

Water Resources Research®



RESEARCH ARTICLE

10.1029/2021WR030582

Optimizing the Management of Multireservoir Systems Under Shifting Flow Regimes

Vahid Espanmanesh¹ and Amaury Tilmant¹ 

¹Department of Civil Engineering and Water Engineering, University of Laval, Quebec City, QC, Canada

Key Points:

- Incorporating climate variables into the state vector of Stochastic Dual Dynamic Programming yields more efficient reservoir operating policies
- Reservoir operating policies based on climate states naturally hedge against adverse hydrologic conditions
- Climate-derived streamflow forecasts are more reliable and accurate than those obtained from an autoregressive model

Correspondence to:

A. Tilmant,
amaury.tilmant@gci.ulaval.ca

Citation:

Espanmanesh, V., & Tilmant, A. (2022). Optimizing the management of multireservoir systems under shifting flow regimes. *Water Resources Research*, 58, e2021WR030582. <https://doi.org/10.1029/2021WR030582>

Received 28 JUN 2021

Accepted 5 JUN 2022

Author Contributions:

Conceptualization: Amaury Tilmant
Data curation: Vahid Espanmanesh
Formal analysis: Vahid Espanmanesh
Funding acquisition: Amaury Tilmant
Investigation: Vahid Espanmanesh, Amaury Tilmant
Methodology: Vahid Espanmanesh, Amaury Tilmant
Project Administration: Amaury Tilmant
Resources: Amaury Tilmant
Software: Vahid Espanmanesh
Supervision: Amaury Tilmant
Validation: Amaury Tilmant
Visualization: Vahid Espanmanesh
Writing – original draft: Vahid Espanmanesh
Writing – review & editing: Amaury Tilmant

© 2022. The Authors.

This is an open access article under the terms of the [Creative Commons Attribution-NonCommercial-NoDerivs License](https://creativecommons.org/licenses/by/4.0/), which permits use and distribution in any medium, provided the original work is properly cited, the use is non-commercial and no modifications or adaptations are made.

Abstract Over the past few decades, significant research efforts have been devoted to the development of tools and techniques to improve the operational effectiveness of multireservoir systems. One of those efforts focuses on the incorporation of relevant hydrologic information into reservoir operation models. This effort is particularly relevant in regions characterized by low-frequency climate signals, where time series of river discharges exhibit regime-like behavior. Failure to properly capture such regime-like behavior yields suboptimal operating policies, especially in systems characterized by large storage capacity such as large multireservoir systems. Hidden Markov Modeling is a class of hydrological models that can accommodate both overdispersion and serial dependence in time series, two essential hydrological properties that must be captured when modeling a system where the climate is switching between different states (e.g., dry, normal, and wet). In terms of reservoir operation, Stochastic Dual Dynamic Programming (SDDP) is one of the few optimization techniques that can accommodate both system and hydrologic complexity, that is, a large number of reservoirs and diverse hydrologic information. However, current SDDP formulations are unable to capture the long-term persistence of the streamflow process found in some regions. In this paper, we present an extension of the SDDP algorithm that can handle the long-term persistence and provide reservoir operating policies that explicitly capture regime shifts. Using the Senegal River Basin as a case study, we illustrate the potential gain associated with reservoir operating policies tailored to climate states.

1. Introduction

Anthropogenic disturbances of the terrestrial water cycle, climate change, and low frequency climate signals all challenge the stationarity assumption, which has been the backbone of water resources planning techniques during much of the 20th century (Loucks & Van Beek, 2017b). Time series of hydrological observations can exhibit nonstationarity in different forms such as increasing or decreasing trends (Bayazit, 2015; Olsen et al., 1999), shifts in the mean (Fortin et al., 2004; Potter, 1991), and shifts in variance (Coulibaly & Burn, 2004; Whitcher et al., 2002).

Detecting and attributing the source of nonstationarity in hydrological records is essential for the efficient management of water resources systems. This paper focuses on nonstationarity due to climate variability, which must not be confused with climate change. Climate variability is random variation from a long-run average distribution, whereas climate change is a trend or a shift in the long-run distribution (Stedinger & Griffis, 2011). More specifically, we investigate the value of water allocation policies that explicitly capture climate variability through shifting flow regimes. Climate variability is often linked to low frequency climate signals such as El Niño-Southern Oscillation (ENSO), Pacific Decadal Oscillations, and North Atlantic Oscillation, which impact the hydrology in many regions worldwide (Akintug & Rasmussen, 2005; Bracken et al., 2014; Jha et al., 2021).

Observations of long-term fluctuations in time series of hydrological variables such as annual flows have triggered the development of methods to quantify this long-term persistence. Among such investigations, the work of Hurst (1951) stands out. While analyzing the time series of river discharges in the Nile River basin, Hurst found that there is a tendency of wet years to cluster into multi-year wet periods or of dry years to cluster into multi-year drought periods (Koutsoyiannis, 2002). In the late 60th, Baum and his colleagues presented the basic concept of a regime-shifting time series model named Hidden Markov Model (HMM) (Rabiner, 1989). HMMs, also known as Markov mixture models or Markov switching models, is a class of probabilistic models for “labeling” the observations. Rather than focusing on shifts in the mean of a process, HMMs estimate shifts in the state of a process (Whiting et al., 2004). HMMs, which allow the probability distribution of each observation to depend on the unobserved (hidden) state of a Markov chain, can accommodate both overdispersion (when there is a greater

variability in the data set than would be expected based on a given statistical model) and serial dependence, which measures the correlation between two successive variables in time series (Loucks & Van Beek, 2017a). The motivation for adopting this type of model in hydrology lies in that the climate regime can be represented by a state variable that can take only a limited number of values. Consequently, along with a time series of historical river discharges, a HMM considers another time series with discrete climate states. These properties make this method a suitable option for modeling the governing hydrology of a system where the climate is switching between different states, say dry, normal, and wet. HMMs have been applied, for example, to analyze time series of daily precipitation data (Zucchini & Guttorp, 1991) or streamflow observations (Akintug & Rasmussen, 2005; Tan et al., 2017; Turner & Galelli, 2016; Yu et al., 2018).

Capturing those low-frequency shifts is particularly important in water resources systems characterized by large storage capacity because of their ability to move water over extended periods of time means that they could potentially take advantage of climate information in order to hedge against prolonged dry or wet periods (You & Cai, 2008; Zhao et al., 2012). However, when it comes to modeling techniques to optimize reservoir operation, only a few solution approaches can address the complexities associated with the presence of: (a) multiple hydrologic information/processes, (b) multiple, often conflicting, operating objectives, (c) stochastic variables, (d) nonlinear relationships, and (e) trade-offs between immediate and future consequences associated with the operation of the system (Pina et al., 2017).

