

Publié par : Faculté des sciences de l'administration
Published by: 2325, rue de la Terrasse
Publicación de la: Pavillon Palasis-Prince, Université Laval
Québec (Québec) Canada G1V 0A6
Tél. Ph. Tel. : (418) 656-3644
Télec. Fax : (418) 656-7047

Disponible sur Internet : <http://www4.fsa.ulaval.ca/la-recherche/publications/documents-de-travail/>
Available on Internet
Disponible por Internet :

DOCUMENT DE TRAVAIL 2015-013

Approximate Stochastic Dynamic Programming for
Hydroelectric Production Planning

Luckny ZÉPHYR
Pascal LANG
Bernard F. LAMOND

Octobre 2015

Dépôt légal – Bibliothèque et Archives nationales du Québec, 2015
Bibliothèque et Archives Canada, 2015

ISBN 978-2-89524-424-0 (PDF)

Approximate Stochastic Dynamic Programming for Hydroelectric Production Planning

Luckny Zéphyr¹, Pascal Lang¹ and Bernard F. Lamond¹

¹Operations and Decision Systems Department, Université Laval, Québec, Canada

October 14, 2015

Abstract

Stochastic dynamic programming (SDP) has been widely used to derive operational policies for hydroelectric production. Several schemes have been proposed to approximate the value function with a view to tackle the curse of dimensionality. This paper presents a simplicial decomposition of the state space for such an approximation. The vertices of the resulting simplices define an irregular grid on which the value function is evaluated. Bounds on the true value function are used to refine the grid. Analytical forms for the expectation of the value function are developed for a specific information-decision process. The methodology is experimented with simulated data and on a real hydropower system.

Keywords: Approximate stochastic dynamic programming; value function; perfect correlation; reservoir trajectory; state space discretization; simplicial partitioning.

1. Introduction

Reservoir systems management partly involves decision making on water utilization, i.e. the volume of water to be released for immediate production of energy and the volume to be stored for future usage. Over the years, deterministic optimization and simulation models have been proposed to provide guidelines (Lee et al., 2008). However, deterministic plans are not always appropriate as they may fail to hedge against extreme situations (Philbrick and Kitanidis, 1999), as future natural inflows and energy demand may be significantly uncertain.

Stochastic optimization has been widely used in order to cope with such uncertainty. However, the computational burden may be considerably larger than for deterministic models since many scenarios may be needed to adequately represent uncertainty (Labadie, 2004). Several stochastic models have been developed for reservoir problems, among others two-stage linear programming with recourse (Seifi and Hipel, 2001; Lee et al., 2006; Xu et al., 2014), multistage stochastic linear programming (Kracman et al., 2006; Etkin et al., 2015), SDP (Kim et al., 2007; Zhao et al., 2011) and chance-constrained programming (Ouarda and Labadie, 2001; Jiekang et al., 2008; Zeng et al., 2013; Van Ackooij et al., 2014).

SDP was one of the first approaches used to address reservoir management. However, this scheme requires discretization of the state space. An optimization is carried out for each state. Thus, the computational burden usually increases exponentially in the dimension of the state space. This limits applications to a small number of reservoirs.

Many numerical methods have been proposed to approximate the value function such as Chebyshev polynomials (e.g. Rust, 1996), splines approximations (e.g. Johnson et al., 1993; Cervellera et al., 2007; Martinez et al., 2015), finite difference approximations (e.g. Candler, 2001). However, these approaches require appropriate smoothness of the value function. See Zéphyr et al. (2015) and Grüne and Semmler (2004) for more accounts on approximation techniques.

To obtain an acceptable precision, the value function is to be evaluated over a sufficiently large number of grid points. However, a common issue is that the computational burden induced by the number of grid points increases exponentially in the dimension of the state space (Rust, 1996; Grüne and Semmler, 2004).

This paper is concerned with an *approximate stochastic dynamic programming* scheme for reservoir management problems. In approximating the value function, we shall seek a balance between precision and computational burden. The value function will be evaluated over an irregular grid constructed by partitioning the state space into simplices. Under concavity of the estimated value function, refinement of the grid is guided by lower and upper bounds on the true value function, thus densifying the grid where the latter is more non-linear. Under additional assumptions on the natural inflows, and the information-decision process, we propose analytical forms for the expectation of the value function.

Section 2 revisits the modeling of a reservoir system and its operation, paying particular attention to the information-decision process. This provides a basis for a fundamental SDP recursion which we then seek to operationalize. Section 3 gives an overview of simplicial partitioning. Section 4 addresses a computational issue which partly contributes to the problem’s complexity, namely evaluating the expectation of the value function. This computation can be simplified by exploiting a “uni-basin” feature shared by many reservoir systems in practical settings. Section 5 deals with optimization. A generalized linear program is used to evaluate the value function. Concavity-based bounds enable us to assess the approximation error and refine the partition via simplex divisions. Finally, Section 6 reports numerical experiments and Section 7 deals with an application in the context of a real hydropower system.

2. Structure and control of a reservoir system

A hydroelectric system is generally composed of several reservoirs and power plants, each of which may be associated with a reservoir or be run-of-the river. The reservoirs’ connexions usually form an arborescence.

Management of reservoir systems is challenging, in large part, because of uncertain natural inflows. Operational decisions are concerned with water releases for energy production and spillage in case of excess water. We may assume that planned release decisions are made before observing natural inflows. The available information for these decisions then are the reservoirs’ initial levels and possibly an observed previous hydrological variable that may be

used to forecast future natural inflows. By contrast, spillage occurs after the realization of the natural inflows process. Therefore, spilled water is conditioned on both releases, and natural inflows.

Energy production is a function of water releases as well as of head effects, namely the difference between the upstream and downstream water levels and friction loss in penstock. Managers are often concerned with maintaining a high head in order to maximize efficiency. However, head effects may be difficult to capture when upstream or downstream water levels significantly depend on the rate of water release. The relationship between release rate and head size may then require careful attention (Pérez-Díaz et al., 2010).

To produce energy, water is sent through turbines that have finite mechanical capacities. Thus, bounds are imposed on the release of water. Reservoir systems also have finite capacities and are often used for other purposes such as irrigation, flood control, and recreational activities. As a consequence, bounds are also imposed on their levels.

A multi-reservoir system may be modeled by a graph where each node represents a reservoir alone, or a run-of-the-river power plant, or a reservoir-power plant group. The nodes may be numbered from 1 to p in a topological order. Figure 1 shows an example with 6 nodes. The triangles represent reservoirs, the rectangles represent power plants, the full arcs represent planned releases or spillage, and the dashed arcs represent natural inflows.

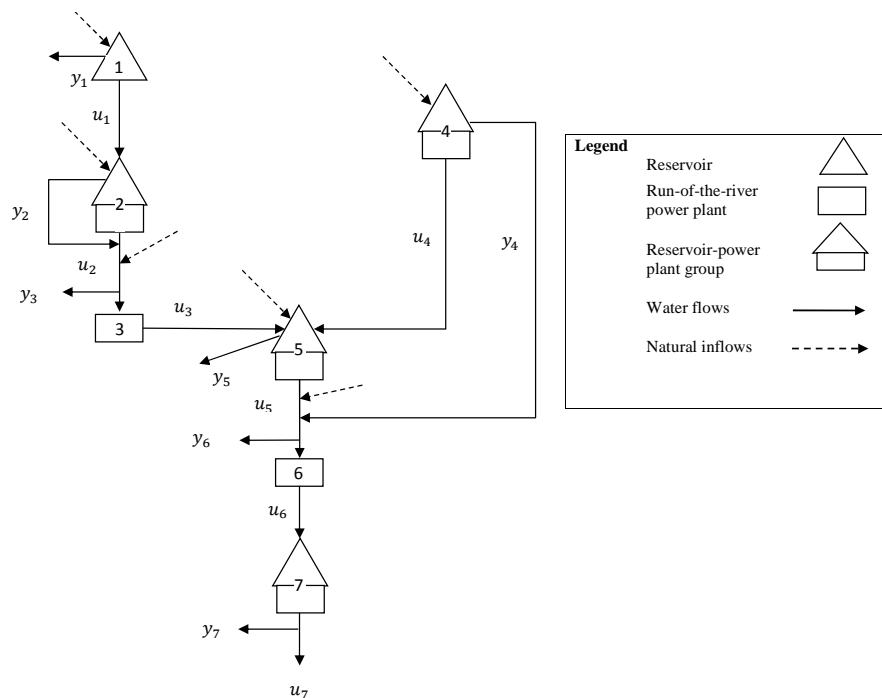


Figure 1: Example of a multi-reservoir system

We distinguish node types via the following sets of indices: (i) I_F : the set of nodes with power plant alone; (ii) I_R : the set of nodes with a reservoir (alone or reservoir-power plant group), with $|I_R| = n$; (iii) I_C : the set of nodes with power plant, with $|I_C| = m$ and (iv) $I_{RC} = I_R \cap I_C$: the set of nodes with reservoir and power plant.

This graph comprises two networks corresponding to planned releases and spillage. The topology of these networks is described via the following $p \times p$ incidence matrices:

$$B_{ij} = \begin{cases} 1 & \text{if } i = j \\ -1 & \text{if the planned releases leaving node } j \text{ arrive at node } i \\ 0 & \text{otherwise} \end{cases}$$

and

$$C_{ij} = \begin{cases} 1 & \text{if } i = j \\ -1 & \text{if the water spilled at node } j \text{ arrives at node } i \\ 0 & \text{otherwise} \end{cases}$$

For instance, in Figure 1:

$$B = \begin{pmatrix} 1 & 0 & 0 & 0 & 0 & 0 \\ -1 & 1 & 0 & 0 & 0 & 0 \\ 0 & -1 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & -1 & -1 & 1 & 0 \\ 0 & 0 & 0 & 0 & -1 & 1 \end{pmatrix} \quad \text{and} \quad C = \begin{pmatrix} 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & -1 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & -1 & 0 & 1 \end{pmatrix}$$

Optimization models are often used to determine releases and spillage at each node over a planning horizon of T periods, typically a year with weekly intervals or a month with daily time steps. For each period t , let u_{it} , $i \in I_C$, denote releases (in hm^3) and y_{it} , $i = 1, \dots, p$, spillage (in hm^3).

