Article:

Rougé, C. and Tilmant, A. (2016) Using stochastic dual dynamic programming in problems with multiple near-optimal solutions. Water Resources Research, 52 (5). pp. 4151-4163. ISSN 0043-1397

https://doi.org/10.1002/2016wr018608

Author-produced version of the paper published in
Water Resources Research, vol.52, doi:10.1002/2016WR018608 (2016).
Original paper available at
http://http://onlinelibrary.wiley.com/doi/10.1002/2016WR018608/full

Using stochastic dual dynamic programming in problems with multiple near-optimal solutions

Charles Rougé[1,2] and Amaury Tilmant[1]

May 3, 2016

Abstract

Stochastic dual dynamic programming (SDDP) is one of the few algorithmic solutions available to optimize large-scale water resources systems while explicitly considering uncertainty. This paper explores the consequences of, and proposes a solution to, the existence of multiple near-optimal solutions (MNOS) when using SDDP for mid- or long-term river basin management. These issues arise when the optimization problem cannot be properly parametrized due to poorly defined and/or unavailable data sets. This work shows that when MNOS exists, 1) SDDP explores more than one solution trajectory in the same run, suggesting different decisions in distinct simulation years even for the same point in the state-space, and 2) SDDP is shown to be very sensitive to even minimal variations of the problem setting, e.g. initial conditions – we call this "algorithmic chaos". Results that exhibit such sensitivity are difficult to interpret. This work proposes a re-optimization method, which simulates system decisions by periodically applying cuts from one given year from the SDDP run. Simulation results obtained through this re-optimization approach are steady-state solutions, meaning that their probability distributions are stable from year to year.
Keywords: Stochastic dual dynamic programming (SDDP), Year-periodic re-optimization, Multiple near-optimal solutions, Limited data availability, Zambezi River Basin, Chaos.

Key points

- SDDP results can be hard to interpret in the presence of multiple near-optimal solutions.
- A year-periodic re-optimization method is proposed to solve this issue.
- Limited data availability favors the presence of multiple near-optimal solutions.

1 Introduction

Stochastic dual dynamic programming (SDDP; *Pereira*, 1989; *Pereira and Pinto*, 1991) is an approximate stochastic optimization algorithm to analyze multistage, stochastic, decision making problems such as reservoir operation, irrigation scheduling, intersectoral allocation, etc. SDDP is one of the

[1]Department of Civil Engineering and Water Engineering, Université Laval, Québec, Canada
[2]Department of Mechanical, Civil and Aerospace Engineering, University of Manchester, United Kingdom

few algorithmic solutions available to handle large-scale problems, i.e. problems characterized by a large state-space, while explicitly considering the hydrologic uncertainty. To achieve this, SDDP constructs a locally-accurate approximation of the decision making problem. As we will see later, this approximation might be a source of concern when the decision making problem cannot be properly parametrized.

Initially, SDDP was developed for short- and mid-term hydropower scheduling in hydropower-dominated systems, e.g. in Brazil (*Pereira*, 1989; *Pereira and Pinto*, 1991; *Maceira and Damázio*, 2004) or in Norway (*Rotting and Gjelsvik*, 1992; *Mo et al.*, 2001). It has then been extended to long-term hydropower scheduling (e.g. *Gjelsvik et al.*, 2010; *Homem-de Mello et al.*, 2011; *Bezerra et al.*, 2012), where the planning horizon extends over several years. Furthermore, its ability to assess the marginal value of water at each stage and place in the basin (*Tilmant et al.*, 2008, 2012) has made it useful for the hydro-economic analysis (*Harou et al.*, 2009) of large river basins. SDDP enables tackling varied issues including risk assessment (*Tilmant and Kelman*, 2007) or cost assessment of noncoordinated irrigation development among riparian countries (*Tilmant and Kinzelbach*, 2012), the restoration of a flow regime through the coordination of multiple reservoirs (*Tilmant et al.*, 2010) or the integrated assessment of possible future developments in the Blue Nile River Basin (*Goor et al.*, 2010; *Arjoon et al.*, 2014).

This work explores the consequences of the presence of multiple near optimal-solutions (MNOS) when using SDDP to analyze water resources allocation problems with limited data. The existence of MNOS has been demonstrated for multiple reservoir systems in the case of deterministic inflow sequences (*Liu et al.*, 2011). Within an implicit stochastic optimization framework, finding MNOS enhances the flexibility of decision rules (*Liu et al.*, 2014; *Zhang et al.*, 2015), thus turning MNOS into an opportunity. This work does not seek to search for MNOS but rather to analyze the existence of MNOS with respect to the algorithmic structure of SDDP and the quality of the dataset.

Since SDDP is an approximate optimization algorithm, it can, in the presence of MNOS, potentially switch between different near-optimal decision rules at each stage where decisions must be taken. Such a behavior would mean that a steady-state operating policy may not exist. In other words, since the statistical distribution of inflows is year-periodic, one would expect that the probability distribution of simulated variables (storage, power production, etc) converges toward a year-periodic steady state.

The basic idea behind SDDP is to approximate the convex benefit-to-go function by Bender's cuts, mathematical objects that can be thought of as hyperplanes. The algorithm then simulates reservoir operation decisions through the use of these hyperplanes approximating the true benefit-to-go functions. To deal with the convergence issue associated with MNOS, a year-periodic re-optimization (YPRE) procedure, named after the very similar re-optimization approach (*Tejada-Guibert et al.*, 1993), is proposed. By repetitively imposing the consecutive cuts from a year-long period over the whole planning period, YPRE yields steady-state solutions

This work also draws a link between limited data availability and the existence of MNOS. The collection of data for both water supplies and demands is one of the challenges associated with the modeling and analysis of large-scale water resources systems. This is especially the case in developing countries where data availability is limited (or simply absent), and/or when not all stakeholders have an incentive to provide the relevant data. For example, because demand curves (marginal net benefit functions) in irrigated agriculture are rarely available, authors are left with no choice but to assume horizontal functions whereby the marginal net benefit is constant over the range of supplies (*Wu and Whittington*, 2006). In the energy sector, the assumption that turbine efficiencies remain constant regardless of the head (*Archibald et al.*, 1999; *Cai et al.*, 2002; *Wallace and Fleten*, 2003; *Tilmant and Kelman*, 2007; *Goor et al.*, 2011) is both more computationally convenient and less data intensive than variable efficiency. Besides, the hydropower companies are assumed to be price takers, i.e. the value of electricity is independent of their power output (*Gjelsvik et al.*, 2010). In practice however, representing the impact of hydroelectric production on electricity prices leads to more realistic reservoir operation policies (*Pereira-Cardenal et al.*, 2015).

These examples of limited data availability often translate into a less convex, and more linear, problem formulation. This results in situations in which the marginal value of water is constant or undergoes little variation when allocating large amounts to competing uses. These quantities may then be allocated to either of the uses with scarcely any impact on the objective function, favoring the existence of MNOS. This effect is sometimes compounded with the physical characteristics of the system, such as the existence of large reservoirs where the head varies little despite large variations in storage.

The rest of this work is as follows. Section 2 presents the SDDP algorithm as well as the case study inspired from the Zambezi River basin in Africa. Section 3 then demonstrates how limited data availability and the presence of MNOS pose challenges to the production of a steady-state solution. Section 4 proposes YPRE to address these issues. Finally, Section 5 discusses some issues raised by the work of Sections 2 to 4. Concluding remarks are given in Section 6.

2 Material and methods

2.1 Principle of the SDDP algorithm

For the sake of concision, this section gives a basic description of the SDDP algorithm, emphasizing the aspects that are of interest for the present paper. For more details please see for instance *Tilmant et al.* (2008).

2.1.1 Objective

SDDP is used to optimize the expected value of a benefit function or a cost function over a given planning horizon involving T stages (weeks, months). :

$$Z = E\left[\sum_{t=1}^{T} f_t(\mathbf{x}_t, \mathbf{q}_t, \mathbf{u}_t) + \nu(\mathbf{x}_{T+1})\right] \tag{1}$$

where $E[.]$ is the expectation operator, $f_t(.)$ represents the benefits to be reaped from system operation at stage t and $\nu(.)$ is a terminal value function. Vector $\mathbf{x}_t$ is the system state, which typically includes beginning-of-period storage $\mathbf{s}_t$ and previous inflow $\mathbf{q}_{t-1}$; vector $\mathbf{q}_t$ represents inflow into the system at stage t, and $\mathbf{u}_t$ is the vector of all decisions to be taken to manage the system, e.g. electricity generation, reservoir release and spillage, water withdrawals, etc. Problem (1) is optimized under a set of hydrological, physical and institutional constraints.