Dynamic Programming (DP) is a well-established optimization method to solve the reservoir operation problem (Labadie, 2004). To achieve this, DP decomposes the original problem into sub-problems that are solved recursively. Stochastic DP (SDP) is an extension of DP that can accommodate stochastic state variables like reservoir inflows, snowpack, and flow forecasts (Stedinger et al., 1984). The incorporation of more hydrological information into the state vector has the potential to enhance SDP-derived policies and thus improve system efficiency. However, due to the curse of dimensionality, that is, the exponential increase in computational effort and volume associated with adding extra dimensions to the state-space, the number of state variables is often limited to three to four. To handle the uncertainties associated with various hydrologic information, variants of SDP have been proposed such as Bayesian SDP Karamouz et al. (2003) or Sampling SDP (Kelman et al., 1990; Kim & Palmer, 1997; Liu et al., 2020; Mujumdar & Nirmala, 2007; Tang et al., 2010).

Although useful, those variants fall short of capturing the long-term persistence that characterizes time series of river discharges exhibiting a regime-like behavior. Recently, Turner and Galelli (2016) proposed a SDP formulation that can handle long-term persistence through an additional state variable representing climate states. This new formulation, called SDP- Φ , uses a Hidden Markov Model (HMM) to partition the streamflow time series into a small number of discrete regime states such as dry, normal or wet. That study confirms that regime-like behavior is a major cause of suboptimal hedging decisions for reservoirs that are vulnerable to prolonged dry spells. While this SDP- Φ has been successfully implemented on single reservoir operation problems, it cannot be applied to larger systems due to the curse of dimensionality.

Our literature review reveals that only few solutions are available to handle both hydrologic and system complexities. One of them is Direct Policy Search, a methodology that has the capability of handling multiple objectives and sources of uncertainties (Libisch-Lehner et al., 2019). However, as pointed out by Castelletti et al. (2013), those simulation-based optimization methods are computationally demanding when the number of objectives increases, and difficult to parametrize when the water resources system is large and the network complex. Another alternative solutions is Stochastic Dual Dynamic Programming (SDDP) (M. V. Pereira & Pinto, 1991) which belongs to the field of Approximate DP (Powell, 2007). The SDDP algorithm relies on an iterative procedure (i.e., backward optimization and forward simulation) to constructs a locally-accurate approximation of the benefit-to-go function through sampling and Benders' decomposition (M. Pereira & Pinto, 1983). The basic principle is to approximate the benefit-to-go function using a piecewise linear function which is progressively refined by iteratively adding new hyperplanes (cuts) until the approximation is found to be statistically acceptable. As we will see later in Section 2.5, these piecewise linear approximations are constructed from the primal and dual solutions of the one-stage optimization problems as the algorithm progresses backward.

Generally speaking, SDDP uses a multisite periodic autoregressive model (MPAR) to capture the hydrologic uncertainty. This model can represent serial and spatial correlations within a river basin and between different basins as well as seasonality. A number of recent researches have examined ways of improving the built-in

hydrological model. Lohmann et al. (2016) suggests a Spatial PAR model for SDDP that can take spatial information (such as distance between hydro plant) into account. The idea is that inflows into “neighboring” plants are correlated and this correlation can be accounted for by using neighboring inflow series to explain the inflow at a given location. Pritchard (2015) introduces an inflow modeling approach that can overcome probabilistic sampling errors introduced when the fitted continuous inflow model is replaced by a discrete approximation obtained through sampling. Poorsepahy-Samian et al. (2016) use the Box-Cox transformation in the inflow modeling of SDDP while meeting the convexity requirement of the algorithm. Raso et al. (2017) present a stochastic streamflow model with a multiplicative component and a nonuniform time step. Pina et al. (2017) propose an approach to incorporate exogenous hydrological variables (e.g., snow water equivalent, ENSO variables) into the SDDP algorithm through a MPARX model. In Treistman et al. (2020), ENSO variables are rather treated as additional state variables modeled by a Markov chain along with the PAR(p) model for the streamflows. More recently, Mbeutcha et al. (2021) incorporate a moving average component in the built-in hydrological model of SDDP.

However, none of the reported extensions of the SDDP algorithm can handle low-frequency signals. The main focus of this study is to extend the SDDP algorithm so as to determine reservoir operating policies that explicitly capture shifting flow regimes due to climate variability. To achieve this, the state-space vector now includes a climate variable whose transition is governed by a HMM and the cuts approximating the benefit-to-go functions are also classified according to their climate state. The proposed SDDP-HMM model is then implemented on the multipurpose, multireservoir system in the Senegal River basin in West Africa, a region characterized by shifting flow regimes (Faye et al., 2015).

This paper is organized as follows: Section 2 presents the methodology to efficiently capture historical streamflow regime changes through HMM followed by a description of the modified SDDP algorithm to solve the multireservoir operation problem. Section 3 then provides a description of the Senegal River Basin. Section 4 discusses the results. Finally concluding remarks, limitations, and potential extensions are discussed in Section 5.

2. Material and Methods

2.1. Unfolding the Hidden Climate States

In numerous river basins around the world, hydrologic time series display characteristic responses to large scale climate patterns (Akintug & Rasmussen, 2005; Bracken et al., 2014). For example, let us consider the time series of annual streamflows at Bakel in the Senegal River basin (West Africa) from 1904 to 2011 (Figure 1). This time series displays a regime-like behavior with wet periods during [1904–1910], [1915–1937], and [1943–1970], and dry periods during [1910–1915], [1937–1943], and [1970–2011].

For multi-year reservoir operation planning, characterizing these lengthy departures from the mean are essential as they stress the system far more than single wet or dry years (Bracken et al., 2014). Capturing such a long-term persistence therefore can be achieved by considering the climate states associated with flow regimes, something that can be achieved using, for example, Hidden Markov Models (HMM) through labeling the observations (Rabiner, 1989). In other words, when dealing with a time series of river discharges, it means that there exists another time series with discrete climate states that can take only a limited number of values (i.e., dry/wet for 2 states; dry/normal/wet for 3 states) (Figure 2). Denote $\{q_1, q_2, \dots, q_T\}$ the time series of monthly flows and let $\{\Phi_1, \Phi_2, \dots, \Phi_T\}$ be the time series of climate states which can only take Φ possible climate states.

The state variable is unobserved and is accordingly referred to as a hidden variable. The key information that we need to extract from the analysis of time series of river discharges are the transition probabilities from one climate state to another in a single time step. The probabilities $P_t^{i,j}$ of moving from climate state i at time $t - 1$ to climate state j at the next time step (t):

$$P_t^{i,j} = Pr(\Phi_t = j | \Phi_{t-1} = i) \quad \text{with } t = 2, \dots, T \quad \text{and } i, j = 1, 2, \dots, \Phi \quad (1)$$

are the elements of the $(\Phi \times \Phi)$ transition probability matrix P between time $t-1$ and t .