The state of the system is described by the reservoir levels at the end of period t , noted s_t . We note respectively \underline{s}_t and \bar{s}_t lower and upper bounds on such levels in period t . Similarly, \underline{u}_t and \bar{u}_t denote lower and upper bounds on the planned releases in each period t .

The natural inflows to the network's nodes in period t , noted Q_t , are modeled by a non-stationary process that depends stochastically on a scalar hydrological variable $\tilde{\epsilon}_t$, which forms an autoregressive stochastic process of order 1. We then have:

$$\tilde{\epsilon}_t = \alpha \tilde{\epsilon}_{t-1} + a_t$$

where a_t is a white noise process (sequence of independent and identically distributed random variables). Such hydrological variables are often used to model persistence phenomena in the natural inflows process (e.g. Mousavi et al., 2004; Kim et al., 2007; Zhao et al., 2011).

Thus, in each period t , the state of the system evolves according the standard water balance equation:

$$s_t = s_{t-1} - B_{I_R} u_t - C_{I_R} y_t + Q_{I_R t}$$

and the water conservation equation:

$$Q_{I_F t} - B_{I_F} u_t - C_{I_F} y_t = 0$$

where, for any matrix $A \in \mathbf{R}^{p \times r}$, any vector $a \in \mathbf{R}^p$ and any subset $I \subset \{1, \dots, p\}$, A_I (a_I) denotes the submatrix (the subvector) obtained by selecting the lines of A (the elements of a) indexed in I .

The last equation reflects the fact that the water that enters a run-of-the-river node in period t leaves in the same period since it is not stored.

Let $x_t = s_{t-1} - B_{I_R}u_t$ be the reservoir levels after the releases u_t and before the realization of the $\{Q_t\}$ process. For any node $i \in I_{RC}$, $\bar{f}_{it}(s_{i,t-1}, x_{it} + Q_t, \tilde{\epsilon}_{t-1})$ denotes its production function in period t . Head effects are taken into account through the reservoir levels at the beginning and the end of each period, and water release during the same period. For any node $i \in I_F$, energy production is a function of u_{it} alone since there is no stock variation; $f_{it}(u_{it})$ denotes such a function.

Optimization models generally seek to maximize the expected energy production over the entire planning horizon. This optimization can be stated as:

$$\text{Max}_{u_t, y_t} \sum_{t=1}^T E_{Q_t | \tilde{\epsilon}_{t-1}} \left[\sum_{i \in I_{RC}} \bar{f}_{it}(s_{i,t-1}, x_{it} + Q_t, \tilde{\epsilon}_{t-1}) + \sum_{i \in I_F} f_{it}(u_{it}) \right]$$

Subject to, for $1 \leq t \leq T$:

$$s_t = s_{t-1} - B_{I_R}u_t - C_{I_R}y_t + Q_{I_R t}$$

$$Q_{I_F t} - B_{I_F}u_t - C_{I_F}y_t = 0$$

$$\underline{s}_t \leq s_t \leq \bar{s}_t$$

$$\underline{u}_t \leq u_t \leq \bar{u}_t$$

$$y_t \geq 0$$

$$u_t \text{ conditioned on } s_{t-1} \text{ and } \epsilon_{t-1}$$

$$y_t \text{ and } s_t \text{ conditioned on } s_{t-1}, u_t \text{ and } Q_t$$

Let $f_{it}(u_{it}, s_{i,t-1}, \epsilon_{i,t-1}) = E_{Q_t | \tilde{\epsilon}_{t-1}} [\bar{f}_{it}(s_{i,t-1}, s_{i,t-1} - B_{I_R}u_t + Q_t, \tilde{\epsilon}_{i,t-1})]$, $\forall i \in I_{RC}$. Within the framework of SDP, we note $V_t(s_t, \epsilon_t)$ the value of the available water at the end of period t under the information state (s_t, ϵ_t) . Under the previous assumptions, for $t = T, T-1, \dots, 1$, an SDP recursion reads:

$$V_{t-1}(s_{t-1}, \epsilon_{t-1}) = \text{Max}_{u_t} \left\{ \sum_{i \in I_{RC}} f_{it}(u_{it}, s_{i,t-1}, \epsilon_{i,t-1}) + \sum_{i \in I_F} f_{it}(u_{it}) + E_{Q_t, \tilde{\epsilon}_t | \epsilon_{t-1}} \left[\text{Max}_{s_t, y_t} V_t(s_t, \tilde{\epsilon}_t) \right] \right\}$$

$$\text{s.t. } s_t = s_{t-1} - B_{I_R}u_t - C_{I_R}y_t + Q_{I_R t}$$

$$Q_{I_F t} - B_{I_F}u_t - C_{I_F}y_t = 0$$

$$\underline{s}_t \leq s_t \leq \bar{s}_t$$

$$\underline{u}_t \leq u_t \leq \bar{u}_t$$

$$y_t \geq 0$$

$$u_t \text{ conditioned on } s_{t-1} \text{ and } \epsilon_{t-1}$$

$$y_t \text{ and } s_t \text{ conditioned on } s_{t-1}, u_t \text{ and } Q_t$$

For a given period t , this optimization problem consists in (i) computing the optimal spillage y_t^* , (ii) computing the expectation, and (iii) computing the optimal planned releases u_t^* . The former recursion can be expressed by the following two steps: an “*ex post*” recourse step consisting in computing the optimal spillage policy y_t^* resulting from water overflows and an “*ex ante*” step consisting in computing the optimal planned releases u_t^* .

First, we assume that after release decisions have been made the levels of the reservoirs, x_t , are known. The natural inflows Q_t and the hydrological variable of the previous period $\tilde{\epsilon}_{t-1}$ are observed. In case of water overflows, spillage decisions are then made to maximize the value (in energy units) of the remaining water V_t . The ex post optimal decisions then are solutions to the following problem:

$$F_t(x_t, Q_t, \epsilon_t) = \text{Max}_{s_t, y_t} V_t(s_t, \epsilon_t) \quad (1)$$

$$\text{S.t. } s_t = x_t - C_{I_R} y_t + Q_{I_R t} \quad (2)$$

$$Q_{I_F t} - B_{I_F} u_t - C_{I_F} y_t = 0 \quad (3)$$

$$\underline{s}_t \leq s_t \leq \bar{s}_t \quad (4)$$

$$y_t \geq 0 \quad (5)$$

Second, the reservoir levels s_{t-1} at the beginning of period t are observed and release decisions are made to maximize immediate energy production and the expected future value of water. The ex ante optimal decisions then solve:

$$V_{t-1}(s_{t-1}, \epsilon_{t-1}) = \text{Max}_{u_t, x_t} \left\{ \sum_{i \in I_{RC}} f_{it}(u_{it}, s_{i,t-1}, \epsilon_{i,t-1}) + \sum_{i \in I_F} f_{it}(u_{it}) + E_{Q_t, \tilde{\epsilon}_t | \epsilon_{t-1}} [F_t(x_t, Q_t, \tilde{\epsilon}_t)] \right\}$$

$$\text{S.t. } x_t = s_{t-1} - B_{I_R} u_t$$

$$B_{I_F} u_t = 0$$

$$x_t \geq \underline{s}_t$$

$$\underline{u}_t \leq u_t \leq \bar{u}_t$$

The feasible domain of the optimal policy of the ex post step is very large. It is nonetheless possible to formulate as a simple rule an optimal spillage policy consistent with the notion of increasing value of water.

Proposition 1. *If all the nodes are equipped with a spillage system, function $V_{t-1}(s_{t-1}, \epsilon_{t-1})$ is non-decreasing in s_{t-1} .*

Given that V_{t-1} is non-decreasing, the optimal spillage decisions for (1-5) can be found via the following simple rule (Lamond and Lang, 2007):

Proposition 2. *Let P_i be the set of predecessors of node i (the set of nodes j such that $C_{ij} = -1, j \in \{1, \dots, p\}$) in the spillage network. Under the assumptions of proposition 1, if $x_t \geq \underline{s}_t$, there exists a unique Pareto-minimal spillage policy given recursively by:*

$$y_{it}^*(x_t + Q_t) = \max \{0, x_{it} + Q_{it} + \sum_{j \in P_i} y_{jt}^*(x_t + Q_t) - \bar{s}_{it}\}$$

with the conventions $x_{it} = 0 \forall i \in I_F$ and $\sum_{j \in \emptyset(\cdot)} = 0$. This policy is optimal for the ex post spillage problem. The resulting final stock is:

$$s_{it}^*(x_t + Q_t) = \min \{\bar{s}_{it}, x_{it} + Q_{it} + \sum_{j \in P_i} y_{jt}^*(x_t + Q_t)\} \quad (6)$$

Note that the optimal final stock satisfies the conditions $s_t^*(x_t + Q_t) \geq \underline{s}_t$. In case of water overflows (spillage), the upper bounds on the stock levels are saturated, i.e. $(\bar{s}_t - s_t^*)^T y_t^* = 0$. This simple spillage policy states that starting with the upstream nodes, we only spill the excess water in each reservoir.

Under this spillage policy, the stochastic dynamic program simplifies to:

$$V_{t-1}(s_{t-1}, \epsilon_{t-1}) = \text{Max}_{u_t, x_t} \left\{ \sum_{i \in I_{RC}} f_{it}(u_{it}, s_{i,t-1}, \epsilon_{i,t-1}) + \sum_{i \in I_F} f_{it}(u_{it}) + E_{Q_t, \tilde{\epsilon}_t | \epsilon_{t-1}} [V_t(s_t^*(x_t + Q_{I_{Rt}}), \tilde{\epsilon}_t)] \right\} \quad (7)$$

$$\text{S.t. } x_t = s_{t-1} - B_{I_R} u_t \quad (8)$$

$$B_{I_F} u_t = 0 \quad (9)$$

$$x_t \geq \underline{s}_t \quad (10)$$

$$\underline{u}_t \leq u_t \leq \bar{u}_t \quad (11)$$

3. Simplicial partitioning of the state space

Problem (7-11) cannot be solved for all feasible values $(s_{t-1}, \epsilon_{t-1})$, so most models require discretizing the state space. Function V_{t-1} is evaluated at the resulting grid points. In order to spare evaluations, we propose a simplicial approximation of V_{t-1} noted \hat{V}_{t-1} . The final stocks space in period $t-1$ is taken to be a hyperrectangle $H_{t-1} = \{s_{t-1} \in \mathbf{R}^n | \underline{s}_{t-1} \leq s_{t-1} \leq \bar{s}_{t-1}\}$. This hyperrectangle is partitioned into n -dimensional simplices whose vertices are the grid points where function V_{t-1} is evaluated.