2.1.2 Decomposition into one-stage linear problems

SDDP breaks down the multi-stage non-linear problem (1) into a series of one-stage linear problems which are solved recursively. This is made possible by assuming that both $f_t(.)$ and all the constraints are linear functions, meaning that the overall problem (1) is convex. In practice, however, many water resource management problems are non-convex. This is for example the case when pumping costs are head-dependent (*Davidsen et al.*, 2016), or where price is an endogenous variable in hydropower scheduling problems (*Mo et al.*, 2001; *Kristiansen*, 2004).

At stage t, the one-stage problem is solved for state $\mathbf{x}_t$, assuming current inflows $\mathbf{q}_t$:

$$F(\mathbf{x}_t) = \max_{\mathbf{u}_t} \{f_t(\mathbf{x}_t, \mathbf{q}_t, \mathbf{u}_t) + F_{t+1}\} \tag{2}$$

where benefits from hydropower generation, benefits from other offstream and instream uses, and penalties all are expressed through linear inequality constraints. When data availability is limited, a piecewise linear approximation of the hydropower production function (*Goor et al.*, 2011) is an improvement over the assumption that turbine efficiency remains constant regardless of the head.

F_{t+1} represents the benefit-to-go function. It is bounded by L Bender's cuts, which are inequality constraints:

$$F_{t+1} \leq \mathbf{a}^l_{t+1}.\mathbf{x}_{t+1} + \beta^l_{t+1} \tag{3}$$

Cuts $l = 1 \dots L$ form an approximation of the benefit-to-go function at stage t. Similar to state $\mathbf{x}_{t+1}$ lumping together variables such as $\mathbf{s}_{t+1}$ and $\mathbf{q}_t$, $\mathbf{a}^l_{t+1}$ lumps together the vectors that represented the marginal values of these variables in earlier presentations of SDDP (*Tilmant and Kelman*, 2007; *Tilmant et al.*, 2008).

The one-stage problem is also subject to physical constraints, such as water balance constraints, upper and / or lower bounds on release and storage decisions.

SDDP then proceeds iteratively, after an initialization phase where a first sequence of state variables x_t is produced for $t = 1 \dots T$. This can be done by simulating system operations in a "myopic" way, in the sense that it disregards future benefits. For each iteration there are two phases, and iteration L ($L \geq 1$) is as follows.

2.1.3 Backward optimization phase

The backward optimization phase builds the Benders' cuts defined in equation (3), recursively from the last stage back to the first one. At a given stage t, there are L sampled states – same as the iteration number. For each sampled state, an approximation of the value function is obtained by solving K times the one-stage problem of equation (2):

$$F^*(\mathbf{x}_t) = \frac{1}{K} \sum_{k=1}^{K} \left[\max_{\mathbf{u}_t} \left\{ f_t(\mathbf{x}_t, \mathbf{q}^k_t, \mathbf{u}_t) + F_{t+1} \right\} \right] \tag{4}$$

where each $\mathbf{q}_t^k$ is called a backward opening. It is linked with $\mathbf{q}_{t-1}$ through a PAR(1) model that also accounts for spatial correlation. Recall that the current state $\mathbf{x}_t$ includes previous inflow $\mathbf{q}_{t-1}$. Coefficients $\mathbf{a}_{t+1}^l$ and β_{t+1}^l that define F_{t+1} in equation (3) come from the dual solutions to the K one-stage linear problems of equation (4) at stage $t+1$. They can be interpreted as the derivatives of the objective with respect to the state variables.

At each stage, F_{t+1} is an upper approximation to the future benefits. Therefore, the backward phase of SDDP yields an upper approximation $\overline{Z}$ to the objective function:

$$\overline{Z} = F^*(\mathbf{x}_0) \tag{5}$$

where $\mathbf{x}_0$ is the initial state.

2.1.4 Forward simulation phase

In the simulation phase, SDDP uses M inflow time-series $(\mathbf{q}_t^m)_{1\leq t\leq T}$, also called simulation sequences. For each sequence and from an initial state $\mathbf{x}_0$, the one-stage linear problem of equation (2) is solved recursively forward from stage 1 to stage T using the benefit-to-go functions derived in the backward optimization phase.

Thus, the forward simulation phase yields all the successive states $\mathbf{x}_t$ and decisions $\mathbf{u}_t$ for each of the M simulation sequences. Since operations are at best optimal, we can define a lower bound for the optimum of the objective Z as follows:

$$\underline{Z} = \frac{1}{M}\sum_{m=1}^{M}\left[\sum_{t=1}^{T} f_t(x_t, q_t^m, u_t) + \nu(x_{T+1}^m)\right] \tag{6}$$

$\underline{Z}$ is an estimate for the expected benefits using the cuts derived in the backward optimization phase, and for $M \geq 30$, a 95% confidence interval for that value can easily be computed. If the upper approximation $\overline{Z}$ falls in that confidence interval, one can consider that SDDP has converged towards an approximately optimal solution. After SDDP has converged, the results of the last simulation phase can be exploited (allocation decisions, marginal water value) and their probability distribution be traced out.

Otherwise, iteration $L+1$ of the algorithm is necessary. Then a new sequence of sample states has to be added to the L existing sequences, to be used in the backward optimization phase of iteration $L+1$. It comes from simulation results of the forward phase of iteration L, usually for a sequence of historical inflows. This additional sequence of states will help refine the approximations of equation (3).

2.2 Illustrative case-study

In this work, the SDDP algorithm described in Section 2.1 is applied to a case-study of the Zambezi River basin (Figure 1). It covers an area of 1.39 million km^2 and is shared by eight riparian countries (Angola, Botswana, Malawi, Mozambique, Namibia, Tanzania, Zambia, Zimbabwe). It is an important regional water resources system in terms of energy generation and food production. Historically, hydropower generation has been the largest economic use of water in the basin. It has also been a

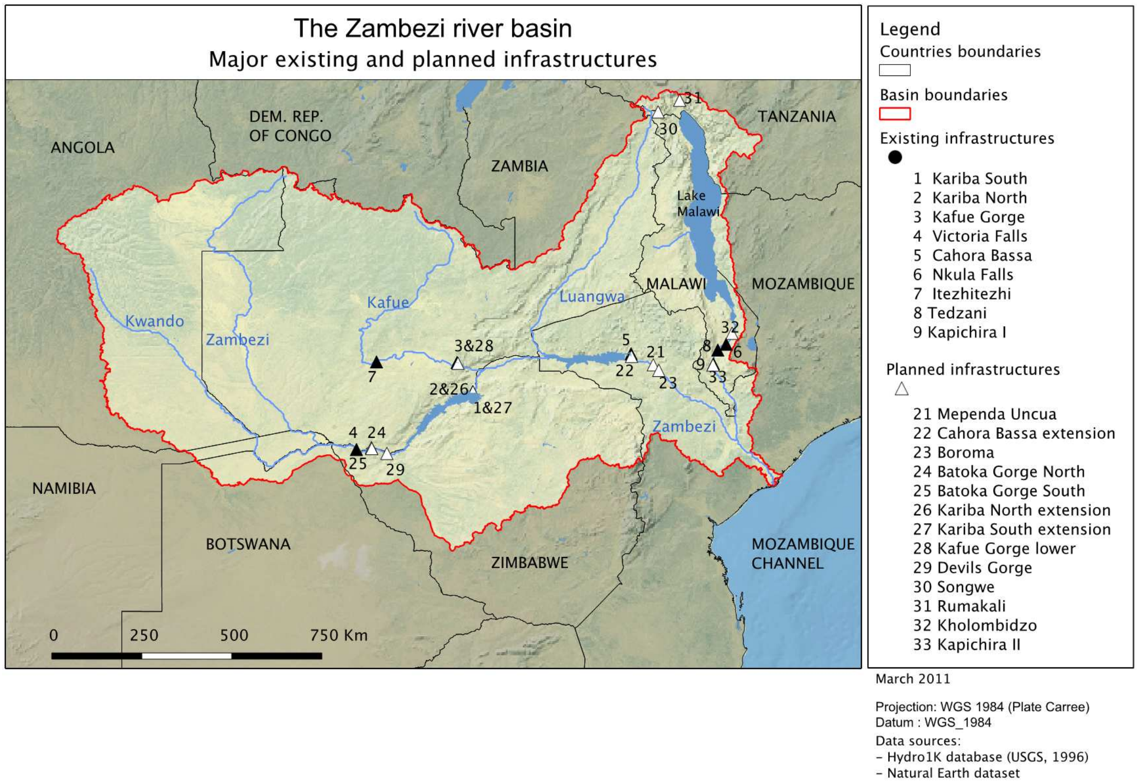

Figure 1: The Zambezi river basin: geographical situation and infrastructure.

factor of regional integration with the establishment of the Zambezi River Authority for the operation of the Kariba reservoir, and the interconnection of most of the hydropower stations with the Southern African Power Pool (SAPP). The second largest economic use of water is irrigated agriculture, which mainly occurs in the lower and middle Zambezi. Despite the evaporation losses from the large man-made reservoirs and irrigation consumptive uses, the Zambezi is still a largely open river basin (with the notable exception of the Kafue, a tributary flowing through Zambia). Like many other regional water resource systems, the development in the Zambezi River Basin requires careful planning to simultaneously achieve water, food and energy security. The presence of rival and non-rival uses associated with a complex topology make this basin sensitive to unilateral developments, therefore increasing the risk of collateral damages downstream. This paper uses an intermediate development scenario presented by *Tilmant and Kinzelbach* (2012), with both existing and planned infrastructure. A schematic representation of the basin is provided by Figure 2, and the main characteristics of infrastructure and irrigation areas can be found in Tables 1 and 2.