The observed variable q_t is assumed to have been drawn from a probability distribution whose parameters are conditional upon the distinct state at time t such that, when Φ_t is known, the distribution of q_t only depends on the current state Φ_t (Figure 3). A HMM is described by (a) the parameters of the Gaussian distributions, that

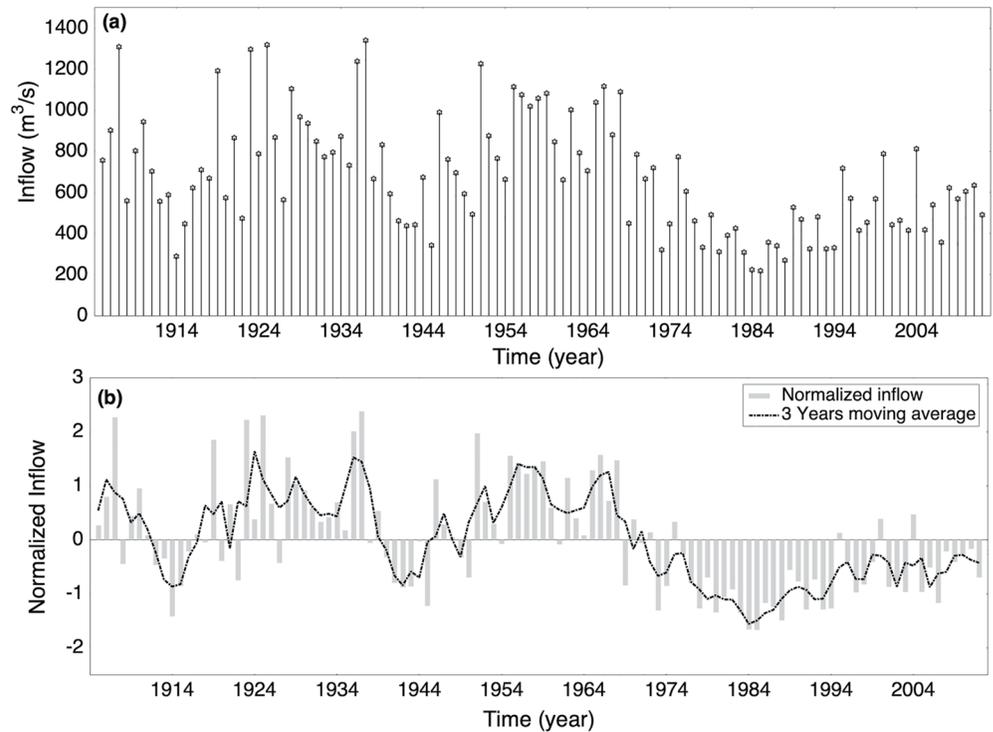


Figure 1. Annual streamflow at Bakel in SRB (1904–2011) (a) Mean annual streamflow, (b) Normalized streamflow (the dashed line represents the 3-year running mean).

is, the mean $\mu = (\mu_1, \mu_2, \dots, \mu_\Phi)$ and standard deviation $\sigma = (\sigma_1, \sigma_2, \dots, \sigma_\Phi)$ associated with Φ states, (b) the $(\Phi \times \Phi)$ transition probability matrix P , and (c) the initial distribution of the Markov chain δ . Consequently, the set of parameters to be estimated is $\theta = \{\mu, \sigma, P, \delta\}$.

Prior to fitting a HMM to the observed sequence (here the time series of monthly inflows), outliers and the seasonal component must be removed. Then, the likelihood of observing that sequence, as calculated under a Φ -state HMM, must be evaluated. In this study, we use the Expectation-Maximization (EM) algorithm, an iterative method for finding the maximum-likelihood estimate of the parameters of an underlying distribution when some of the data are missing. In the context of HMM, the EM algorithm is known as the Baum-Welch algorithm (Welch, 2003) and the hidden climate states are treated as missing data (Bilmes, 1998; Zucchini et al., 2017). The EM algorithm consists of two main phases: an expectation phase called “E step,” followed by a maximization phase called “M step.” Given the current estimate of the HMM parameters θ , the following steps are repeated until acceptable convergence is achieved: The “E step” phase of the algorithm computes the expected value of unobserved data (i.e., hidden climate states) using the current estimate of the parameters and the observed data. The “M step” phase of the algorithm then provides a new estimate of the parameters by using the data from the “E step” phase as if they were actually measured data. These parameters are then used to calculate the distribution of unobserved data in the next “E step” phase of the algorithm. The resulting value of θ is then the stationary point of the likelihood of the observed data.

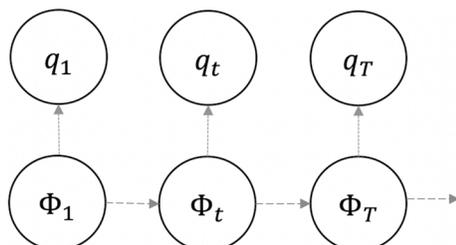


Figure 2. Time series of river discharges and hidden climate states.

Next, we want to determine the sequence of hidden climate states $\{\Phi_1, \Phi_2, \dots, \Phi_T\}$ that has most likely (under the fitted HMM) given rise to the sequence of observations (here the time series of monthly river discharges). In the literature, this is well-known as the decoding problem. In this study, we use the Viterbi algorithm (Viterbi, 1967) to unfold the sequence of hidden climate states (called the Viterbi path). This, in turn, enables us to breaking the observations into Φ distinct regimes and to calculating the probabilities of transitioning between inflow regime states. As we shall see in the following

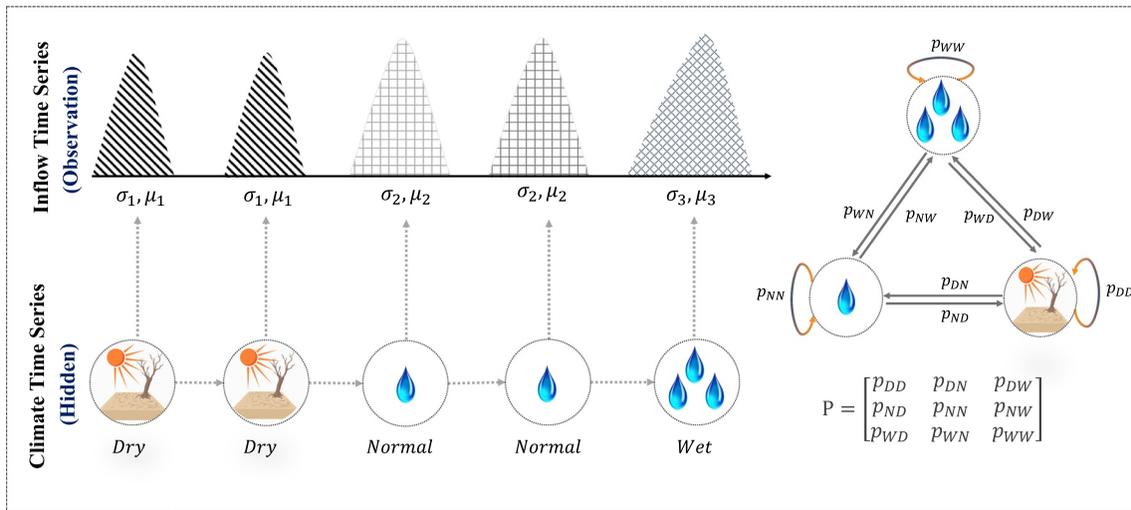


Figure 3. Structure of Hidden Markov Model with 3 hidden climate states Dry (D), Normal (N), and Wet (W).

sections, these transition probabilities can then be used in the SDDP algorithm to derive water allocation policies tailored to discrete climate states such as a dry, normal, or wet.