We present an overview of common operations on simplices, namely (i) testing whether a point belongs to a simplex, (ii) dividing a simplex, (iii) finding a smallest simplex containing a given point, and (iv) moving across simplices.

A simplex $\Sigma(S) \subset H_{t-1}$ is the convex hull of its vertices s^1, s^2, \dots, s^{n+1} :

$$\Sigma(S) = \left\{ s \in \mathbf{R}^n \mid \begin{pmatrix} s \\ 1 \end{pmatrix} = \begin{pmatrix} S \\ e^T \end{pmatrix} \lambda, \lambda \geq 0 \right\},$$

where $S = (s^1, s^2, \dots, s^{n+1})$ is a $n \times (n+1)$ square matrix, and $e^T = (1, \dots, 1)$. Since these vertices are affinely independent, $\begin{pmatrix} S \\ e^T \end{pmatrix}$ is of full rank.

The minimal number of simplices into which H_{t-1} can be divided is $n!$. An algorithm developed independently by Coxeter, Freudenthal and Kuhn achieves this minimum (see Moore, 1992). Except for the initial choice of an opposite pair of points, this algorithm leaves no room for discretionary choices. Zéphyr et al. (2015) present a method that allows more degrees of freedom and a better coverage of the stocks space. This approach decomposes any hypercube into $2^n n!$ simplices. Hence, the counterpart is a higher computational complexity.

Suppose H_{t-1} is partitioned into initial simplices. We may want to obtain a finer partition. This can be achieved by dividing existing simplices one at a time.

For any simplex $\Sigma(S)$, the *division point* s^* will be located on some k -dimensional face F of the hypercube, yielding a division into $k+1$ subsimplices. To maintain the overall

integrity of the partition, any other simplex $\Sigma(S')$ such that $\Sigma(S') \cap \Sigma(S) = F$ will similarly be divided. The possible resulting cases are illustrated in Figure 2. In case (a), since s^* is an interior point of $\Sigma(S)$, the simplex is divided into three subsimplices. In case (b), simplex $\Sigma(S)$ is divided into two subsimplices given that the division point is located on the boundary of the state space. In case (c), both $\Sigma(S)$ and $\Sigma(S')$ are divided into two subsimplices with s^* a common vertex, since this latter is located on the edge $\{x, w\}$ common to simplices $\Sigma(S)$ and $\Sigma(S')$.




Figure 2: Illustration of the division of simplices (Zéphyr et al., 2015)

Since some simplices are divided into smaller simplices, the overall partition of the state space has a nested structure, which can be represented as an arborescence. The nodes of this arborescence are simplices. Its arcs represent division. A simplex which is not divided (a leaf in the arborescence) will be called *active*. The set of active simplices forms the finest available partition of the state space, the one of interest to us. We nonetheless keep track of the arborescence to be able to locate a point efficiently.

Any point $s \in H_{t-1}$ belongs to at least one active simplex. From the definition above, we see that s belongs to simplex $\Sigma(S)$ if and only if $\lambda = \begin{pmatrix} S \\ e^T \end{pmatrix}^{-1} \begin{pmatrix} s \\ 1 \end{pmatrix} \geq 0$. To locate the smallest simplex containing s , we may apply a depth-first search procedure in the arborescence.

Our last operation is *simplex traversal*. Starting from a point $s \in \Sigma(S)$, our task is to move as far as possible in a given direction $\delta \in \mathbf{R}_+^n \setminus \{0\}$ without leaving the simplex. We say that simplex $\Sigma(S)$ *properly contains* the point s relatively to the direction δ , which we note $s \in_\delta \Sigma(S)$, if a positive step is possible, that is if for any sufficiently small $\varepsilon > 0$, $s + \varepsilon\delta \in \Sigma(S)$. It follows from the definition that $s \in_\delta \Sigma(S)$ if and only if $\forall 1 \leq j \leq n+1, \lambda_j = 0 \Rightarrow \sigma_j \geq 0$, where $s \in \Sigma(S), \lambda = \begin{pmatrix} S \\ e^T \end{pmatrix}^{-1} \begin{pmatrix} s \\ 1 \end{pmatrix}$ and $\sigma = \begin{pmatrix} S \\ e^T \end{pmatrix}^{-1} \begin{pmatrix} \delta \\ 0 \end{pmatrix}$. This property holds since by definition, $s \in_\delta \Sigma(S)$ if and only if $\forall \varepsilon > 0, \lambda + \varepsilon\sigma \geq 0$.

The maximal displacement will get us to a point $s' \neq s$ located on the boundary of $\Sigma(S)$ and such that $s' \notin_\delta \Sigma(S)$. Formally, we seek the step length

$$\theta^* = \max\{\theta \in \mathbf{R}_+ | s + \theta\delta \in \Sigma(S)\}.$$

It is readily verified that $\theta^* = \min\left\{-\frac{\lambda_j}{\sigma_j} | \sigma_j < 0\right\}, 1 \leq j \leq n+1$.

4. Computing the expectation of the value function

The SDP recursion (7-11) has two major sources of complexity, namely (i) taking expectation and (ii) optimizing. The object of this section is the computation of the ex-

pectation $E_{Q_t, \tilde{\epsilon}_t | \epsilon_{t-1}} [V_t(s_t^*(x_t + Q_{I_{Rt}}), \tilde{\epsilon}_t)]$ for all discrete states (x_t, ϵ_{t-1}) . Without loss of generality, we assume that the hydrological variable evolves according a known transition probability distribution $P(\tilde{\epsilon}_t = \epsilon_t | \tilde{\epsilon}_{t-1} = \epsilon_{t-1})$. This is akin to consider that the hydrological variable is Markovian, which is a common approach in reservoir management (see Tejada-Guibert et al., 1995; Faber and Stedinger, 2001; Goor et al., 2010). We then have:

$$\begin{aligned} E_{Q_t, \tilde{\epsilon}_t | \epsilon_{t-1}} [V_t(s_t^*(x_t + Q_{I_{Rt}}), \tilde{\epsilon}_t)] &= E_{\tilde{\epsilon}_t | \epsilon_{t-1}} E_{Q_t | \epsilon_{t-1}} [V_t(s_t^*(x_t + Q_{I_{Rt}}), \tilde{\epsilon}_t)] \\ &= \sum_{\tilde{\epsilon}_t} P(\tilde{\epsilon}_t = \epsilon_t | \tilde{\epsilon}_{t-1} = \epsilon_{t-1}) E_{Q_t | \epsilon_{t-1}} [V_t(s_t^*(x_t + Q_{I_{Rt}}), \epsilon_t)]. \end{aligned}$$

Taking the expectation then amounts to computing:

$$J_t(x_t, \epsilon_{t-1}) = E_{Q_t | \epsilon_{t-1}} [V_t(s_t^*(x_t + Q_{I_{Rt}}), \epsilon_t)]$$

for a discrete set of values (x_t, ϵ_{t-1}) . This may again be simplified if the natural inflows to the reservoirs are spatially perfectly correlated. This “uni-bassin” hypothesis is approximately verified in many real situations, where the reservoirs are fed by a common hydrological basin. From a practical point of view, the dimension of the natural inflows’ support decreases from n to 1. This assumption is often used by hydrologists in case of planning horizon with monthly or weekly time step (see Turgeon, 2007). For instance this hypothesis will be used later in the case study. Our approach will be compared to a “classical” dynamic programming model developed by the company that also uses the same assumption to compute the expectation of the value function.

In the case of perfect correlation, we may consider a scalar variate $\tilde{\tau}_t \geq 0$ such that

$$Q_t = q_t \tilde{\tau}_t,$$

where $q_t \in \mathbf{R}^n$ represents the relative contributions of natural inflows to the reservoirs, and $\tilde{\tau}_t \geq 0$.

Numerical integration might be used to compute function $J_t(x_t, \epsilon_{t-1})$ whenever the probability distribution of $\tilde{\tau}_t$ is known. However, the final stocks space being partitioned into simplices, we are able to obtain analytical forms for the expectation. We will first analyze the trajectory of the final stocks in the partitioned space. Second, we will address the evaluation of function \hat{V}_t over this trajectory.

The vertices of the simplices form an irregular grid over which \hat{V}_t has been evaluated. For any simplex $\Sigma(S)$, let $z_i = \hat{V}_t(s^i, \epsilon_t)$, $1 \leq i \leq n + 1$. For any other point $s \in \Sigma(S)$, \hat{V}_t is interpolated over $\Sigma(S)$. The interpolation is uniquely determined as the affine form:

$$\hat{V}_t(s, \epsilon_t) = z^T \begin{pmatrix} S \\ e^T \end{pmatrix}^{-1} \begin{pmatrix} s \\ 1 \end{pmatrix}, \text{ where } z^T = (z_1, \dots, z_{n+1}).$$

Under the uni-bassin hypothesis, the optimal spillage and final stocks trajectories are unidimensional paths in \mathbf{R}^n . From (6) we see that the final stocks’ trajectory $\hat{s}_t(\tau) \triangleq s_t^*(x_t + q_t \tau)$ is linear until attaining a reservoir upper bound. We call *corner* a point at which one new upper bound becomes saturated. At each new corner, the trajectory direction changes, the new bound remaining saturated until the end of the trajectory, \bar{s}_t . The trajectory $\hat{s}_t(\tau)$ is illustrated in Figure 3.

In addition to corners, the trajectory $\hat{s}_t(\tau)$ crosses the simplices at intermediate *traversal points* (see Figure 3). We may define a set of *nodes* $\Theta = \{\theta_i | i = 1, \dots, w\}$ of the parameter $\tilde{\tau}_t$ such that $\hat{s}_t(\theta_i)$ is a corner or a traversal point.