This case-study features some characteristics put forward in the introduction as contributing factors to the presence of MNOS: A) both the marginal net value of water for irrigation and the marginal price of hydropower are assumed to be constant, and B) presence of large reservoirs that display minimal head variations even for some large variations in storage. These two large reservoirs have already been built, and they are Kariba and Cahora Bassa, respectively at nodes 4 and 6 in Figure 2.

SDDP is run with a monthly time-step over a ten-year planning period, so that $T = 120$. Energy

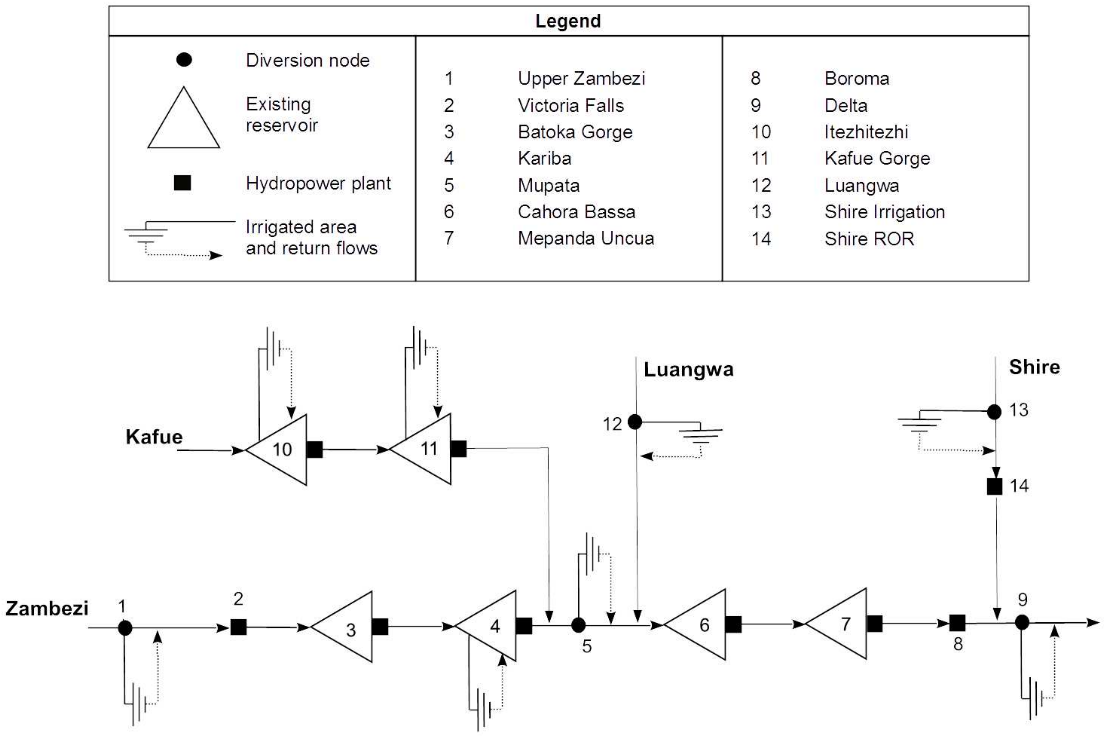

Figure 2: Schematic representation of the Zambezi River Basin.

Name	Node number	Capacity (MW)	Storage volume (km^3)
Kariba	4	1,980	180.6
Cahora Bassa	6	2,925	77.7
Kafue Gorge	11	1,500	9.5
Itezhitezhi	10	120	5.6
Mepanda Uncua	7	1,500	2.3
Batoka Gorge	3	1,600	1.7
Victoria Falls	2	108	0
Boroma	8	160	0
Nkula + Kedzani + Kapichira	14	279	0

Table 1: Hydropower plants in the chosen Zambezi case study, by decreasing storage capacity of corresponding reservoir.

Name	Node number	Irrigated area ($\times 10^3$ ha)
Upper Zambezi	1	77
Zimbabwe	4	116
Mupata	5	1
Delta	9	120
Upper Kafue	10	97
Kafue Flats	11	40
Luangwa River	14	108
Shire River	231	160

Table 2: Irrigated in the chosen Zambezi case study. Figures correspond to scenarios of future developments as reported in *Tilmant et al.* (2012) and *Tilmant and Kinzelbach* (2012).

Scenario	Cahora Bassa (km^3)	Kariba (km^3)
$S1$	57.0	154.7
$S2$	56.9	154.8

Table 3: Initial storage differences that define scenarios $S1$ and $S2$, used to highlight the existence of algorithmic chaos.

prices are fixed at 40$/MWh. Similar to *Tilmant and Kinzelbach* (2012), the net benefit from irrigation is estimated as 500$ per ha and per year. Some environmental constraints have been added, by imposing an environmental pulse of $7,500m^3/s$ in February in the Delta (node 9), as recommended in *World Bank* (2010a); and a smaller pulse in the Kafue Flats (upstream of node 11) of $300m^3/s$ in March. The penalty for not meeting thee requirements is set at 30$ per $1,000m^3$, a price within the range given in *Tilmant et al.* (2012) for the valuation of environmental benefits.

3 Challenges to producing a steady-state solution

3.1 Evidence: algorithmic chaos

Robustness of simulation results to small variations in inputs is a crucial issue in modeling, e.g., hydrological modeling (e.g. *Schulz et al.*, 1999; *Perrin et al.*, 2003). Yet, this section document how with SDDP, differences in average output trajectories can sometimes be extremely large for minimal variations in input.

This phenomenon is demonstrated through a minimal change in the initial storage of the two major reservoirs between two scenarios $S1$ and $S2$. All other SDDP inputs, including the $K = 20$ hydrological backward openings and the $M = 30$ simulation sequences, are identical. In scenario $S1$, initial storage is set at 60% of live storage capacity in all reservoirs. In scenario $S2$, the initial quantity stored in the system is the same, but initial storage is increased by $0.1km^3$ (0.06% of total storage capacity) at Kariba (node 4) and decreased by the same amount at Cahora Bassa (node 6). This modification is three orders of magnitude below both the live and total storage capacity of these reservoirs. Other reservoirs have the same initial storage in both scenarios, so that the only differences between scenarios $S1$ and $S2$ are summarized in Table 3.

SDDP is run for both $S1$ and $S2$, while being programmed to stop when the upper bound $\overline{Z}$ falls within the 95% confidence interval of the lower boud $\underline{Z}$. $S1$ and $S2$ converged after 8 and 10 iterations, respectively. Cuts from SDDP runs of scenarios $S1$ and $S2$ are then applied to a hundred synthetic ten-year time series representative of current regional climate, a sample large enough to prevent a large bias in the results. After that, storage trajectories and mean storage values obtained for the same simulated sequences can be compared in Figure 3 for Kariba. The difference in mean storage gets two orders of magnitude greater than the initial difference in storage, a remark that holds true for some of the individual storage trajectories.