2.2. The Multireservoir Operation Problem

The operation of a multireservoir system is a multistage decision-making problem. The problem is to determine a sequence of allocation decisions \mathbf{x}_t that maximize the expected net benefit from system operation over a given planning period while meeting operational and/or institutional constraints (Labadie, 2004).

$$Z = \mathbb{E}_{\mathbf{h}_t | \mathbf{h}_{t-1}} \left[\sum_{t=1}^T \alpha_t b_t(\mathbf{S}_t, \mathbf{x}_t) + \alpha_{T+1} \nu(\mathbf{S}_{T+1}, \mathbf{x}_t) \right] \quad (2)$$

where Z is the expected sum of one-stage net benefits from system operation, T is the planning horizon, \mathbf{x}_t is the vector of allocation (decision) variables, $\nu(\cdot)$ is the terminal value function, α_t is the discount factor at stage t , and $\mathbb{E}[\cdot]$ is the expectation operator. In general, the vector of the state variables \mathbf{S}_t includes the volume of water in storage \mathbf{s}_t , as well as hydrologic state variables \mathbf{h}_t . The allocation decisions can be reservoir storages, turbined outflows, spillage losses, and water withdrawals for off-stream uses, $b_t(\cdot)$ represents the aggregated net benefits at stage t , which includes here the net benefit from irrigated agriculture and hydropower generation (the two dominant uses). Note that other operating objectives (e.g., navigation, ecological flows) are handled using penalty coefficients: $b_t(\cdot)$ is penalized when the desired objective is not met.

The most common constraints are:

1. Water balance equations:

$$\mathbf{s}_{t+1} - \mathbf{C}^R(\mathbf{r}_t + \mathbf{l}_t) - \mathbf{C}^I(\mathbf{i}_t) + \mathbf{e}_t(\mathbf{s}_t, \mathbf{s}_{t+1}) = \mathbf{s}_t + \mathbf{q}_t \quad (3)$$

where \mathbf{s}_t is the vector of storage at the beginning of the period t , \mathbf{r}_t is the vector of releases, \mathbf{l}_t is the vector of spills, \mathbf{i}_t is the vector of water withdrawals, \mathbf{q}_t is inflow during time t , \mathbf{C}^R is the reservoir system connectivity matrix, $\mathbf{C}_{n_1, n_2}^R = 1$ (-1) when reservoir n_1 receives (release) water from (to) reservoir n_2 , \mathbf{C}^I is the connectivity matrix for return flows, and \mathbf{e}_t is the vector of evaporation losses.

2. Lower and upper bounds on storage levels:

$$\underline{\mathbf{s}}_{t+1} \leq \mathbf{s}_{t+1} \leq \bar{\mathbf{s}}_{t+1} \quad (4)$$

where $\underline{\mathbf{s}}_{t+1}$, and $\bar{\mathbf{s}}_{t+1}$ are the lower and upper bound on storage level respectively.

3. Limits on reservoir releases:

$$\underline{\mathbf{r}}_t \leq \mathbf{r}_t \leq \bar{\mathbf{r}}_t \quad (5)$$

Maximum release $\bar{\mathbf{r}}_t$ is introduced to account for the maximum carrying capacity downstream of the reservoir, depending on the capacity of the hydraulic structures (e.g., penstock or spillway). Minimum allowable release $\underline{\mathbf{r}}_t$ is given to maintain a desired downstream minimum flow for water quality, navigation, and etc.

4. Water withdrawals can be limited by the pumping station or channel capacity:

$$\underline{\mathbf{i}}_t \leq \mathbf{i}_t \leq \bar{\mathbf{i}}_t \quad (6)$$

As indicated earlier, several features make the reservoir operation problem Equations 2–6 computationally challenging to solve with optimization techniques: the presence of stochastic variables, multiple, often conflicting, objectives, and nonlinear functions such as the hydropower production function, evaporation, and other losses (Rani & Moreira, 2010). Karamouz et al. (2003) suggests a time decomposition approach to deal with the complexity of the reservoir optimization problem. A time decomposition approach breaks down the original problem into long, mid, and short-term planning periods, each having a specific model. Long-term planning models (e.g., monthly time step and over-year planning horizons) are run first and provide strategic and tactical policy information. This information (e.g., future marginal water value and storage levels) then constitutes the boundary conditions for the mid- and short-term optimization tools (e.g., weekly to daily time horizons), whose outputs finally feed real-time operation models. This study focuses on the long-term reservoir operation problem and seeks to identify steady-state, monthly, reservoir operating policies that explicitly capture regime shifts.

2.3. One-Stage SDDP Problem

SDDP solves the optimization problem Equations 2–6 by decomposing it into a sequence of one-stage problems that are solved recursively. Let F_t be the benefit-to-go function (BTF) from stage t to the end of the planning period T associated with decision \mathbf{x}_t , \mathbf{h}_t be the vector of hydrological state variables, the one-stage SDDP optimization problem (OSOP) is a linear program:

$$F_t(\mathbf{s}_t, \mathbf{h}_t) = \max_{\mathbf{x}_t} [\{b_t(\cdot) + \alpha_{t+1} F_{t+1}\}] \quad (7)$$

Subject to constraints Equations 3–6, as well as:

1. Approximation of the hydropower production functions:

$$\begin{cases} \hat{P}_t - \boldsymbol{\psi}^1 \mathbf{s}_{t+1}/2 - \boldsymbol{\omega}^1 \mathbf{r}_t \leq \boldsymbol{\delta}^1 + \boldsymbol{\psi}^1 \mathbf{s}_t/2 \\ \vdots \\ \hat{P}_t - \boldsymbol{\psi}^H \mathbf{s}_{t+1}/2 - \boldsymbol{\omega}^H \mathbf{r}_t \leq \boldsymbol{\delta}^H + \boldsymbol{\psi}^H \mathbf{s}_t/2 \end{cases} \quad (8)$$

where H is the number of planes approximating the true non-linear hydropower functions, $\boldsymbol{\omega}$, $\boldsymbol{\delta}$, and $\boldsymbol{\psi}$ are the parameters derived from the corresponding convex hulls (Goor et al., 2010).