Figure 3: Example of final stocks' trajectory

Function \hat{V}_t is evaluated at each node $\theta_i, i = 1, \dots, w$. Let $h_t(\theta_i, \epsilon_t) = \hat{V}_t(\hat{s}_t(\theta_i), \epsilon_t)$. Function h_t 's slopes change at the nodes in Θ . Thus, h_t is piecewise affine. For each segment (θ_{i-1}, θ_i) , the slope is constant and is given by:

$$\alpha_i = \frac{\omega_i - \omega_{i-1}}{\theta_i - \theta_{i-1}}, \quad i = 2, \dots, w + 1,$$

where $\omega_i = \hat{V}_t(\hat{s}_t(\theta_i), \epsilon_t)$. In particular, $\alpha_{w+1} = 0$. Function h_t is constant for $\tau \geq \theta_w$, the stocks of water attaining their maximal levels. We have

$$h_t(\theta_i, \epsilon_t) = \begin{cases} \omega_i & i = 2, \dots, w \\ \omega_w & i = w + 1. \end{cases}$$

Then, function h_t is of the form:

$$h_t(\tau, \epsilon_t) = \begin{cases} \omega_{i-1} + \alpha_i(\tau - \theta_{i-1}) & \theta_{i-1} \leq \tau \leq \theta_i, \quad i = 2, \dots, w, \\ \omega_w & \tau \geq \theta_w. \end{cases}$$

4.1 Analytical form

Function $J_t(x_t, \epsilon_{t-1})$ is calculated over each segment $(\theta_{i-1}, \theta_i), i = 2, \dots, w + 1$. Let $F_t(\tau)$ be the cumulative distribution of $\tilde{\tau}_t$ conditioned on the value ϵ_{t-1} . Over each interval $[\theta_{i-1}, \theta_i]$, the conditional expectation is computed as $\int_{\theta_{i-1}}^{\theta_i} h_t(\tau, \epsilon_t) dF_t(\tau)$. Using the linearity property of the expectation, we then get

$$J_t(x_t, \epsilon_{t-1}) = \sum_{i=2}^{w+1} \int_{\theta_{i-1}}^{\theta_i} h_t(\tau, \epsilon_t) dF_t(\tau),$$

where

$$\begin{aligned}
\sum_{i=2}^{w+1} \int_{\theta_{i-1}}^{\theta_i} h_t(\tau, \epsilon_t) dF_t(\tau) &= \sum_{i=2}^{w+1} \int_{\theta_{i-1}}^{\theta_i} (\omega_{i-1} + \alpha_i(\tau - \theta_{i-1})) dF_t(\tau) \\
&= \sum_{i=2}^{w+1} \alpha_i E[\tilde{\tau}_t | \epsilon_{t-1}]_{\theta_{i-1}}^{\theta_i} + \sum_{i=2}^{w+1} \gamma_{i-1} [F_t(\theta_i) - F_t(\theta_{i-1})] \\
&= \sum_{i=2}^{w+1} \alpha_i E[\tilde{\tau}_t | \epsilon_{t-1}]_{\theta_{i-1}}^{\theta_i} + \sum_{i=2}^w \gamma_{i-1} [F_t(\theta_i) - F_t(\theta_{i-1})] + \omega_w [1 - F_t(\theta_w)],
\end{aligned}$$

with $\gamma_i = \omega_i - \alpha_{i+1}\theta_i$ and $\theta_{w+1} = +\infty$.

We provide more details for the specific cases of normal and log-normal distributions.

4.2 Case of the truncated normal distribution

Assume $\tilde{\tau}_t$ is a normal random variable with mean $\mu(\epsilon_{t-1})$ and standard deviation $\sigma(\epsilon_{t-1})$, we have:

$$\begin{aligned}
F_t(\tau) &= P(\tilde{\tau}_t \leq \tau) = \Phi\left(\frac{\tau - \mu(\epsilon_{t-1})}{\sigma(\epsilon_{t-1})}\right) \\
f_t(\tau) &= F_t'(\tau) = \frac{1}{\sigma(\epsilon_{t-1})} \phi\left(\frac{\tau - \mu(\epsilon_{t-1})}{\sigma(\epsilon_{t-1})}\right),
\end{aligned}$$

where $\phi(v) = \frac{1}{\sqrt{2\pi}} e^{-\frac{v^2}{2}}$ and $\Phi(v) = \int_{-\infty}^v \phi(t) dt$ are respectively the density function and the cumulative distribution of the standard normal variable. Let $G(\tau) = \mu(\epsilon_{t-1})F_t(\tau) - (\sigma(\epsilon_{t-1}))^2 f_t(\tau)$. The standard normal distribution has the property: $\phi'(v) = -v\phi(v)$. It follows that

$$\int \tau f_t(\tau) d\tau = G(\tau).$$

Therefore

$$E[\tilde{\tau}_t | \epsilon_{t-1}]_{\theta_{i-1}}^{\theta_i} = G(\theta_i) - G(\theta_{i-1})$$

and

$$J_t(x_t, \epsilon_{t-1}) = \frac{1}{1 - F_t(0)} \sum_{i=2}^{w+1} \{\alpha_i [G(\theta_i) - G(\theta_{i-1})] + \gamma_{i-1} [F_t(\theta_i) - F_t(\theta_{i-1})]\}.$$

4.3 Case of the log-normal distribution

$\tilde{\tau}_t$ is a log-normal variable with parameters $\mu(\epsilon_{t-1})$ and $\sigma(\epsilon_{t-1})$ if $\ln \tilde{\tau}_t$ is a normal variable with mean $\mu(\epsilon_{t-1})$ and variance $(\sigma(\epsilon_{t-1}))^2$. The density function of $\tilde{\tau}_t$ is: $f_t(\tau; \mu(\epsilon_{t-1}), \sigma(\epsilon_{t-1})) = \frac{1}{\sqrt{2\pi}\sigma(\epsilon_{t-1})\tau} e^{-\frac{(\ln \tau - \mu(\epsilon_{t-1}))^2}{2(\sigma(\epsilon_{t-1}))^2}}$; its mean is $\nu = E[\tilde{\tau}_t] = e^{\mu(\epsilon_{t-1}) + \frac{(\sigma(\epsilon_{t-1}))^2}{2}}$, and its k^{th} moment about the origin $\nu_k = E[\tilde{\tau}_t^k] = e^{k\mu(\epsilon_{t-1}) + \frac{k^2(\sigma(\epsilon_{t-1}))^2}{2}}$.

As previously, let Φ be the cumulative distribution of the standard normal variable. Let $\rho(v) = \frac{\ln v - \mu(\epsilon_{t-1}) - \frac{(\sigma(\epsilon_{t-1}))^2}{2}}{\sigma(\epsilon_{t-1})}$. We conclude with the following result:

Proposition 3. For $0 \leq a < b$, $E[\tilde{\tau}_t | \epsilon_{t-1}]_a^b = \nu [\Phi(\rho(b)) - \Phi(\rho(a))]$.

Let us observe that the cumulative distribution of $\tilde{\tau}_t$ is : $F_t(\tau) = \Phi(\rho(\tau) + \sigma(\epsilon_{t-1}))$. Then

$$J_t(x_t, \epsilon_{t-1}) = \sum_{i=2}^{w+1} \{ \nu \alpha_i [\Phi(\rho(\theta_i)) - \Phi(\rho(\theta_{i-1}))] + \gamma_{i-1} [\Phi(\rho(\theta_i) + \sigma(\epsilon_{t-1})) - \Phi(\rho(\theta_{i-1}) + \sigma(\epsilon_{t-1}))] \}.$$

5. Approximating the value function

Once the expectation of function V_t is computed, problem (7-11) will be solved to find the optimal release policy u_t^* for a discrete set of values $(s_{t-1}, \epsilon_{t-1})$. This will require evaluating the production functions. These are sometimes highly non-concave. Hence problem (7-11) may be non-convex. Finding global optima may then not be guaranteed.

Non-concavity may have several causes. The power delivered by a turbine varies non-linearly with the flow rate. A minimum rate is needed for setting the turbine into motion. The turbine achieves its maximum efficiency at a given flow rate. Beyond an upper threshold, turbulence may actually decrease the turbine's efficiency. Second, the total energy produced depends on the number of turbines in operation. Depending on production goals, plants use established *unit commitment* policies. Last but not least comes head effect. Each turbine's power depends on the head size. An additional non-linear effect occurs when the overall release rate affects the head size. Figure 4 depicts a hypothetical energy production function together with a concave piecewise affine approximation. Such approximation is particularly acceptable in the case where the operational policies derived from a mid term model are input to a short term model aimed at deriving the daily or hourly operational policies. This is the case, for instance, at the Rio Tinto Alcan company, our case study.




Figure 4: Example of a plant's energy production function and its approximation

5.1 Optimization via generalized linear programming

We now turn to the solution of problem (7-11). Since its feasible domain is a convex polyhedron, this problem can in principle be solved to optimality if the expected value function J_t is concave in x_t and the production functions f_{it} are concave in the release rates u_{it} . However, even in this convex case, not all constrained optimization methods are equally appropriate. The mere fact that the feasible domain is a polyhedron implies that neither the optimal

policy nor the value function are smooth. Inner *generalized linear programming* (GLP, e.g. Shapiro, 1979) performs interpolations of functions over a discrete sample of points. If the original problem is convex, well-conceived densification of the sample will ultimately converge to the original problem. In non-convex cases (e.g. here if the production functions are not concave), GLP performs a convex approximation (a “convexification”) of the original problem. For instance, the master program in column generation is in GLP format.