Moreover, these results are not an artefact of scenarios converging after different numbers of iterations; rather they arise during the SDDP run. At each iteration, a sequence of states $(\mathbf{x}_t)$ is added to the pool of sample states, and these states include storage values. $S1$ and $S2$ use the same sequence of historical flows to generate new sequences of sample states; yet, the maximal difference between

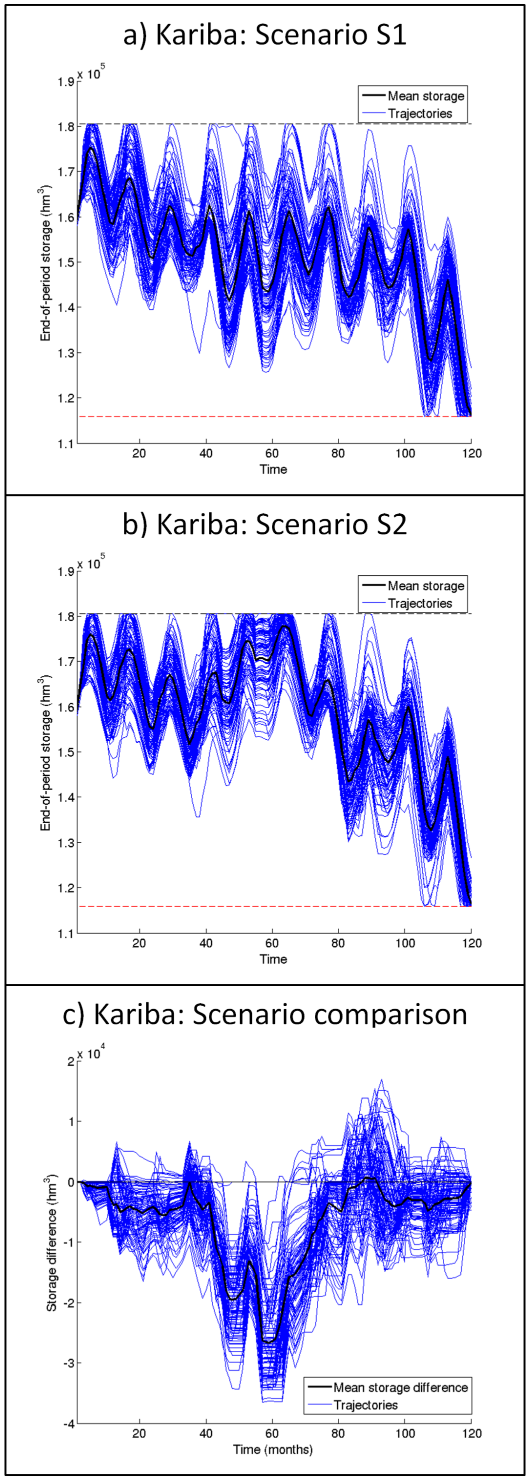

Figure 3: Comparaison of end-of-period storage levels at Kariba for the base-case scenario ($S1$), and scenario $S2$ in which $100hm^3$ of initial storage have been transferred upstream from Cahora Bassa (node 6) to Kariba (node 4).

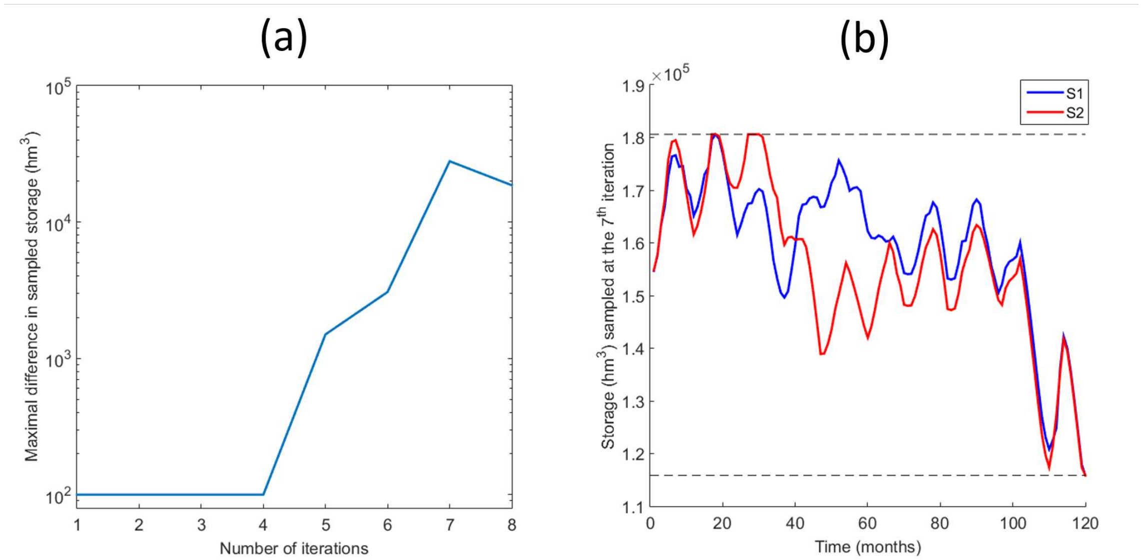

Figure 4: Differences between $S1$ and $S2$ in state sampling during the SDDP run. (a) Maximal difference between two sampled storage values at Kariba for each iteration (log scale). (b) Sampled storage trajectories at Kariba during iteration 7 of the respective SDDP runs.

sampled storage values at Kariba starts increasing markedly starting at iteration 5, and by iteration 7 storage trajectories are very different (Figure 4).

Minimal changes in sample states between runs for $S1$ and $S2$ gradually translate into different approximations of the benefit-to-go function. In the presence of MNOS, these different approximations can lead to different operation policies, and therefore into larger changes in state trajectories – and in sampled states. Using an analogy with physics where the term of "chaos" is used to describe sensitivity to initial conditions (*Rössler*, 1976), this paper calls "algorithmic chaos" the phenomenon observed in Figures 3 and 4. Indeed, "chaotic" behavior is not a staple of large scale water resources systems, but appears along with iterations of the SDDP algorithm.

3.2 Consequences

Very small differences in inputs leading to excessively large difference in simulation outputs are not an automatic occurrence with SDDP. Rather, the case documented in Section 3.1 proves that such situations can happen, and therefore, that limited data availability can make SDDP results hard to interpret. For instance, comparing results from two distinct scenarios can prove challenging.

Besides, in Figure 3, storage probability distributions change from one year to the next. This issue applies to all simulated quantities, including economic outputs from the exploitation of water resources. For instance, judging from this year-to-year variability of mean simulated power production (Figure 5), the actual probability distribution of annual power production at Kariba cannot be deduced from the empirical simulation results of any given year, nor of any sequence of several years. By using cuts generated by SDDP runs from both $S1$ and $S2$ on samples of synthetic time series of different sizes, we also verified that results from Figures 3 and 5 do not depend on the number of such time series used

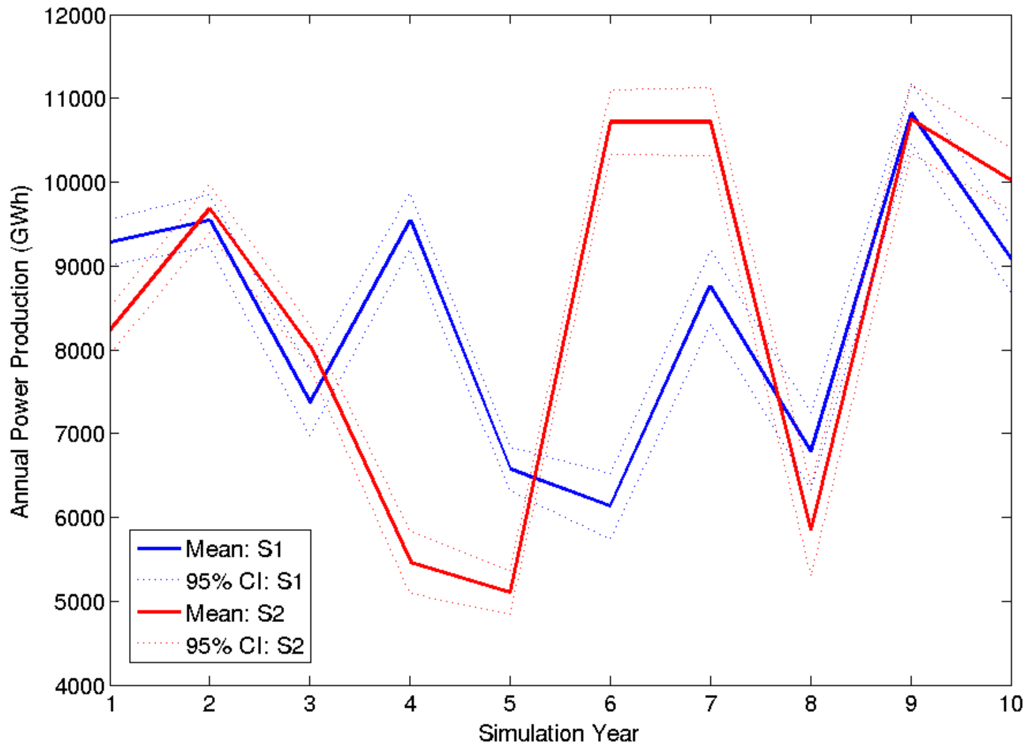

Figure 5: Simulated mean annual power production for each simulation year at Kariba for scenarios $S1$ and $S2$, with its 95% confidence interval (CI).

to get the probability distribution of simulation results. What is more, a run with the same inputs as $S1$, but with $K = M = 100$ also yielded a strong year-to-year variability of outputs, further suggesting that hydrological sampling is not to blame for the failure to produce steady-state conditions.