2. Benefit-to-go F_{t+1} is now a scalar bounded from above by the following inequalities (cuts):

$$\begin{cases} F_{t+1} - \boldsymbol{\varphi}_{t+1}^1 \mathbf{s}_{t+1} \leq \boldsymbol{\gamma}_{t+1}^1 \mathbf{h}_t + \boldsymbol{\beta}_{t+1}^1 \\ \vdots \\ F_{t+1} - \boldsymbol{\varphi}_{t+1}^L \mathbf{s}_{t+1} \leq \boldsymbol{\gamma}_{t+1}^L \mathbf{h}_t + \boldsymbol{\beta}_{t+1}^L \end{cases} \quad (9)$$

where L is the number of cuts (hyperplanes), $\boldsymbol{\varphi}_{t+1}$, $\boldsymbol{\gamma}_{t+1}$, and $\boldsymbol{\beta}_{t+1}$ are $1 \times n$ vectors of cut parameters calculated at stage $t + 1$ (Tilmant et al., 2008), and \mathbf{h}_t typically includes the natural inflows observed during the last p periods (\mathbf{q}_{t-1} , \mathbf{q}_{t-2} , ..., \mathbf{q}_{t-p}).

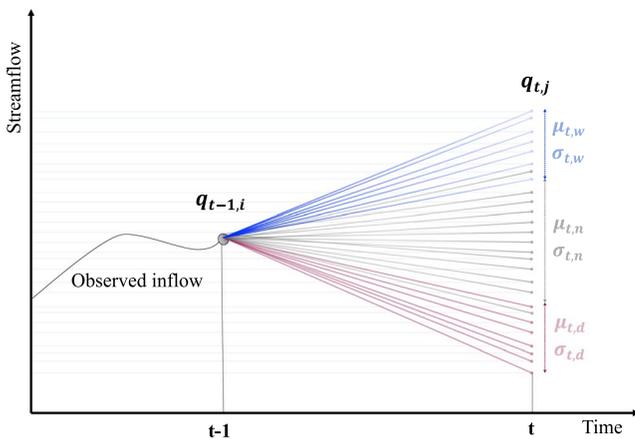


Figure 4. Schematization of inflow modeling considering 3 climate states: dry (d), normal (n), and wet (w).

2.4. Incorporating the Hidden Climate States Into SDDP

A common approach to capture the inflow uncertainty in SDDP is through a built-in multi-site periodic autoregressive model MPAR. The parameters of the MPAR model are estimated from historical flow records (Pina et al., 2017). Suppose the periodic process is modeled by an MPAR of order p - MPAR(p) - then at each site n , the incremental inflows can be derived from:

$$\left(\frac{\mathbf{q}_t - \boldsymbol{\mu}_{\mathbf{q}_t}}{\boldsymbol{\sigma}_{\mathbf{q}_t}} \right) = \sum_{i=1}^p \phi_{t,i} \left(\frac{\mathbf{q}_{t-i} - \boldsymbol{\mu}_{\mathbf{q}_{t-i}}}{\boldsymbol{\sigma}_{\mathbf{q}_{t-i}}} \right) + \boldsymbol{\epsilon}_t \quad (10)$$

where \mathbf{q}_t is the incremental flow at time t , $\boldsymbol{\mu}_{\mathbf{q}_t}$ is periodic mean of \mathbf{q}_t , $\boldsymbol{\sigma}_{\mathbf{q}_t}$ is standard deviation of \mathbf{q}_t , ϕ_t is autoregressive parameter of the periodic model, and $\boldsymbol{\epsilon}_t$ is a time dependent-spatially correlated stochastic noise.

Incorporating the (hidden) climate state variables into the state space vector of SDDP offers the potential to improve the operating policies derived by SDDP since the long-term persistence will be captured. As mentioned earlier in Section 2.1, the historical flow record can be broken into Φ distinct classes by a fitted HMM, providing Φ classes of inflows for each month (e.g., a

HMM with 3 states, say dry, normal, and wet, yields 3 classes of inflow for each month). Instead of having a single time series of historical observations, the HMM now provides another time series with discrete climate states. At stage $t - 1$, assume (a) that $\mathbf{q}_{t-1,i}^o$ is sampled from i th climate state and (b) that, for notational simplicity, an autoregressive model of order one is adjusted:

$$\left(\frac{\mathbf{q}_t - \boldsymbol{\mu}_{\mathbf{q}_{t,j}}}{\boldsymbol{\sigma}_{\mathbf{q}_{t,j}}} \right) = \phi_t \left(\frac{\mathbf{q}_{t-1,i}^o - \boldsymbol{\mu}_{\mathbf{q}_{t-1,i}}}{\boldsymbol{\sigma}_{\mathbf{q}_{t-1,i}}} \right) + \boldsymbol{\epsilon}_t \quad (11)$$

where $\boldsymbol{\mu}_{\mathbf{q}_{t,j}}$, and $\boldsymbol{\sigma}_{\mathbf{q}_{t,j}}$ are respectively the vectors of periodic means and standard deviations of j th climate state (Figure 4). In the backward phase of SDDP, at the beginning of stage t , we can use the modified hydrological model to (a) generate the K inflow scenarios $q_{t,j}$ from the sampled previous flow $q_{t-1,i}$ and (b) analytically calculate the cuts' parameters for stage $t - 1$ (see Appendix A). Note that the K inflow scenarios correspond to an ensemble forecast with K members and with a forecast horizon of only one period (here 1 month). These inflow scenarios are also called backward openings in SDDP (M. V. Pereira & Pinto, 1991). In the forward phase, we can either use the synthetically generated scenarios or historical records to simulate the system over the planning period. The proposed approach can therefore handle both the short-term and the long-term persistence, the former through autoregressive models, and the latter via the fitted HMM.

2.5. SDDP Algorithm With Hidden Climate States

When hidden climate states are included in the state-space vector of SDDP, the objective function of the corresponding OSOP becomes:

$$F_{t,i}(\mathbf{s}_t, \mathbf{q}_{t-1,i}) = \max_{\mathbf{x}_t} \left[\sum_j P_t^{i,j} \{ b_t(\cdot) + \alpha_{t+1} F_{t+1,j} \} \right] \quad (12)$$

where $P_t^{i,j}$ are the transition probabilities.

Compared to the traditional SDDP formulation, the cuts are now categorized based on their climate state j :

$$F_{t+1,j} - \boldsymbol{\varphi}_{t+1,j}^l \mathbf{s}_{t+1} \leq \boldsymbol{\gamma}_{t+1,j}^l \mathbf{q}_{t,j} + \boldsymbol{\beta}_{t+1,j}^l \quad (13)$$

where $\boldsymbol{\varphi}_{t+1,j}^l$, $\boldsymbol{\gamma}_{t+1,j}^l$, and $\boldsymbol{\beta}_{t+1,j}^l$ are $1 \times n$ vectors of cuts' parameters when the system is in climate state j . These parameters are analytically derived from the primal and dual information that becomes available as the algorithm progresses backward (see Appendix A for more details on the mathematical procedure to derive the cuts' parameters). Note that the other constraints Equations 3–8 of the standard SDDP model remain the same.