Since the feasible domain is a polyhedron, GLP is an appropriate approximation for problem (7-11). To implement GLP, we shall need (i) a grid of states $\{\hat{x}_{jt}|j \in \Upsilon_t\}$ over which function $\hat{J}_t(x_t, \epsilon_{t-1}) = E_{Q_t, \tilde{\epsilon}_t|\epsilon_{t-1}} [\hat{V}_t(s_t^*(x_t + Q_{I_{Rt}}), \tilde{\epsilon}_t)]$ has been evaluated and, (ii) for each power plant i , a sample of points $\{\hat{u}_{ijt}|j \in \Gamma_{it}\}$ where the production function is assessed. Furthermore, we distinguish power plants associated with a reservoir ($i \in I_{RC}$) from run-of-the river ones ($i \in I_F$). Given these fixed samples, the GLP seeks a “best” interpolation :

$$\hat{V}_{t-1}(s_{t-1}, \epsilon_{t-1}) = \text{Max}_{u_t, x_t, \lambda, \mu} \left\{ \sum_{i \in I_{RC}} \sum_{j \in \Gamma_{it}^1} f_{it}(\hat{u}_{ijt}^1, s_{i,t-1}) \lambda_{ij}^1 + \sum_{i \in I_F} \sum_{j \in \Gamma_{it}^2} f_{it}(\hat{u}_{ijt}^2) \lambda_{ij}^2 + \sum_{j \in \Upsilon_t} \hat{J}_t(\hat{x}_{jt}, \epsilon_{t-1}) \mu_j \right\} \quad (12)$$

$$\text{S.t. } x_t = s_{t-1} - B_{I_R} u_t \quad (13)$$

$$B_{I_F} u_t = 0 \quad (14)$$

$$x_t = \sum_{j \in \kappa_t} \hat{x}_{jt} \mu_j \quad (15)$$

$$\sum_{j \in \kappa_t} \mu_j = 1 \quad (16)$$

$$x_t \geq \underline{s}_t \quad (17)$$

$$\underline{u}_t \leq u_t \leq \bar{u}_t \quad (18)$$

$$u_{it} = \sum_{j \in J_{it}^1} \hat{u}_{ijt}^1 \lambda_{ij}^1 \quad i \in I_{RC} \quad (19)$$

$$\sum_{j \in J_{it}^1} \lambda_{ij}^1 = 1 \quad i \in I_{RC} \quad (20)$$

$$u_{it} = \sum_{j \in J_{it}^2} \hat{u}_{ijt}^2 \lambda_{ij}^2 \quad i \in I_F \quad (21)$$

$$\sum_{j \in J_{it}^2} \lambda_{ij}^2 = 1 \quad i \in I_F \quad (22)$$

$$\lambda, \mu \geq 0 \quad (23)$$

where $u_t = (u_t^1, u_t^2)^T$, $\hat{u}_t = (\hat{u}_t^1, \hat{u}_t^2)^T$ and $\lambda = (\lambda^1, \lambda^2)^T$.

The following simple property will be critical to refining the state space grid, as later presented.

Proposition 4. \hat{V}_{t-1} is concave in $s_{t-1} \forall \epsilon_{t-1}$.

On which side does the GLP approximation err? This cannot be predicted in general. However, as a basis for comparison, we have a definite answer in the convex case.

Proposition 5. *If the production functions f_i are concave in u_{it} , and if $\hat{V}_t \leq V_t$, then GLP (12-23) underestimates the actual energy production in period t , and $\hat{V}_{t-1} \leq V_{t-1}$.*

Finally, it must be noted that, whereas the optimal policy $u_t^*(s_{t-1}, \epsilon_{t-1})$ is exactly computed at a finite set of grid points, it can be approximately extended over the state space's continuum. Indeed, assume that a particular initial stock s_{t-1} is not a grid point. An approximate policy $\hat{u}_t(s_{t-1}, \epsilon_{t-1})$ can be recovered as follows.

Suppose s_{t-1} belongs to the active simplex $\Sigma(S)$, with $S = (s_{t-1}^1, \dots, s_{t-1}^{n+1})$. For each vertex of this simplex $s_{t-1}^i, 1 \leq i \leq n+1$, we have an exact evaluation of function $\hat{V}_{t-1}(s_{t-1}^i, \epsilon_{t-1})$ as well as the corresponding optimal policy $u_t^{i*}(s_{t-1}^i, \epsilon_{t-1})$. The approximated policy associated with the state $(s_{t-1}, \epsilon_{t-1})$ may be obtained by interpolation over the simplex vertices:

$\hat{u}_t(s_{t-1}, \epsilon_{t-1}) = \sum_{i=1}^{n+1} \lambda_i u_t^{i*}(s_{t-1}^i, \epsilon_{t-1})$, where $\lambda = \begin{pmatrix} S \\ e^T \end{pmatrix}^{-1} \begin{pmatrix} s_{t-1} \\ 1 \end{pmatrix}$. This interpolated policy is feasible since (13-23) is a convex domain.

5.2 Refinement of the state space grid

Suppose we have an initial set of grid points over which function \hat{V}_{t-1} has been evaluated. These points are the vertices of simplices. We may want to densify the grid to improve the approximation. Additional grid points may be obtained by iteratively dividing the existing simplices (as explained in Section 3) until meeting a desired level of precision.

Suppose simplex $\Sigma(S)$ is to be divided. We have evaluations of function \hat{V}_{t-1} at the vertices s_{t-1}^i of $\Sigma(S)$: $z_i = V_{t-1}(s_{t-1}^i, \epsilon_{t-1}), 1 \leq i \leq n+1$, and a subgradient of \hat{V}_{t-1} at each vertex: $g^i \in \partial \hat{V}_{t-1}(s_{t-1}^i, \epsilon_{t-1}), 1 \leq i \leq n+1$, chosen for instance as the Lagrange multipliers associated with Eq. (13). The division of $\Sigma(S)$ will contribute one additional grid point. To improve the approximation, the division point will be chosen where the imprecision over simplex $\Sigma(S)$ is maximal (e.g. Zéphyr et al., 2015).

To estimate the imprecision, we construct lower and upper bounds \underline{V}_{t-1} and \bar{V}_{t-1} such that $\underline{V}_{t-1}(s_{t-1}, \epsilon_{t-1}) \leq V_{t-1}(s_{t-1}, \epsilon_{t-1}) \leq \bar{V}_{t-1}(s_{t-1}, \epsilon_{t-1}) \forall s_{t-1} \in \Sigma(S)$ and $\underline{V}_{t-1}(s_{t-1}^i, \epsilon_{t-1}) = z_i = \bar{V}_{t-1}(s_{t-1}^i, \epsilon_{t-1}), 1 \leq i \leq n+1$. At any point $s_{t-1} \in \Sigma(S)$, the imprecision is then measured by: $\bar{V}_{t-1}(s_{t-1}, \epsilon_{t-1}) - \underline{V}_{t-1}(s_{t-1}, \epsilon_{t-1})$. Further, under concavity assumption on V_{t-1} , a lower bound \underline{V}_{t-1} is obtained by interpolation over the simplex vertices. Similarly, by concavity of V_{t-1} , $\bar{V}_{t-1}(s_{t-1}, \epsilon_{t-1}) = \underset{1 \leq i \leq n+1}{Min} \{z_i + g^i(s_{t-1} - s_{t-1}^i)\}, s_{t-1} \in \Sigma(S)$, is an upper bound on V_{t-1} .

The approximation gap Δ together with an associated division point s^0 then solve the linear program:

$$\Delta = \underset{\mu, s_{t-1}, \lambda}{Max} \mu - z^T \lambda \quad (24)$$

$$\text{S.t. } \mu \leq z^i + g^i(s_{t-1} - s_{t-1}^i), \quad 1 \leq i \leq n+1 \quad (25)$$

$$s_{t-1} = S\lambda \quad (26)$$

$$e^T \lambda = 1 \quad (27)$$

$$\lambda \geq 0 \quad (28)$$

Figure 5 provides geometrical illustrations of the computation of s_{t-1}^0 together with Δ .




Figure 5: Illustration of the computation of Δ (Zéphyr et al., 2015)

In summary, for each active simplex, problem (24-28) is solved to compute the maximal imprecision Δ together with a corresponding division point s_{t-1}^0 . We choose an active simplex $\Sigma(S_j)$ with maximal imprecision Δ_j . The division point s_{t-1}^j of $\Sigma(S_j)$ contributes one additional grid point. Simplex $\Sigma(S_j)$ is divided at s_{t-1}^j and becomes *inactive*. The resulting subsimplices are added to the list of active simplices. This refinement process is repeated until meeting a prescribed criterion such as a threshold on Δ , a fixed number of iterations or a threshold on the computation time.

6. Numerical experiments

We analyzed the trade-off between solution time and accuracy of the proposed method for 45 randomly generated problem instances. We used a one-year planning horizon with weekly time steps. In each week, the algorithm for the decomposition of simplices presented in Section 3 was used to construct the state space grid points over which the value function has been evaluated by solving, for each grid point, the generalized linear program (12-23), but without hydrological variable. Several thresholds on the relative imprecision $\eta = \frac{\Delta_k}{\Delta_1}$, where k is a number of iteration, were used, varying between 0.8 and 0.01, to refine the stock grids. The model parameters were generated by Monte Carlo simulations. The numerical values of these parameters, reported in Table 1, were drawn from uniform and independent distributions with support of the form $[a, b]$.

Table 1: Bounds on the model parameters

Parameter	a	b
$\underline{v}_{it}, i \in I_R$	150	600
$\bar{v}_{it}, i \in I_R$	800	7000
$\underline{u}_{it}, i \in I_R$	0	0
$\bar{u}_{it}, i \in I_R$	$0.05\bar{v}_{it}$	$1.5\bar{v}_{it}$
$\bar{u}_{it}, i \in I_F$	950	3500

Table 2 shows the overall simulation plan. The instance problems were solved using the 5-reservoir configuration illustrated in Figure 1 (see Section 2).

Table 2: Number of replications per value of η

η	# of replications
0.8	10
0.6	10
0.4	10
0.2	9
0.1	5
0.01	1
Total number of replications	45

The production functions were supposed to obey the theoretical forms:

$$f_i(u_{it}) = \beta_i[(u_{it} + \gamma_i)^{\alpha_i} - \gamma_i^{\alpha_i}], i \in I_F, 0 \leq \alpha_i \leq 1, \beta_i > 0, \gamma_i \geq 0$$

and

$$f_i(u_{it}, v_{it}) = \beta_i[(u_{it} + \gamma_i)^{\alpha_i} - \gamma_i^{\alpha_i} + (u_{it}v_{it})^{\theta_i}], i \in I_{RC}, 0 \leq \alpha_i \leq 1, \beta_i > 0, \gamma_i \geq 0, 0 \leq \theta_i \leq 1.$$

In each period, the expectation of the value function was computed as earlier presented in Section 4; we assumed that the natural inflows obeyed log-normal distributions and used a stationary vector of relative contributions to the reservoirs $q = [0.20 \ 0.10 \ 0.30 \ 0.25 \ 0.15]$.