The case presented in this section uses a "flat" benefit-to-go function $\nu(x^m_{T+1})$ at the end of the algorithm's horizon, which means that in equation (3), we have $\mathbf{a}_{T+1} = \beta_{T+1} = 0$. In other words, storage is set to have no value at the end of a simulation. This provides an incentive for SDDP to take decisions that tend to shortsightedly empty reservoirs near the end of the horizon in order to produce electricity. This end-effect adversely affects the year-periodicity of the results. However, test results show that applying cuts from a first run at the final period, therefore artificially increasing the length of the planning period, does not fix that issue.

Results from this section give insights as to why the problem of finding a representative year cannot be fixed merely by fine-tuning the boundary conditions at the beginning and end of the planning period. On one hand, the presence of algorithmic chaos means that modifying these limit conditions may have unpredictable effects on simulation results. On the other hand, both algorithmic chaos and the strong year-to-year variability of simulated power production levels in Figure 5 show that issues are not caused by inappropriate boundary conditions. Otherwise, the effect of these conditions would be felt considerably less during the middle years of the simulation, e.g., years 5 and 6.

3.3

4 Producing a steady-state near-optimal solution

4.1 Year-periodic re-optimization

In order to avoid the challenges due to the presence of different decision rules from year to year, year-periodic re-optimization (YPRE) is introduced in order to simulate management decisions every year using the cuts generated at a single year denoted y. For instance, the cuts from January of year y shall be used for every January of the simulation period. The same is true from the February cuts from year y, and so on. The final boundary condition is provided by the cuts from December of year y.

The performance of YPRE with the cuts from year y should not be worse than simulation results from SDDP, and this comparison should not be biased by the initial and final conditions. Thus after an SDDP run, the performance of YPRE and that of SDDP can only be compared over stages that are towards the middle of the simulation period (upper half Figure 6). Within this exploitable period, any sequence of consecutive cuts spanning a year is a suitable candidate for y. This is why cuts from a suitable stage – e.g., a December month in the case of a monthly time step – can serve as a final condition for a second SDDP run, thus extending the exploitable period (lower half Figure 6). This makes the exploitable period larger and the comparison between SDDP and YPRE results more meaningful.

The procedure for YPRE is presented by Figure 7. Finding the year y which cuts should be used for YPRE is not an objective of this work. One can select y that maximizes the expected value of system-wide benefits. Yet, this supposes that all sequences of cuts spanning a year should be tested, which can be computationally expensive. For instance, let us consider a ten-year simulation period with a monthly time step ($T = 120$), and assume that the exploitable period starts at the beginning of the fourth year ($\tau = 37$ on Figure 6) and ends at T. Then, there are seventy-three candidates for y. This number can be reduced to seven by considering only a subset of these candidates, e.g., any set of consecutive periods spanning a year. Yet, y can also be chosen through a totally different rationale. For instance, in the backward phase of SDDP, inflows are sampled using a PAR(1) model while using the historical flow record as antecedent flow. Therefore, one can select a year with average historical inflows because sampling may be more representative of hydroclimatic conditions.

Regardless of the methodology for selecting year y, results are exploited after verifying either that YPRE with that year outperforms the reference SDDP run, or that there is no such possible choice for y. Application to a few different case-studies suggests that it is generally possible to find simulation years such that the YPRE outperforms the sequence of all cuts generated by SDDP. An intuitive explanation to this is that the cuts generated by SDDP for different simulation years are unequal approximations of the optimal expected value of future benefits. Then, operating the system by using the same superior cuts at all simulation years may be better than using a mix of superior and inferior cuts as when simulating the system in the forward (re-optimization) phase of SDDP.

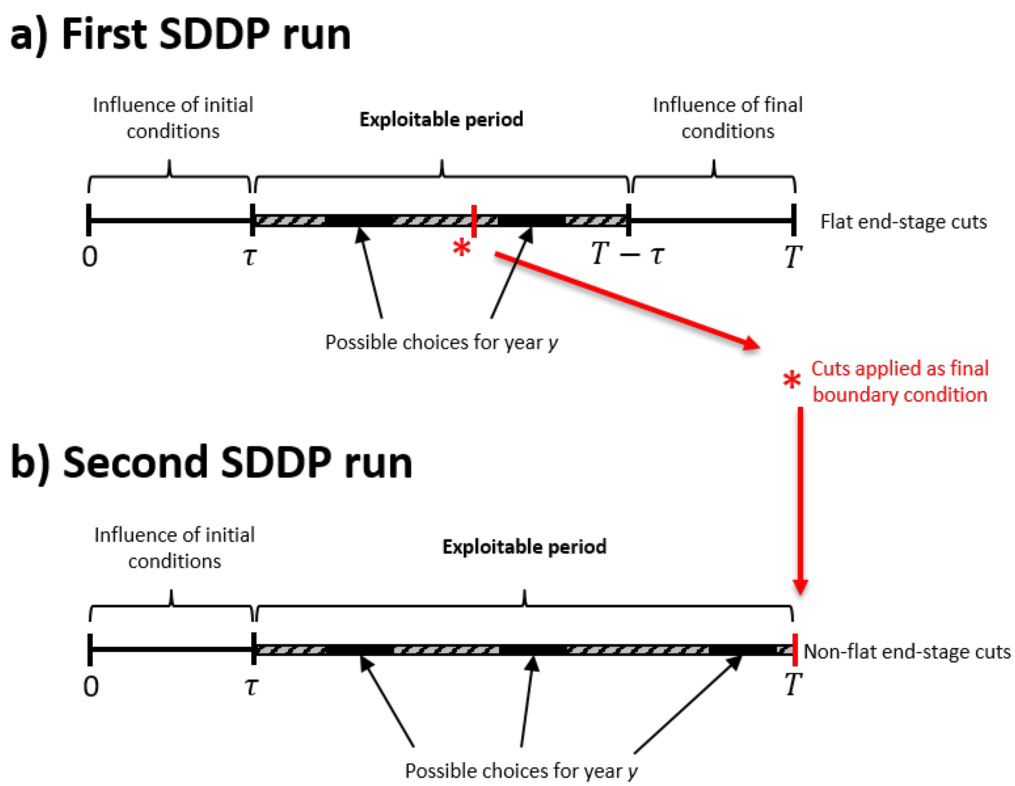

Figure 6: Extension of the exploitable period of an SDDP run.