Algorithm 1: Backward Optimization

```

1 while not converged do
2   Initialize the BTF of the last stage
3    $F_T(S_{T+1}, q_T) = 0$ 
4   for  $t = T, T-1, \dots, 1$  do
5     for each sample storage  $s_t = s_t^1, \dots, s_t^L$  do
6       for each regime state  $j = 1, \dots, \Phi$  do
7         Retrieve cuts parameters generated at stage  $t + 1$ 
8         for each inflow scenario conditioned upon regime state
9            $q_{t,j}^k = q_t^1, \dots, q_t^K$  do
10            Solve SDDP OSOP considering initial storage and inflow
11            scenarios
12            Save dual variables
13          end
14        Aggregate dual variables and calculate cut parameters:
15         $\varphi_{t+1,j}^l, \gamma_{t+1,j}^l, \beta_{t+1,j}^l$ 
16      end
17    end
18  end
19  Calculate the upper bound to  $Z$ 
20 end

```

Since the cuts Equation 13 only offer an approximation overestimating the true benefit-to-go and thus Z , a forward simulation of the system with those cuts will automatically provide a lower bound to Z . In each forward simulation phase, the system is simulated using a set of M hydrologic sequences, which can be historical or synthetically-generated using, for example, Equation 10 or Equation 11. At the end of the simulation phase, the average total return over the M hydrologic sequences is therefore a lower bound to Z , which must be compared to the upper bound calculated at the end of the backward optimization phase. When the difference between the upper bound falls within the confidence interval around the lower bound, the algorithm stops (Tilmant et al., 2008).

The pseudo-codes of the backward optimization and the forward simulation phases with hidden climate states are given next:

Algorithm 2: Forward Simulation

```

1 while not converged do
2   Alternately generate  $M$  inflow sequences (synthetic inflows or historical
3   records) from  $\Phi$  climate state  $q_{t,j}^m = \begin{cases} q_{1,j}^1, \dots, q_{T,j}^1 \\ \dots \\ q_{1,j}^M, \dots, q_{T,j}^M \end{cases} \quad j = 1, \dots, \Phi;$ 
4   Initialize state variables at  $t = 1$ 
5   for  $m = 1, \dots, M$  do
6     for  $t = 1, \dots, T$  do
7       Retrieve vector of flows  $q_{t,i}$ 
8       Retrieve cuts parameters  $\varphi_{t+1,j}^l, \gamma_{t+1,j}^l, \beta_{t+1,j}^l$ 
9       Retrieve transition probabilities  $P_{t+1}^{i,j}$ 
10      Solve SDDP OSOP
11      Save simulated variables
12    end
13  end
14  Calculate the lower bound to  $Z$ 
15 end

```

Compared to the traditional SDDP formulation, the backward optimization phase comprises an additional loop on the hidden climate states (Φ) and the cuts are stored according to their climate state. In the forward simulation phase, at stage t , the benefit-to-go F_{t+1} is then the weighted sum of $F_{t+1,j}$ where the weights are the transition probabilities $P_{t+1}^{i,j}$, that is, the probabilities of moving from, say, the climate state i at time t to the different climate states j at time $t + 1$ ($j = 1, 2, \dots, \Phi$). Since this additional loop in the backward optimization phase automatically increases the computation time, it is one factor to be considered when selecting the number of discrete climate states alongside with the length and inherent variability of streamflow records.

The next section describes the case study that was used to illustrate the gain associated with SDDP-derived allocation policies tailored to specific climates/regimes.

3. Case Study

The Senegal River has a flow regime characterized by multiyear dry, normal, and wet periods (Bader et al., 2014). The basin covers an area of about 337 000 km² in West Africa and is shared by four countries: Guinea, Mali, Mauritania, and Senegal (Figure 5). Three tributaries, the Bafing, Faleme, and Bakoye all have their source in the Fouta Djallon, a high plateau located mainly in the Guinean portion of the SRB. Together, these three tributaries contribute to about 90% of the water flowing in the Senegal River. The basin has three distinct stretches: the mountainous upper basin, the valley, and the delta, which is rich in biological diversity and wetlands.

In the 1940s, the first major attempt to control the Senegal River discharges were made in order to cultivate rice in the delta (United Nations Report, 2003). Today, irrigation is still the motor of development in the basin, particularly in the valley and in the delta. Fishing, is also an important activity in the basin, specially in the valley and the delta. More recently, Senegal River flows have been used to spin the turbines of two hydropower plants: Manantali, which is a 200-MW power station, and Félou, a 62-MW run-of-river power plant. Both supply energy to the West African Power Pool.

The climatic regime of the basin is defined by three seasons, a rainy season from June to September, a cold, dry off-season from October to February, and a hot, dry off-season during the rest of the year. The river flow regime depends, for the most part, on the rain that falls in the upper basin in Guinea (about 2000 mm/year). In the valley and the delta, however, the rainfall is rarely more than 500 mm/year. The year-to-year variability of the river discharges during the high flow season is significant and exposes water users to a high hydrological risk (Tilmant et al., 2020). As indicated above, the flow regime is also characterized by periods of low, normal, and high flows, which can persist for several years. The multiyear drought of the 80ies, for example, is well documented and has had devastating consequences on local communities whose livelihood depends on the flow and the banks of the river for fishing and/or subsistence agriculture (Bruckmann & Beltrando, 2016).

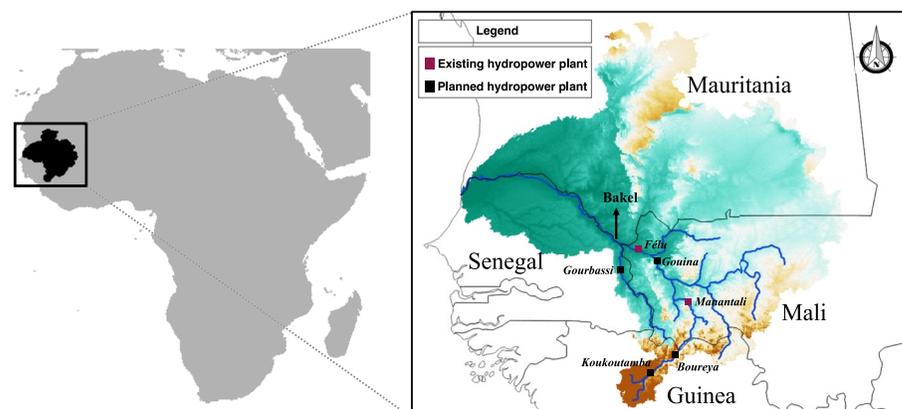


Figure 5. The Senegal river basin.