The results of these experiments are discussed in Section 6.1.

6.1 Results of the experiments

For each simulation run, we measured the overall execution time in seconds to construct the 52 value functions. Afterwards, we computed the overall average execution time for each value of the imprecision (η). Table 3 reports descriptive statistics for the execution time.

These results provide significant insight on the trade-off between computational burden and accuracy. Indeed, with an initial imprecision $\eta = 0.8$, the solution time varied between 51.210 and 257.759 s , with an overall average of 120.205 s . By reducing the tolerance of the initial error by 10% ($\eta = 0.7$), the average solution time was multiplied by about 1.438. A reduction of 40% of the initial imprecision involved a multiplication by about 7 of the total execution time. By reducing the initial imprecision by 70%, on average, the computational burden was multiplied by 42 (about 5 000 s).

Figure 6 shows a semi-log plot of the overall average execution time as a function of the precision ($1 - \eta$). One may observe a piecewise linear relationship, which might suggest that the computation time increases exponentially with the precision, but at different local increasing rates.

The trade-off between solution time and accuracy may also be analyzed by the average number of grid points per period as exemplified in Table 4. With an initial relative error $\eta = 0.8$, overall the 45 simulation runs, the average grid points per week ranged from 55 to 176. Decreasing the initial error to 0.7, the average size of the grids per week ranged between 72 and 297. With $\eta = 0.4$, the minimal average grid points per week more than doubled and the maximal value was multiplied by about 4. Setting the relative error to 10%, the minimal

Table 3: Descriptive statistics for the execution time (in seconds) of the simulations

η	Mean	Minimum	Maximum
0.8	120.205	51.210	257.759
0.7	172.911	74.166	362.775
0.6	440.639	70.088	1 161.306
0.4	842.463	137.58	1 711.899
0.2	1 498.090	982.471	2 073.618
0.1	4 995.312	3 292.068	8 089.491
0.01	14 622.483	14 622.483	14 622.483




Figure 6: Semi-log plot of the average computation time as a function of the precision

and maximal average size of the grid points per week varied between 301 and 1153, i.e. a multiplication by about five and seven, respectively.

Table 4: Descriptive statistics for the average number of grid points per week

η	Mean	Minimum	Maximum
0.8	103.229	55.481	176.115
0.7	156.431	72.346	297.019
0.6	278.871	95.615	430.808
0.4	409.002	123.115	714.673
0.2	412.546	266.096	568.346
0.1	656.035	301.192	1151.750
0.01	675.115	675.115	675.115

7. Case study: Rio Tinto Alcan

Rio Tinto Alcan (RTA) is an international company that produces and sells aluminum on the international market. RTA has power plants located in Saguenay, a city of the province of Quebec (Canada). The hydroelectric system is composed of six power plants, of which three are run-of-the river. The installed capacity is 3100 MW. However, because of insufficient natural inflows at some times during the year, the average production is about 2100 MW. The system feeds four aluminum plants for an overall load of about 2300 MW. To compensate for load shedding, RTA has several supplying contracts with Hydro-Quebec, of which a 3 TWh contract available from December 1st to November 30th.

Section 7.1 describes the system under study. Sections 7.2 and 7.3 report the experimental framework adopted to solve RTA’s problem as well as the results of these experiments, respectively.

7.1 Description of the system under study

For operational reasons, RTA is interested in studies on the south portion of the system as illustrated in Figure 7. The full arcs represent turbinéd flows, the dash line are spillage and the dotted lines illustrate natural inflows. This subsystem is undersized compared to the volume of natural inflows, and is more difficult to manage compared with the upstream portion of the system. Furthermore, the upstream turbines have low capacity. The downstream subsystem, of interest to us, comprises two reservoirs: Chute-Du-Diable and Lac Saint-Jean associated with the power plants Chute-Du-Diable and Isle-Maligne, respectively, and two run-of-the-river power plants: Chute-Savanne and Shipshaw.




Figure 7: Representation of the RTA’s downstream subsystem

In the following, we will use the same notation as earlier.

Power generation is a function of the turbine’s feeds and the level of the reservoirs for the power plants Chute-Du-Diable (CCD, node 1) and Isle-Maligne (CIM, node 3) (associated with the similarly labeled reservoirs). Though Chute-Savanne (CCS, node 2) is a run-of-the power plant, besides the turbine’s feeds, its power generation also is a function of the level of the reservoir Lac Saint-Jean (LSJ, node 3), since this latter affects the level of its tailbay. Lastly, as concerned the power plant Shipshaw (CSH) (node 4), power generation only depends on the water release rate.

The maximal water release rate in the power plants chute-Du-Diable and Isle-Maligne depends on the level of reservoir in the beginning of each period, whereas it is a constant for the power plants Chute-Savanne and Shipshaw. $\tilde{B}_i(s_{it}), i = 1, 3$, denote these maximal release rate functions. Furthermore, in the four power plants, any excess of water is spilled and causes loss of power by increasing the level of the tailbay. $\tilde{A}_i(y_{it}), i = 1, \dots, 4$ denotes linear approximations of the loss of power functions. The reservoir LSJ is equipped with another spillway (PDECH). The water flowing through this spillway does not influence the power generation, however this spillway has a limited capacity that depends on the storage of water.

In period t , $y_{3,1t}$ denotes spillage through the power plant IM, and $y_{3,2t}$ spillage through the outlet PDECH. $\tilde{D}(s_{3t})$ denotes the maximal water rate through the spillway PDECH.

During the summer, the bounds on the reservoir LSJ are very tight, as illustrated in Figure 8, in order to allow navigation and recreational activities. If at the end of the spring the natural inflows are low, it may not be possible to release water from this reservoir during the summer. An optimization problem may then be infeasible. To overcome this, we may penalize bounds violation in the objective function. We will use penalty terms of the form $-\phi \max\{0, \underline{s}_{t+1} - s_{t+1}\}$ and $-\phi \max\{0, s_{t+1} - \bar{s}_{t+1}\}$; where each hm^3 of water under or over the regulatory limit is penalized by the amount ϕ .




Figure 8: Bounds on the reservoir LSJ for each week of the year

The planning horizon of RTA's problem is the year with weekly time steps. In week t , we use the previous natural inflows Q_{t-1} as hydrological variable. Thus, for $t = 52, 51, \dots, 1$, the SDP recursion is $V_t(s_t, Q_{t-1}) = E_{Q_t|Q_{t-1}}[G_t(s_t, Q_t)]$, with

$$G_t(s_t, Q_t) = \max_{u_t, y_t} \left\{ \sum_{i=1,3} f_i(s_{it}, s_{i,t+1}, u_{it}) + f_2(s_{3t}, s_{3,t+1}, u_{2t}) + f_4(u_{4t}) - \sum_{j=1,2,4} \tilde{A}_j(y_j) - \tilde{A}_3(y_{3,1t}) - \phi \tilde{d}_{1t} - \phi \tilde{d}_{2t} + V_{t+1}(s_{t+1}, Q_t) \right\} \quad (29)$$

$$\text{S.t. } s_{t+1} = s_t + Q_{IRt} - B_{IR}u_t - C_{IR}y_t \quad (30)$$

$$\tilde{d}_{1t} \geq \underline{s}_{t+1} - s_{t+1} \quad (31)$$

$$\tilde{d}_{2t} \geq s_{t+1} - \bar{s}_{t+1} \quad (32)$$

$$0 \leq u_{it} \leq \tilde{B}_i(s_{it}), i = 1, 3 \quad (33)$$

$$u_{1t} + y_{1t} + Q_{2t} - u_{2t} - y_{2t} = 0 \quad (34)$$

$$u_{3t} + y_{3,1t} + y_{3,2t} - u_{4t} - y_{4t} = 0 \quad (35)$$

$$0 \leq u_{it} \leq \bar{u}_i, i = 2, 4 \quad (36)$$

$$y_{3,2t} \leq \tilde{D}(s_{3t}) \quad (37)$$

$$y_t \geq 0 \quad (38)$$

$$\tilde{d}_{1t}, \tilde{d}_{2t} \geq 0 \quad (39)$$

7.2 Experimental framework

We tested the simplicial methodology on RTA's subsystem through a two-phase approach: estimation of the value functions and simulation of the optimal operational policy. In each week, we used the simplicial decomposition scheme to construct the state space grids. As discussed in Section 5.1, in each week, a linear approximation of problem (29–39) (\hat{G}_t) was solved for each point of the grid. Recall that to implement GLP, moreover state space grid points, we also need evaluations of the production functions over a sample of points. These evaluations were provided by RTA, but we used interpolation error analysis to reduce the size of the original samples.

RTA's model assumes that release decisions are made after observation of the natural inflows. These latter were modeled as Markovian processes and, based on recommendations from RTA, were supposed to obey log-normal distributions. In period t , assume that $\{q_t^i | 1 \leq i \leq l_t\}$ are the possible realizations of the process Q_t , and for each value of Q_{t-1} , let $P(Q_t = q_t^i | Q_{t-1})$ be the associated conditional transition probabilities. We then have $\hat{V}_t(s_t, Q_{t-1}) = \sum_{i=1}^{l_t} P(Q_t = q_t^i | Q_{t-1}) \hat{G}_t(s_t, Q_t)$.

A multiple regression model with seasonal variables was constructed to estimate the conditional parameters of the log-normal distributions, which we used to estimate the conditional probabilities. The regression parameters were estimated through a 33-year sample of historical inflows. We adopted a similar scheme as in Karamouz and Vasiliadis (1992) to discretize the inflow processes. In week t , suppose the possible realizations of the inflows process are divided into l_t classes for which $q_t^1, \dots, q_t^{l_t}$ are the respective center. Let $F_t(Q_t = q_t^i | Q_{t-1})$ be the cumulative distribution of the log-normal distribution of $(Q_t = q_t^i | Q_{t-1})$ of parameters $\mu_t(Q_{t-1})$ et $\sigma^2(Q_{t-1})$. For each $Q_t = q_t$, we took $P(Q_t = q_t^1 | Q_{t-1}) = F_t(q_t^1)$, $P(Q_t = q_t^i | Q_{t-1}) = F_t(q_t^i) - F_t(q_t^{i-1})$, $i = 2, \dots, l_t - 1$, and $P(Q_t = q_t^{l_t} | Q_{t-1}) = 1 - F_t(q_t^{l_t-1})$. The

conditional probability of each value of Q_t then was estimated as the conditional probability of its class.