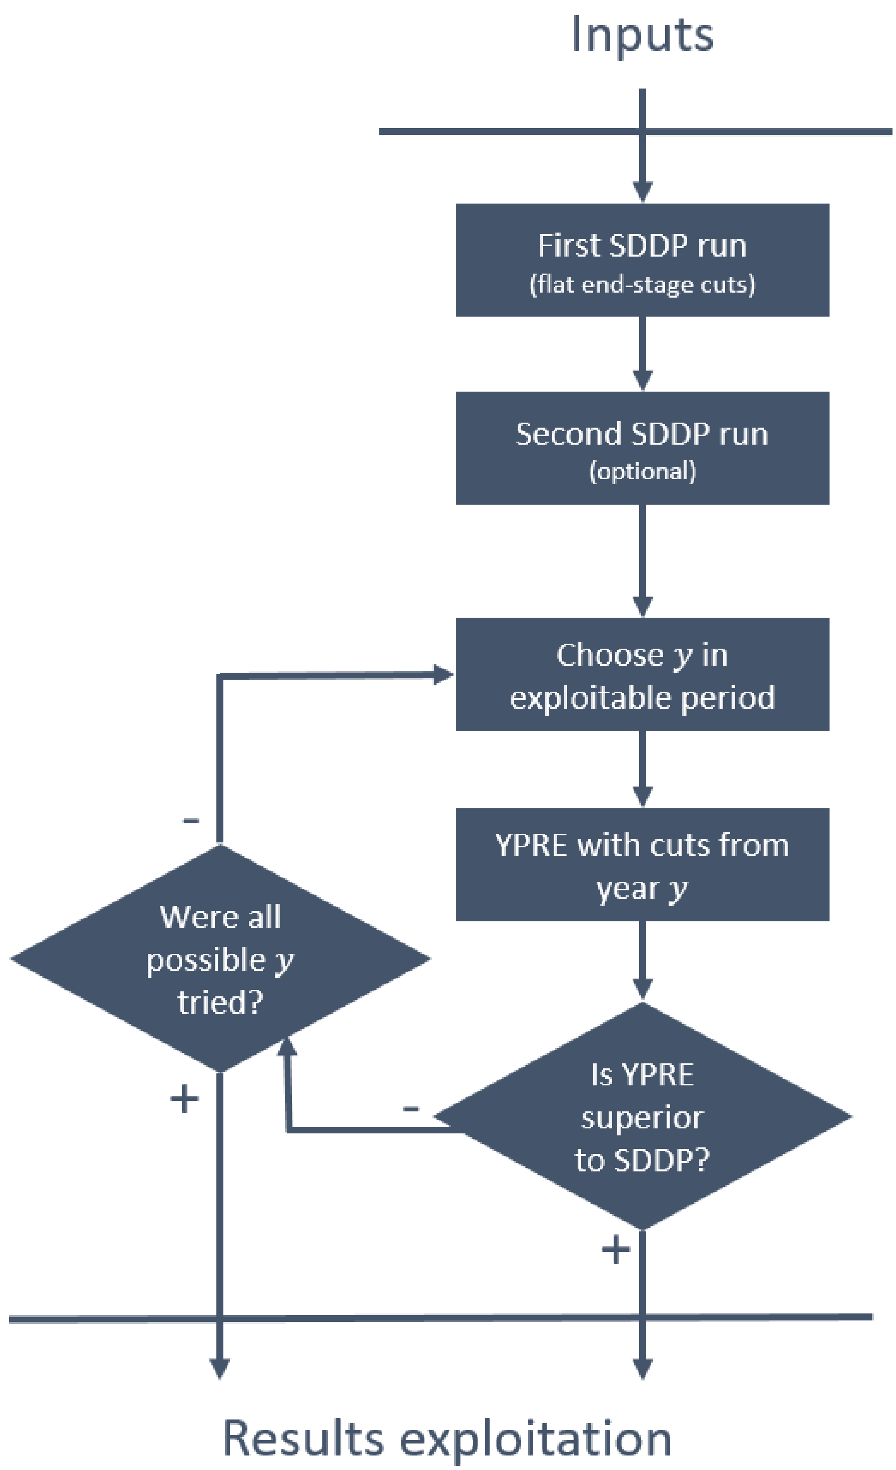

Figure 7: Procedure for year-periodic re-optimization (YPRE).

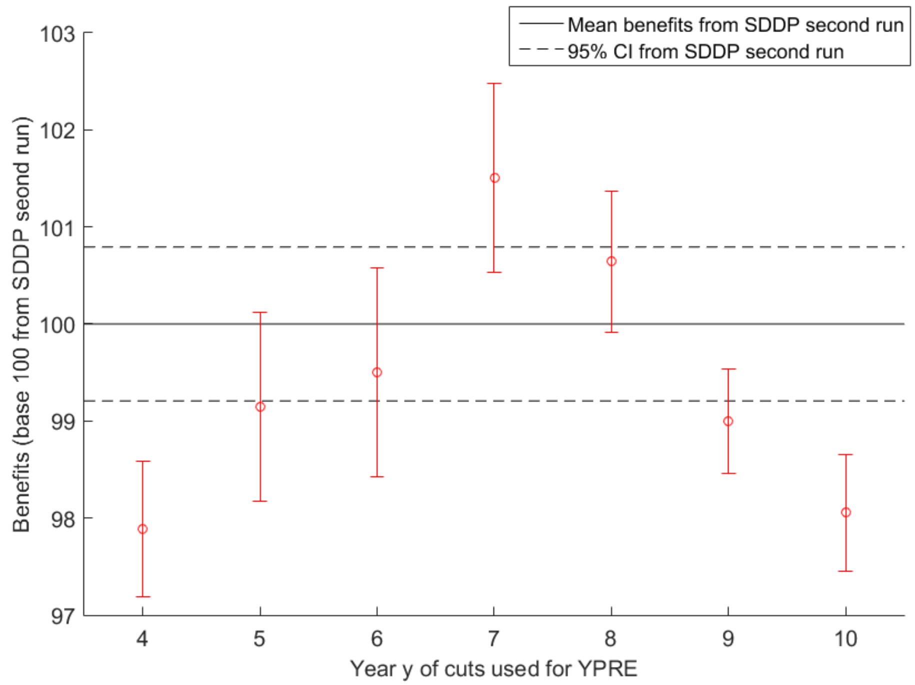

Figure 8: System-wide operational benefits over the last eight simulation years using YPRE, for each of these simulation years. Results from the second SDDP run are used as a base 100.

4.2 Application to the Zambezi case-study

The procedure for YPRE using SDDP is applied to the Zambezi case-study with the same settings as in Section 3, using the same hundred synthetic time-series as in Section 3.1 to test the performance of different sets of cuts. First, two SDDP runs are performed consecutively as to have an exploitation period that is as large as possible: it starts at the beginning of the fourth year and extends to the final stage $T = 120$. Then, y is chosen as the civil year whose cuts yield the best performances, so the performance of each year is measured on the last seven years on the simulation period. System performance varies depending on the simulation year y used for YPRE, and is sometimes greater than with the cuts from the second SDDP run (Figure 8).

Cuts from the simulation year $y = 7$ outperform those of other civil years as well as those from the second SDDP run, so they are used in order to assess how YPRE addresses the issues outlined in Section 3. After a transition period due to initial storage conditions throughout the basin, the probability distribution of storage at Kariba becomes year-periodic (Figure 9), and so is any statistic derived from that distribution, such as the empirical value of the mean storage level. This is the case for all other simulated hydro-economic quantities. For instance, the mean annual power production at Kariba stabilizes around 7,500 GWh, and at almost 65 TWh for the whole basin (Figure 10). Those results are consistent with recent studies on the Zambezi (*World Bank*, 2010b). The year-to-year variations of production are within the 95% confidence interval, and they are particularly negligible when compared to the uncertainty associated with limited data availability. Therefore with YPRE,

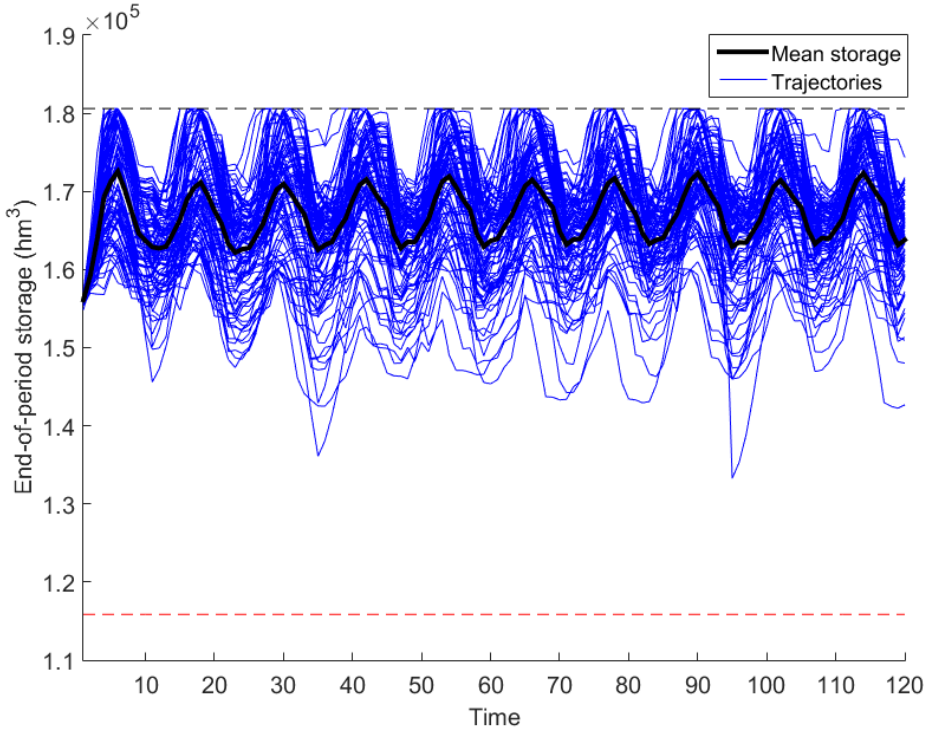

Figure 9: 10-year storage trajectories and mean storage at Kariba after YPRE.

one can use the probability distribution of results from any simulation year towards the end of the simulation period.

5 Discussion

One might wonder whether using cuts from a single year is an ad-hoc solution, since the exploitable period is much longer. Yet, this kind of cut aggregation is actually far from straightforward, which explains why it is outside the scope of this paper, and should rather be addressed by later work. Recall that equation (3), which gives benefit-to-go function F_{t+1} under the form of L inequality constraints written as $F_{t+1} \leq \mathbf{a}^l_{t+1}.\mathbf{x}_{t+1} + \beta^l_{t+1}$. These L inequality constraints must be satisfied simultaneously, but only one of them is active for each state $\mathbf{x}_{t+1}$. In dual programming, what really matters for decision making at a given state is the vector $\mathbf{a}^l_{t+1}$ belonging to the active inequality constraint. Decisions taken equate present and future marginal benefits from water use. Then, even though aggregating different years means increasing the number of inequality constraints, only one of them will be active for each state. Therefore, aggregating is only meaningful if it is done in a way where the active constraint is the one whose coefficients of $\mathbf{a}^l_{t+1}$ are closest to the "true" marginal value of water at that state.

