoir systems. *European J. Oper. Res.* 129, pp. 619–626.

4 A. Boukhtouta and B.F. Lamond (2001). La gestion des barrages en série : Variation de la hauteur de chute et aléas d'apports. *Actes de la 3e Conférence francophone de modélisation et simulation MOSIM'01*, Vol. 2, Troyes, France, 25–27 avril 2001, pp. 631–636.

5 V.C.P. Chen, D. Ruppert and C.A. Shoemaker (1999). Applying experimental design and regression splines to high-dimensional continuous-state stochastic dynamic programming. *Operations Research* 47, pp. 38–53.

6 C. De Boor (1978). *A Practical Guide to Splines.* Springer-Verlag, New York.

7 N. Drouin, A. Gautier, B.F. Lamond and P. Lang (1996). Piecewise affine approximations for the control of a one-reservoir hydroelectric system. *European J. Oper. Res.* 89, pp. 53–69.

8 B.S. Everitt (1987). *Introduction to Optimization Methods and Their Application in Statistics.* Chapman and Hall, London.

9 E. Foufoula-Georgiou and P.K. Kitanidis (1988). Gradient dynamic programming for stochastic optimal control of multidimensional water resources systems. *Water Resources Res.* 24, pp. 1345–1359.

10 J. Gessford and S. Karlin (1958). Optimal policy for hydroelectric operation. *Studies in the Mathematical Theory of Inventory and Production*, Stanford University Press, Stanford, CA, pp. 179–200.

11 D.P. Heyman and M.J. Sobel (1984). *Stochastic Models in Operations Research*, Vol. II, McGraw-Hill, New York.

12 S.A. Johnson, J.R. Stedinger, Y. Li, C.A. Shoemaker and J. A. Tejada-Guibert (1993). Numerical solution of continuous-state dynamic programs using linear and spline interpolation. *Operations Research* 41, pp. 484–500.

13 P.K. Kitanidis and E. Foufoula-Georgiou (1987). Error analysis of conventional discrete and gradient dynamic programming. *Water Resources Res.* 23, pp. 845–858.

14 B.F. Lamond (2001). On the concavity of cumulative rewards in stochastic reservoir optimization. Working paper, Faculté des sciences de l'administration, Université Laval.

15 B.F. Lamond and B. Bachar (1998). Une étude numérique de la discrétisation des apports aléatoires pour un réservoir non saisonnier. *INFOR* 36, pp. 247–260.

16 B.F. Lamond and A. Boukhtouta (2001). Water reservoir applications of Markov decision processes. In E. Feinberg and A. Schwartz, Eds., *Handbook of Markov Decision Processes: Methods and Applications*, Kluwer, pp. 519–540.

17 B.F. Lamond and A. Boukhtouta (1996). Optimizing long-term hydro-power production using Markov decision processes. *International Transactions in Operational Research* 3, pp. 223–241.

18 B.F. Lamond and P. Lang (1996). Lower bounding aggregation and direct computation for an infinite horizon one-reservoir model. *European J. Oper. Res.* 95, pp. 404–410.

19 B.F. Lamond, S.L. Monroe and M.J. Sobel (1995). A reservoir hydroelectric system: Exactly and approximately optimal policies. *European J. Oper. Res.* 81, pp. 535–542.

20 P. Lang (1994). Private communication.

21 M.A. Mariño and H.A. Loaiciga (1985). *Optimal Operation of a Multiple Reservoir System.* Technical Report, California Water Resources Center, University of California, Davis, Calif.

22 M.V.F. Pereira and L.M.V.G. Pinto (1991). Multi-stage stochastic optimization applied to energy planning. *Math. Prog.* 52, pp. 359–375.

23 C.R. Philbrick and P.K. Kitanidis (2001). Improved dynamic programming methods for optimal control of lumped-parameter stochastic systems. *Operations Research* 49, pp. 398–412.
24 M.L. Puterman (1994). *Markov Decision Processes: Discrete Stochastic Dynamic Programming*. Wiley, New York.
25 S. Soares and A.A.F.M. Carneiro (1990). Optimal operation of reservoirs for electric generation. *Proceedings of the IEEE Power Engineering Society 1990 Summer Meeting*, Minneapolis, Minnesota.
26 J.A. Tejada-Guibert, S.A. Johnson and J.R. Stedinger (1995). The value of hydrologic information in stochastic dynamic programming models of a multi-reservoir system. *Water Resources Res.* 31, pp. 2571–2579.
27 A. Turgeon (1980). Optimal operation of multireservoir power systems with stochastic inflows. *Water Resources Res.* 16, pp. 275–283.
28 A. Turgeon (1981). A decomposition method for the long term scheduling of reservoirs in series. *Water Resources Res.* 17, pp. 1565–1570.

Bernard F. Lamond est professeur titulaire au département d'opérations et systèmes de décision de l'Université Laval. Il possède un doctorat de la University of British Columbia et a été président de la Société canadienne de recherche opérationnelle (SCRO) en 2000–2001. Il a présidé le comité organisateur du congrès conjoint de la SCRO et des Journées de l'optimisation de 2001 à Québec, et il a été membre du comité organisateur de FRANCORO III. Ses travaux de recherche portent sur les méthodes d'optimisation, particulièrement en contexte stochastique, avec applications en gestion des barrages, des files d'attente et des systèmes manufacturiers flexibles.