The terminal value of water V_{53} was first considered to be null. The value functions were estimated over 52 weeks three consecutive times. At the end of this procedure, we chose $V_{53} = V_1$ since it is a cycle. The 52 value functions were estimated again using such terminal value of water.

In the second phase, the 52 value functions thus estimated were used to simulate the operational policy for another 25-year sample of historical inflows. Given an observed state (s_t, Q_t) , the optimal policy is $[u_t, s_{t+1}, y_t] \in \text{Argmax}\{\hat{G}_t(s_t, Q_t)\}$, which was obtained by solving the linear approximation of (29–39). An initial storage s_1 was chosen corresponding to the maximal level of the first reservoir and about 98% of the second one. In the first period the linear approximation of (29–39) is solved for the state (s_1, Q_1) . Afterwards, the approximation is solved for (s_2, Q_2) , where s_2 is the optimal final stock at the end the first period. The process was repeated until the last week of the 25-year simulation horizon. For comparison purposes, in the simulation phase we used the same functions as in RTA’s model to evaluate the production corresponding to the optimal releases.

The results of the simulations were compared to those obtained with a classical SDP model constructed by RTA, which we will refer to as RTA’s model. To construct the value functions, in each week, RTA’s model used a fixed $10 \times 30 \times 7$ grid points (these numbers are the number of discretization points of the first and second reservoirs’ levels, and the number of discretization values of the inflows, respectively). We discretized the inflow processes into the same number as in RTA’s model.

7.3 Results and analysis

The two models (RTA and the simplicial scheme) were compared based on the average effectiveness of the simulated policy defined as the ratio of the average generated power to the total outflows (water release and spillage). We do not report the raw results for confidentiality purposes.

Table 5 reports the average effectiveness for each power plant for the two models as well as statistics pertaining to the number of grid points per week to construct the value functions. In our model, the value functions were estimated for three different thresholds on the imprecision ($\eta = \frac{\Delta_\nu}{\Delta_1}$); ν is an iteration counter. The solution time (to estimate the 52 value functions) is also reported for each threshold.

Overall, on average, the effectiveness of the policy is very similar in both cases, with a slight edge in favor of our model. However, the computational burden was significantly reduced with our model. In each week, our method performed fewer evaluations of the value function than did RTA’s model. In each week, the latter model performed 2100 evaluations of the value function (300 stock grid points \times 7 discrete values of inflows) for a total of 109 200 evaluations over the 52 weeks. By contrast, with an imprecision $\eta = 0.8$, in each week, on average, our model performed 413 evaluations of the value functions (59 stock grid points \times 7 discrete values of inflows), for a total of 21 476 evaluations over the entire year, which roughly represents 20% of the total evaluations performed by RTA’s model.

Let us also observe that, in our model, while the solution time increased significantly with the precision $1 - \eta$, the effectiveness of the policy remained quite similar. These results reveal

Table 5: Comparison between the two models: effectiveness of the simulated policy and number of grid points per week

Model	Effectiveness in $MW/m^3/s$				
	CCD	CCS	CIM	CSH	Overall average
RTA	0.278	0.295	0.252	0.566	0.348
Simplicial ($\eta = 0.8$)	0.268	0.308	0.255	0.566	0.349
Simplicial ($\eta = 0.5$)	0.268	0.308	0.254	0.566	0.349
Simplicial ($\eta = 0.1$)	0.268	0.308	0.254	0.566	0.349
Model	Number of stock grid points per week				
	Min	Max	Average	–	–
RTA	300.000	300.000	300.000	–	–
Simplicial ($\eta = 0.8$)	57.000	65.000	59.000	–	–
Simplicial ($\eta = 0.5$)	62.000	78.000	68.327	–	–
Simplicial ($\eta = 0.1$)	117.000	553.000	188.238	–	–
Solution time for the simplicial model in seconds					
	$\eta = 0.8$	$\eta = 0.5$	$\eta = 0.1$	–	–
	175.032	200.144	662.166	–	–

that densifying the stock grids, i.e. performing significant number of evaluations of the value function in each week, does not necessarily improve the operational policy of the system. Therefore, the problem may be reasonably approximated by evaluating the value function at few appropriately selected grid points and by performing interpolations elsewhere.

Figure 9 shows the average optimal trajectory of the reservoirs as prescribed by both models (RTA and simplicial). In both cases, at the beginning of the year, the level of the reservoirs are high. Then the reservoirs are progressively emptied in anticipation of the spring run off. Then refill occurs. Except for the summer, RTA’s model operates the reservoir CD with higher reservoir levels than does our model. However, in general, this is the contrary for LSJ.




Figure 9: Comparison of the average optimal trajectory of the reservoirs

8. Conclusions

We presented an approximate stochastic dynamic programming methodology based on simplicial partitioning of the state space. The vertices of the simplices form irregular grid points where the value function is approximated. We exploited concavity assumptions on the value function to develop bounds to provide local measures of the approximation error. These measures are exploited to refine the grids where the curvature of the value function is higher. Furthermore, we showed that a concave approximation of the value function may be obtained through generalized linear programming.

In a particular context of information-decision process, we also showed that the expectation of the value function may be simplified to analytical forms if the inflows to the reservoirs may be assumed to be spatially correlated.

We illustrated the trade-off between solution time and accuracy allowed by the method through 45 randomly generated instance of problems. Our scheme was also applied in an industrial context. We first estimated the value functions using a 33-year sample of natural inflows. On the other hand, these value functions constituted inputs for the simulation of the operational policy over a 25-year horizon.

The performance of our method (ratio of the average power to the total outflows) was similar to that of a classical stochastic dynamic programming. However, the computational burden significantly decreased as in each week our model significantly spared evaluations of the value function compared to the classical model. Thus, the problem may be solved for few appropriately selected grid points and interpolated elsewhere with an acceptable precision.

Appendix

Proof of proposition 1. The spillage network is arborescent. Without loss of generality, assume the nodes are labeled in a topological order (as in Figure 1). The incidence matrix C of this network is then lower triangular. Since each node is equipped with a spillage system, the diagonal elements of this matrix are all equal to 1. It follows that C is non-singular and $C^{-1} \in \{0, 1\}^{p \times p}$ (Bazaraa et al., 1990).

Now, let us show that for all $ds \geq 0$, $V_{t-1}(s_{t-1} + ds, \epsilon_{t-1}) \geq V_{t-1}(s_{t-1}, \epsilon_{t-1})$. To simplify the notation, we omit the hydrological variable ϵ_{t-1} .

Define

$$\phi_t(s_{t-1}, u_t, Q_t, y_t) = \sum_{i \in I_{RC}} f_{it}(u_{it}, s_{i,t-1}, \epsilon_{t-1}) + \sum_{i \in I_F} f_{it}(u_{it}) + V_t(s_{t-1} - Bu_t - Cy_t + Q_t, \epsilon_t).$$

Function ϕ_t measures the value of water for given reservoirs' levels s_{t-1} , release decisions u_t , observed inflows Q_t and spillage decisions y_t . Given an initial stock s_{t-1} , in period t , the optimal policy consists in a rule $u_t^*(s_{t-1})$ followed by a rule $y_t^*(s_{t-1} - Bu_t + Q_t)$. It follows that

$$V_{t-1}(s_{t-1}) = E_{Q_t} [\phi_t(s_{t-1}, u_t^*(s_{t-1}), Q_t, y_t^*(s_{t-1} - Bu_t^*(s_{t-1}) + Q_t))].$$

Suppose the initial stock increases from s_{t-1} to $\hat{s}_{t-1} = s_{t-1} + ds$ ($ds \geq 0$). Consider a new policy where such stock increases are spilled:

$$\hat{u}_t(\hat{s}_{t-1}) = u_t^*(s_{t-1}),$$

$$\hat{y}_t(\hat{s}_{t-1} - B\hat{u}_t(\hat{s}_{t-1}) + Q_t) = y_t^*(s_{t-1} - Bu_t^*(s_{t-1}) + Q_t) + dy,$$

with $Cdy = ds$.

This policy is feasible. Indeed (i) the planned releases do not change and satisfy their bounds, (ii) the final stock does not change, (iii) the spillage increases by $dy = C^{-1}ds \geq 0$. In addition, one verifies that

$$\begin{aligned} \phi_t(s_{t-1}, \hat{u}_t(\hat{s}_{t-1}), Q_t, \hat{y}_t(\hat{s}_{t-1} - B\hat{u}_t(\hat{s}_{t-1}) + Q_t)) = \\ \phi_t(s_{t-1}, u_t^*(s_{t-1}), Q_t, y_t^*(s_{t-1} - Bu_t^*(s_{t-1}) + Q_t)). \end{aligned}$$

Since the new policy is feasible but not necessarily optimal for the initial stock level $s_{t-1} + ds$, we get:

$$\begin{aligned} V_{t-1}(s_{t-1} + ds) &\geq E_{Q_t} [\phi_t(s_{t-1}, \hat{u}_t(\hat{s}_{t-1}), Q_t, \hat{y}_t(\hat{s}_{t-1} - B\hat{u}_t(\hat{s}_{t-1}) + Q_t))] \\ &= E_{Q_t} [\phi_t(s_{t-1}, u_t^*(s_{t-1}), Q_t, y_t^*(s_{t-1} - Bu_t^*(s_{t-1}) + Q_t))] \\ &= V_{t-1}(s_{t-1}). \end{aligned}$$

□

Proof of proposition 3. We have

$$\begin{aligned} E[\tilde{\tau}_t^b]_a &= E[e^Z]_{\ln a}^{\ln b}, \text{ where } Z \sim N(\mu(\epsilon_t), \sigma(\epsilon_t)^2) \\ &= \frac{1}{\sqrt{2\pi\sigma(\epsilon_t)}} \int_{\ln a}^{\ln b} e^{-\frac{1}{2}\left(\frac{z-\mu(\epsilon_t)}{\sigma(\epsilon_t)}\right)^2 + z} dz. \end{aligned}$$