Yet, the constant coefficient β^l_{t+1} largely controls which equation may be active. Since it is the constant part of the future benefits, its value decreases at each stage. Therefore, if one aggregates all cuts without modifying with the β coefficients, only the cuts from the year when these coefficients are

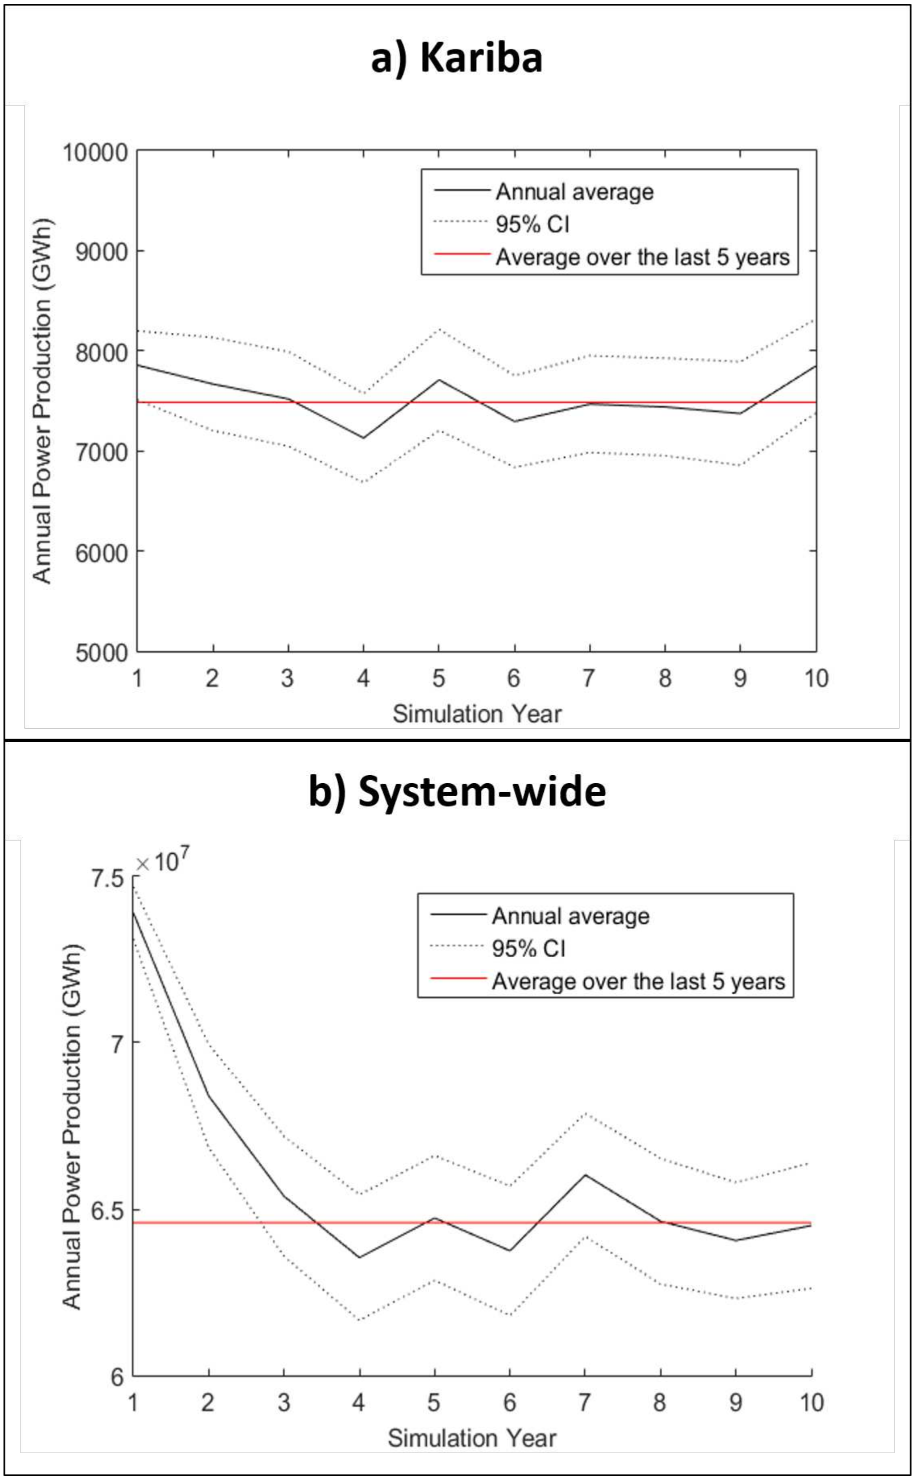

Figure 10: 10-year storage trajectories and mean storage at Kariba after YPRE.

smallest (the last year) will be considered. Relating constant coefficients β from different years with each other is challenging because of the year-to-year irregularity which we demonstrate in Section 3, and which is the very reason to propose YPRE in the first place.

Besides, this work has highlighted how the presence of MNOS could pose difficulties when using SDDP to analyze water resources allocation problems with scarce data. However, the presence of MNOS in water resources system may also be viewed as an opportunity. For instance, *Liu et al.* (2011) state that the search for MNOS offers more flexibility than that of a single optimum; besides, the search for a single optimum ought to be replaced by that of a limited number of alternatives that can later be compared, possibly using criteria that were not accounted for in the objective function to be optimized (*Loucks and van Beek*, 2005). For instance, the presence of MNOS in a deterministic hydropower production problem enables the exploration of the solutions to find which are best for the reliability and safety of the production units (*Liu et al.*, 2012). In this work, a NOS is given by a sequence of Bender's cuts that covers a year, and that lies within a desired range of any other solution, e.g., one percent. This tentative definition can be viewed as a step towards finding a set of MNOS with a desired tolerance in an explicit stochastic context, even though actually finding all the solutions lies outside the scope of this work.

A connection between the presence of MNOS and the concept of equifinality has been suggested in previous research (*Liu and Cai*, 2010). Equifinality is a term that exists in hydrological modeling (*Beven*, 1993; *Schulz et al.*, 1999; *Beven*, 2006) to describe how different model parameters lead to similar goodness-of-fit against observed data during the calibration phase. In that case, parameter values chosen by the modelers can be viewed as decision variables, while goodness-of-fit is quantified through an objective to optimize. This is in a sense very similar to the case of MNOS described in this paper, since these are described by benefit-to-go functions that serve as a basis for taking decisions. These functions lead to very similar values of the objective over the whole horizon $[1, T]$, even though the sequence of decision rules can lead to very different simulated trajectories in the same conditions.

6 Concluding remarks

In the presence of MNOS, approximate optimization algorithms such as SDDP can have an unstable and unpredictable behavior if the problem cannot be adequately parametrized. This was illustrated through a model of the Zambezi River Basin. SDDP proved unstable because it contradicted the trivial prediction that simulation results had to be year-periodic if the inputs were. It proved unpredictable through the discovery of algorithmic chaos. These two traits are related to the fact that SDDP relies on a state sampling that is unpredictable itself, and leads to explore different near-optimal policies in a single run.

YPRE, obtained by iterating cuts from simulation year y over the whole planning period, fixes this instability issue. It proposes policies that are unchanged from one year to the next, and that therefore lead to periodic probability distribution of all system-related quantities. Results suggest that it is relatively straightforward to select cuts that are among the better near-optimal policies, because YPRE using these cuts can outperform the simulation results obtained using SDDP. YPRE should not be thought of as the only possible way to circumvent the difficulties posed by limited data availability.

For instance, the objective can be made more convex by specifying convex demand curves for water demand and hydropower production. However, this would imply for modelers to make more complex assumptions which they would have to then justify.

Acknowledgements

The authors wish to thank the three reviewers and the associate editor, whose constructive comments helped improve this work in a significant way. Data presented in this manuscript can be accessed by contacting Pr. Amaury Tilmant at the email address amaury.tilmant@gci.ulaval.ca.

References

Archibald, T. W., C. S. Buchanan, K. I. M. McKinnon, and L. C. Thomas (1999), Nested benders decomposition and dynamic programming for reservoir optimisation, *The Journal of the Operational Research Society*, *50*(5), 468–479.

Arjoon, D., Y. Mohamed, Q. Goor, and A. Tilmant (2014), Hydro-economic risk assessment in the eastern nile river basin, *Water Resources and Economics*, *8*(0), 16 – 31, doi: http://dx.doi.org/10.1016/j.wre.2014.10.004, economics of Transboundary Water Resources Management.