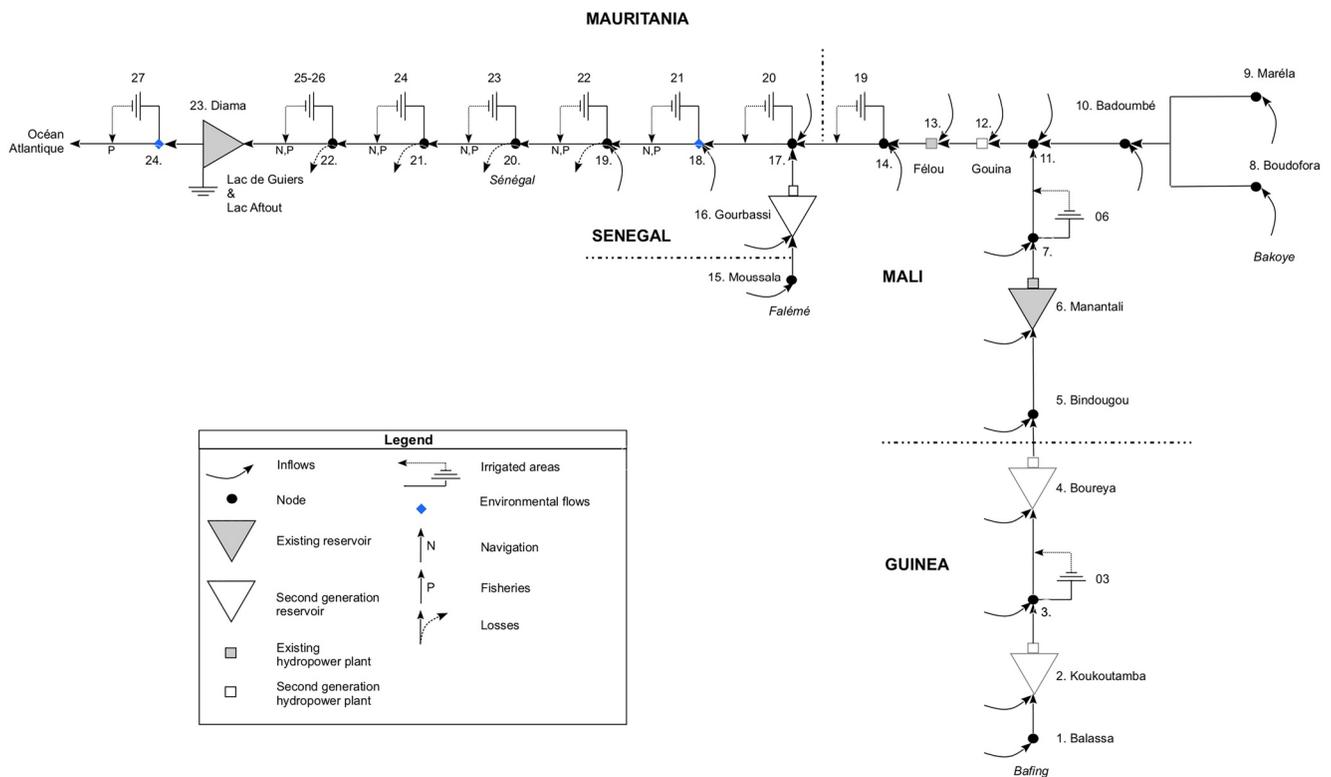


Figure 6. Schematization of the SRB for an intermediate development scenario (Tilmant et al., 2020).

The Senegal is one of the few transboundary river basins managed by a genuine active joint-management organization (OMVS, the Senegal River Basin Authority) with personnel coming from the four riparian countries (Kliot et al., 2001). OMVS’ mandates involve not only the planning but also the operation of water resources in the basin. In this study, the intermediate development of the SRB is considered (Figure 6). It involves 5 reservoirs, 6 power plants, and about 200.000 irrigated hectares. In the next section, optimal allocation policies for that particular future development scenario will be determined and then simulated using historical streamflow data over the 1904–2011 period.

Table 1 lists the main characteristics of the system while agricultural water demands are presented in Table 2. The future configuration of the system has a total storage capacity of 22.3 km³, which corresponds to a 88% increase compared to the current storage capacity. The power plants listed in Table 1, will add 626 MW (+310%) to the installed capacity of the current system. The reservoir inflow-to-storage ratios computed on the mean of the cumulative inflow are 2.09, 3.86, 1.65, and 2.47 for Koukoutamba, Boureya, Manantali, and Gourbassi, respectively, indicating that the system has a seasonal carryover capacity.

Name	Node ID	Storage (km ³)	Installed capacity (MW)
Koukoutamba	2	3.60	294
Boureya	4	4.75	114
Manantali	6	11.30	200
Gouina	12	ROR ^a	140
Felou	13	ROR ^a	60
Gourbassi	16	2.10	18
Diamas	23	0.58	–

^aROR: run of river power plant.

The monthly incremental inflows at the different nodes of the network are based on naturalized flows extracted from Bader et al. (2014) for the 1904–2011 period.

4. Results and Discussion

4.1. Climate Sequences and HMM Classifications

To unfold the hidden climate states attached to the inflow time series in the SRB, we first fit a HMM to the time series of monthly streamflow data measured at Bakel (Senegal River). Bakel marks the transition between the upper (humid) part of the river basin and the lower valley where the contribution of tributaries to the river discharges is much smaller; as mentioned earlier, about 90% of the flows in the SRB goes through Bakel. Prior to fitting a HMM,

Table 2
Crop Water Demands (Intermediate Development Scenario)

Node ID	Withdrawals (hm ³ /y)	Area (ha)
3	369	19926
7	30	1562
14	208	10948
17	17	862
18	140	7147
19	1174	59936
20	553	28219
21	1097	44402
22	1537	78498
24	75	3827
Total	5198	255327

outliers and the seasonal component are removed (Shoelson, 2021). Then, model parameters are fitted using the EM algorithm. The optimum order of HMM (Φ) is not known a priori and must be selected based on objective criteria. Both the Akaike Information Criteria (AIC) (Akaike, 1974) and the Bayesian information Criteria (BIC) (Schwarz, 1978) can be used to determine the optimum order for a HMM. Note that a trade-off must be found between the model complexity (i.e., the number of HMM parameters to be estimated) and the number of climate state that can properly represent historical climate shifts. In this study the added value for both AIC and BIC when considering more than 3 climate states was constant, indicating that there is no added-value to more than three states. The analysis is therefore carried out with $\Phi = 3$, representative of dry, normal, and wet climate regimes as illustrated in Figure 7.

Results show that the hidden climate state with the lowest mean, denoted by “dry,” dominates the period between 1971 and 1996, which corresponds to the extended drought of the late 20th century, which has been analyzed, for example, by Faye et al. (2015) and Sagna et al. (2021). The monthly river discharges are then classified based on the fitted HMM. Figure 8 compares the statistics of the climate states for each month. The red, black, and blue boxes are associated with dry, normal, and wet climates/regimes respectively. We can see that the climate states’ statistics are distinct, which validate the choice of the underlying structure of the HMM (Figure 8).

Table 3 presents the transition probability matrix associated with the HMM. As we can see, the transition probabilities are close to zero or one, suggesting again that the states are clearly distinct. The values close to one on the diagonal indicate that when the climate is in a particular state, it will likely remain in that state at the next time period (month). In other words, shifting from one climate state to another is far less likely. For the dry climate state, for example, there is more than 90% probability of remaining in that state and less than 1% probability of transitioning to a wet state.