We also have

$$\begin{aligned}
-\frac{1}{2} \left(\frac{z - \mu(\epsilon_t)}{\sigma(\epsilon_t)} \right)^2 + z &= -\frac{1}{2\sigma(\epsilon_t)^2} (z^2 - 2\mu(\epsilon_t)z + \mu(\epsilon_t)^2 - 2\sigma(\epsilon_t)^2 z) \\
&= -\frac{1}{2\sigma(\epsilon_t)^2} (z^2 - 2(\mu(\epsilon_t) + \sigma(\epsilon_t)^2)z + \mu(\epsilon_t)^2) \\
&= -\frac{1}{2\sigma(\epsilon_t)^2} \left[(z - (\mu(\epsilon_t) + \sigma(\epsilon_t)^2))^2 - (\mu(\epsilon_t) + \sigma(\epsilon_t)^2)^2 + \mu(\epsilon_t)^2 \right] \\
&= -\frac{(z - (\mu(\epsilon_t) + \sigma(\epsilon_t)^2))^2}{2\sigma(\epsilon_t)^2} + \mu(\epsilon_t) + \frac{\sigma(\epsilon_t)^2}{2},
\end{aligned}$$

Hence

$$E[\tilde{\tau}_t | \epsilon_t]_a^b = \frac{\nu}{\sqrt{2\pi}\sigma(\epsilon_t)} \int_{\ln a}^{\ln b} e^{-\frac{1}{2} \left(\frac{z - \mu(\epsilon_t) - \sigma(\epsilon_t)^2}{\sigma(\epsilon_t)} \right)^2} dz.$$

By the change of variable $t = \frac{z - \mu(\epsilon_t) - \sigma(\epsilon_t)^2}{\sigma(\epsilon_t)}$, we get

$$E[\tilde{\tau}_t | \epsilon_t]_a^b = \nu \int_{\rho(a)}^{\rho(b)} \frac{1}{\sqrt{2\pi}} e^{-\frac{1}{2}t^2} dt.$$

The integrand is the density function of the standard normal variable, which proves the proposition. \square

Proof of proposition 4. The objective is linear and is maximized. The set $D_t = \{u_t, x_t, \lambda, \mu | (13 - 23)\}$ is convex. It follows that \hat{V}_{t-1} is concave in $s_{t-1} \forall \epsilon_{t-1}$. \square

Proof of proposition 5. Since the production functions f_t are concave in u_t , for the same releases, the production value estimated as the best interpolation (convex combination) over samples points $\{\hat{u}_{jt} | j \in \Gamma_t\}$ is a lower bound on the actual production. In addition, since $\hat{V}_t \leq V_t$, we have $\hat{J}_t(x_t, \epsilon_{t-1}) \leq E_{Q_t, \tilde{\epsilon}_t | \epsilon_{t-1}} [V_t(s_t^*(x_t + Q_{I_{Rt}}), \tilde{\epsilon}_t)] \forall x_t$ (by linearity of the expectation). Therefore, for any post-release reservoir levels x_t , the expectation evaluated as the best interpolation over sample points $\{\hat{x}_{jt} | j \in \Upsilon_t\}$ is underestimated. Lastly, let us observe that the polyhedron (13-23) (feasible domain of \hat{V}_{t-1}) is included in the polyhedron (8-11) (feasible domain of V_{t-1}). Since the objective is maximized, these prove our claim. \square

References

- Bazararaa, M., Jarvis, J., and Sherali, H. *Linear programming and network flows*. Wiley, New York, 1990.
- Candler, G. V. *Finite-difference methods for continuous-time dynamic programming*. Computational Methods for the Study of Dynamic Economies, Oxford University Press, 2001.
- Cervellera, C., Wen, A., and Chen, V. C. Neural network and regression spline value function approximations for stochastic dynamic programming. *Computers & Operations Research*, 34(1):70–90, 2007.
- Etkin, D., Kirshen, P., Watkins, D., Roncoli, C., Sanon, M., Some, L., Dembele, Y., Sanfo, J., Zoungrana, J., and Hoogenboom, G. Stochastic Programming for Improved Multiuse Reservoir Operation in Burkina Faso, West Africa. *Journal of Water Resources Planning and Management*, pp. 141, 04014 056. doi:10.1061/(ASCE)WR.1943-5452.0000 396, 2015.
- Faber, B., and Stedinger, J. Reservoir optimization using sampling SDP with ensemble streamflow prediction (ESP) forecasts. *Journal of Hydrology*, 249(1):113–133, 2001.
- Goor, Q., Kelman, R., and Tilmant, A. Optimal multipurpose-multireservoir operation model with variable productivity of hydropower plants. *Journal of Water Resources Planning and Management*, 137(3):258–267, 2010.
- Grüne, L., and Semmler, W. Using dynamic programming with adaptive grid scheme for optimal control problems in economics. *Journal of Economic Dynamics and Control*, 28(12):2427–2456, 2004.
- Jiekang, W., Jianquan, Z., Guotong, C., and Hongliang, Z. A hybrid method for optimal scheduling of short-term electric power generation of cascaded hydroelectric plants based on particle swarm optimization and chance-constrained programming. *IEEE Transactions on Power Systems*, 23(4):1570–1579, 2008.
- Johnson, S. A., Stedinger, J. R., Shoemaker, C. A., Li, Y., and Tejada-Guibert, J. A. Numerical solution of continuous-state dynamic programs using linear and spline interpolation. *Operations Research*, 41(3):484–500, 1993.
- Karamouz, M., and Vasilidis, H. V. Bayesian stochastic optimization of reservoir operation using uncertain forecasts. *Water Resources Research*, 28(5):1221–1232, 1992.
- Kim, Y.-O., Eum, H.-I., Lee, E.-G., and Ko, I. H. Optimizing operational policies of a Korean multireservoir system using sampling stochastic dynamic programming with ensemble streamflow prediction. *Journal of Water Resources Planning and Management*, 133(1):4–14, 2007.
- Kracman, D. R., McKinney, D. C., Watkins, D., and Lasdon, L. Stochastic optimization of the highland lakes system in Texas. *Journal of Water Resources Planning and Management*, 132(2):62–70, 2006.

- Labadie, J. W. Optimal operation of multireservoir systems: State-of-the-art review. *Journal of Water Resources Planning and Management*, 130(2):93–111, 2004.
- Lamond, B. F., and Lang, P. Stochastic optimization of multi-reservoir systems with power plants and spillways. *WIT Transactions on Ecology and the Environment*, 104:31–40, 2007.
- Lee, Y., Kim, S., and Ko, I. Multistage stochastic linear programming model for daily coordinated multi-reservoir operation. *Journal of Hydroinformatics*, 10(1):23–41, 2008.
- Lee, Y., Kim, S.-K., and Ko, I. H. Two-stage stochastic linear programming model for coordinated multi-reservoir operation. In *Operating Reservoirs in Changing Conditions*, pp. 400–410. ASCE, 2006.
- Martinez, D. L., Shih, D. T., Chen, V. C., and Kim, S. B. A convex version of multivariate adaptive regression splines. *Computational Statistics & Data Analysis*, 81:89–106.doi:10.1016/j.csda.2014.07.015, 2015.
- Moore, D. W. Simplicial mesh generation with applications. Ph.D. thesis, 1992.
- Mousavi, S. J., Karamouz, M., and Menhadj, M. B. Fuzzy-state stochastic dynamic programming for reservoir operation. *Journal of Water Resources Planning and Management*, 2004.
- Ouarda, T., and Labadie, J. Chance-constrained optimal control for multireservoir system optimization and risk analysis. *Stochastic Environmental Research and Risk Assessment*, 15(3):185–204, 2001.
- Pérez-Díaz, J. I., Wilhelmi, J. R., and Arévalo, L. A. Optimal short-term operation schedule of a hydropower plant in a competitive electricity market. *Energy Conversion and Management*, 51(12):2955–2966, 2010.
- Philbrick, C. R., and Kitanidis, P. K. Limitations of deterministic optimization applied to reservoir operations. *Journal of Water Resources Planning and Management*, 125(3):135–142, 1999.
- Rust, J. Numerical dynamic programming in economics. *Handbook of computational economics*, 1:619–729, 1996.
- Seifi, A., and Hipel, K. W. Interior-point method for reservoir operation with stochastic inflows. *Journal of Water Resources Planning and Management*, 127(1):48–57, 2001.
- Shapiro, J. F. *Mathematical Programming: Structures and Algorithms*. Wiley New York, 1979.
- Tejada-Guibert, J. A., Johnson, S. A., and Stedinger, J. R. The value of hydrologic information in stochastic dynamic programming models of a multireservoir system. *Water Resources Research*, 31(10):2571–2579, 1995.
- Turgeon, A. Stochastic optimization of multireservoir operation: The optimal reservoir trajectory approach. *Water Resources Research*, 43(5):W05 420, 2007.

- Van Ackooij, W., Henrion, R., Möller, A., and Zorgati, R. Joint chance constrained programming for hydro reservoir management. *Optimization and Engineering*, 15(2):509–531, 2014.
- Xu, W., Zhang, C., Peng, Y., Fu, G., and Zhou, H. A two stage Bayesian stochastic optimization model for cascaded hydropower systems considering varying uncertainty of flow forecasts. *Water Resources Research*, 50(12):9267–9286. doi:10.1002/2013WR015181, 2014.
- Zeng, Y., Wu, X., Cheng, C., and Wang, Y. Chance-Constrained Optimal Hedging Rules for Cascaded Hydropower Reservoirs. *Journal of Water Resources Planning and Management*, 140(7):04014010, 2013.
- Zéphyr, L., Lang, P., and Lamond, B. F. Controlled approximation of the value function in stochastic dynamic programming for multi-reservoir systems. *Computational Management Science*, DOI 10.1007/s10287-015-0242-1, published Online : 01 October, 2015.
- Zhao, T., Cai, X., Lei, X., and Wang, H. Improved dynamic programming for reservoir operation optimization with a concave objective function. *Journal of Water Resources Planning and Management*, 138(6):590–596, 2011.