Beven, K. (1993), Prophecy, reality and uncertainty in distributed hydrological modelling, *Advances in Water Resources*, *16*(1), 41 – 51, doi:http://dx.doi.org/10.1016/0309-1708(93)90028-E, research Perspectives in Hydrology.

Beven, K. (2006), A manifesto for the equifinality thesis, *Journal of Hydrology*, *320*, 18 – 36, doi:http://dx.doi.org/10.1016/j.jhydrol.2005.07.007, the model parameter estimation experiment {MOPEX} {MOPEX} workshop.

Bezerra, B., A. Veiga, L. Barroso, and M. Pereira (2012), Assessment of parameter uncertainty in autoregressive streamflow models for stochastic long-term hydrothermal scheduling, in *Power and Energy Society General Meeting, 2012 IEEE*, pp. 1–8, doi:10.1109/PESGM.2012.6345322.

Cai, X., D. McKinney, and L. S. Lasdon (2002), A framework for sustainability analysis in water resources management and application to the syr darya basin, *Water Resources Research*, *38*(6), 1085, doi:10.1029/2001WR000214.

Davidsen, C., S. Liu, X. Mo, D. Rosbjerg, and P. Bauer-Gottwein (2016), The cost of ending groundwater overdraft on the north china plain, *Hydrology and Earth System Sciences*, *20*(2), 771–785, doi:10.5194/hess-20-771-2016.

Gjelsvik, A., B. Mo, and A. Haugstad (2010), Long- and medium-term operations planning and stochastic modelling in hydro-dominated power systems based on stochastic dual dynamic programming, in *Handbook of Power Systems I*, edited by S. Rebennack, pp. 33–55, Springer-Verlag Berlin Heidelberg, doi:10.1007/978-3-642-02493-1 2.

Goor, Q., C. Halleux, Y. Mohamed, and A. Tilmant (2010), Optimal operation of a multipurpose multireservoir system in the eastern nile river basin, *Hydrology and Earth System Sciences*, *14*, 1895–1908, doi:10.5194/hess-14-1895-2010.

Goor, Q., R. Kelman, and A. Tilmant (2011), Optimal multipurpose-multireservoir operation model with variable productivity of hydropower plants, *Journal of Water Resources Planning and Management*, *137*(3), 258–267, doi:10.1061/(ASCE)WR.1943-5452.0000117.

Harou, J. J., M. Pulido-Velazquez, D. E. Rosenberg, J. Medellín-Azuara, J. R. Lund, and R. E. Howitt (2009), Hydro-economic models: Concepts, design, applications, and future prospects, *Journal of Hydrology*, *375*, 627–643, doi:10.1016/j.jhydrol.2009.06.037.

Homem-de Mello, T., V. de Matos, and E. Finardi (2011), Sampling strategies and stopping criteria for stochastic dual dynamic programming: a case study in long-term hydrothermal scheduling, *Energy Systems*, *2*(1), 1–31, doi:10.1007/s12667-011-0024-y.

Kristiansen, T. (2004), Financial risk management in the electric power industry using stochastic optimization, *Advanced Modeling and Optimization*, *6*(2), 17–24.

Liu, P., and X. Cai (2010), *Deriving Near-Optimal Solutions to Deterministic Reservoir Operation Problems*, chap. 267, pp. 2602–2618, World Environmental and Water Resources Congress 2010, doi:10.1061/41114(371)268.

Liu, P., X. Cai, and S. Guo (2011), Deriving multiple near-optimal solutions to deterministic reservoir operation problems, *Water Resources Research*, *47*, W08506, doi:10.1029/2011WR010998.

Liu, P., T.-D. Nguyen, X. Cai, and X. Jiang (2012), Finding multiple optimal solutions to optimal load distribution problem in hydropower plant, *Energies*, *5*(5), 1413, doi:10.3390/en5051413.

Liu, P., L. Li, G. Chen, and D. E. Rheinheimer (2014), Parameter uncertainty analysis of reservoir operating rules based on implicit stochastic optimization, *Journal of Hydrology*, *514*, 102 – 113, doi:http://dx.doi.org/10.1016/j.jhydrol.2014.04.012.

Loucks, D. P., and E. van Beek (2005), *Water Resources Systems Planning and Management*, UNESCO.

Maceira, M., and J. Damázio (2004), The use of par(p) model in the stochastic dual dynamic programming optimization scheme used in the operation planning of the brazilian hydropower system, in 8^{th} *International Conference on Probabilistic Methods Applied to Power Systems, Iowa State University, Ames, Iowa, September 12-16*, pp. 397–402.

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

Pereira, M. (1989), Optimal stochastic operations scheduling of large hydroelectric systems, *International Journal of Electrical Power & Energy Systems*, *11*(3), 161 – 169, doi:http://dx.doi.org/10.1016/0142-0615(89)90025-2.

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

Pereira-Cardenal, S. J., B. Mo, N. D. Riegels, K. Arnbjerg-Nielsen, , and P. Bauer-Gottwein (2015), Optimization of multipurpose reservoir systems using power market models, *Journal of Water Resources Planning and Management*, *141*(8), 04014,100, doi:10.1061/(ASCE)WR.1943-5452.0000500.

Perrin, C., C. Michel, and V. Andréassian (2003), Improvement of a parsimonious model for streamflow simulation, *Journal of Hydrology*, *279*(1 - 4), 275 – 289, doi:http://dx.doi.org/10.1016/S0022-1694(03)00225-7.

Rössler, O. (1976), Different types of chaos in two simple differential equations, *Zeitschrift fur Naturforschung - Section A Journal of Physical Sciences*, *31*(12), 1664–1670, doi:10.1515/zna-1976-1231.

Rotting, T., and A. Gjelsvik (1992), Stochastic dual dynamic programming for seasonal scheduling in the norwegian power system, *Power Systems, IEEE Transactions on*, *7*(1), 273–279, doi:10.1109/59.141714.

Schulz, K., K. Beven, and B. Huwe (1999), Equifinality and the problem of robust calibration in nitrogen budget simulations, *Soil Science Society of America Journal*, *63*(6), 1934–1941, cited By 40.

Tejada-Guibert, J. A., S. A. Johnson, and J. R. Stedinger (1993), Comparison of two approaches for implementing multireservoir operating policies derived using stochastic dynamic programming, *Water Resources Research*, *29*(12), 3969–3980, doi:10.1029/93WR02277.

Tilmant, A., and R. Kelman (2007), A stochastic approach to analyze trade-offs and risks associated with large-scale water resources systems, *Water Resources Research*, *43*, W06425, doi:10.1029/2006WR005094.

Tilmant, A., and W. Kinzelbach (2012), The cost of noncooperation in international river basins, *Water Resources Research*, *48*, W01,503,doi:10.1029/2011WR011,034, doi:10.1029/2011WR011034.

Tilmant, A., D. Pinte, and Q. Goor (2008), Assessing marginal water values in multipurpose multireservoir systems via stochastic programming, *Water Resources Research*, *44*, W12,431, doi:10.1029/2008WR007,024.

Tilmant, A., L. Beevers, and B. Muyunda (2010), Restoring a flow regime through the coordinated operation of a multireservoir system: The case of the zambezi river basin, *Water Resources Research*, *46*(7), n/a–n/a, doi:10.1029/2009WR008897.

Tilmant, A., W. Kinzelbach, D. Juizo, L. Beevers, D. Senn, and C. Casarotto (2012), Economic valuation of benefits and costs associated with the coordinated development and management of the zambezi river basin, *Water Policy*, *14*, 490–508, doi:10.2166/wp.2011.189.

Wallace, S., and S.-E. Fleten (2003), *Stochastic programming models in energy*, pp. 637–677, (Handbooks in Operations Research and Management Science), doi:10.1016/S0927-0507(03)10010-2.

World Bank (2010a), Basin development scenarios. vol. 2 of the zambezi river basin : a multi-sector investment opportunities analysis., *Tech. rep.*, Washington, DC: World Bank.

World Bank (2010b), State of the basin. vol. 3 of the zambezi river basin : a multi-sector investment opportunities analysis., *Tech. rep.*, Washington, DC: World Bank.

Wu, X., and D. Whittington (2006), Incentive compatibility and conflict resolution in international river basins: A case study of the nile basin, *Water Resources Research*, *42*(2), n/a–n/a, doi:10.1029/2005WR004238.

Zhang, J., P. Liu, H. Wang, X. Lei, and Y. Zhou (2015), A bayesian model averaging method for the derivation of reservoir operating rules, *Journal of Hydrology*, *528*, 276 – 285, doi: http://dx.doi.org/10.1016/j.jhydrol.2015.06.041.