4.2. Analysis of Simulation Results

To assess the gains associated with climate-derived allocation policies, three SDDP formulations, each having a specific hydrologic model, are implemented and their performances compared. The first formulation attempts at capturing the temporal persistence of the hydrologic processes through a MPAR(p) model (de Matos & Finardi, 2012; Jasson et al., 2017) whose order p varies in space (site) and time (month). This version is denoted SDDP-MPAR. The second formulation is the proposed extension in which the state-space vector now includes as

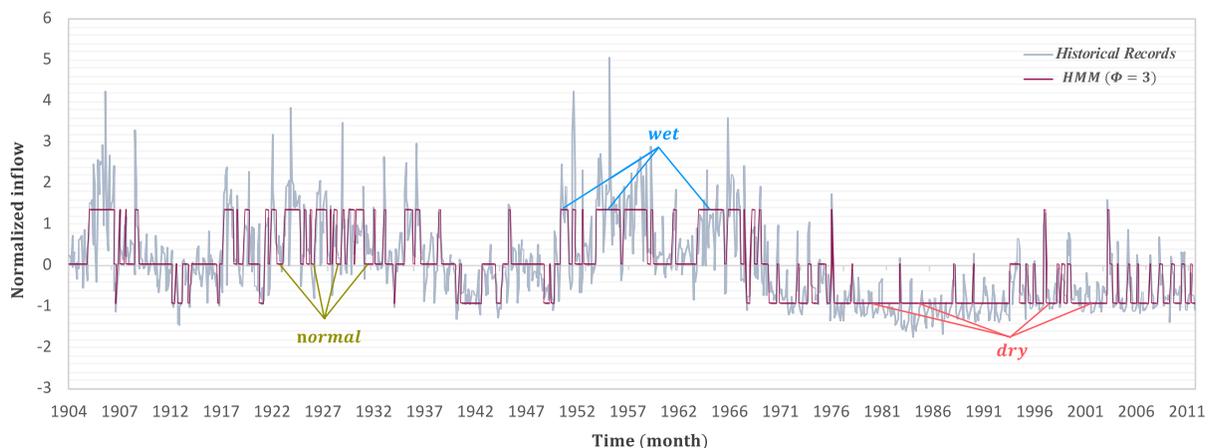


Figure 7. Persistent structure of the underlying regime process at Bakel in SRB considering three possible climate states (i.e., dry, normal, and wet).

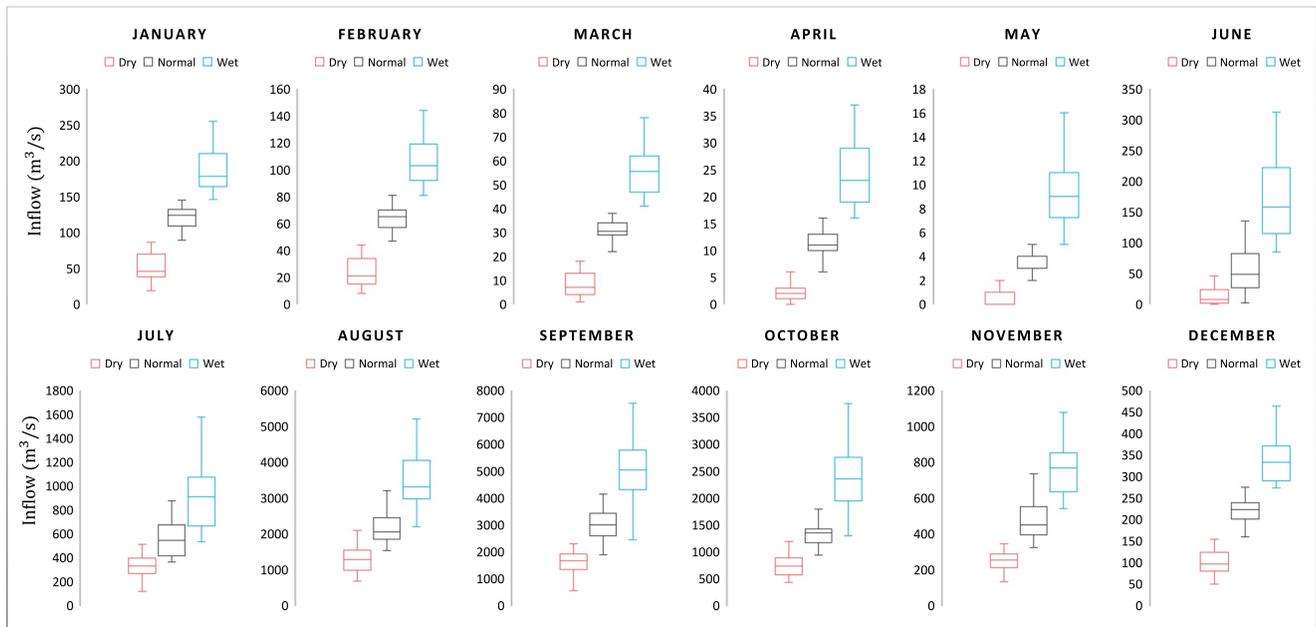


Figure 8. Classified monthly inflows at Bakel derived from Hidden Markov Model when $\Phi = 3$.

hydrologic variables the previous flow q_{t-1} and the climate state Φ_t with the latter capturing the long term persistence through three different states (dry, normal, and wet). This second version is denoted SDDP-HMM. The third formulation is actually a deterministic dual DP (DDDP) in which one assumes perfect foresight of future inflows, therefore overestimating the performance of the system. The comparison between SDDP-HMM and SDDP-MPAR will reveal the gains associated with allocation policies that explicitly capture the long-term persistence and that are therefore tailored to shifting flow regimes. The comparison between SDDP-HMM and DDDP will indicate the extent of the losses associated with an imperfect knowledge of future hydrologic conditions. In other words, the difference in terms of performances between SDDP-HMM and DDDP is the opportunity cost associated with imperfect flow forecasts. Table 4 lists the main characteristics of the three dual DP formulations.

Both SDDP formulations (with and without climate states) are implemented with the following hydrologic configuration: the number of backward openings is set to forty (i.e., $K = 40$), and forward simulations are carried out on 40 historical sequences of 5 years (i.e., $M = 40$). This setup is in line with previous studies (e.g., Pina et al. (2017); Tilmant et al. (2009) and was found to be a good compromise between computation time and the accurate representation of the hydrologic uncertainty. Both SDDP built-in hydro models are calibrated using 107 years of monthly historical record extracted from Bader et al. (2014) for the 1904–2011 period.

For both SDDP formulations, the planning period is 60 months. Since the boundary conditions imposed on the state variables influence the allocation decisions during the first (by initial storages) and last stages (by zero terminal value function) (Rougé & Tilmant, 2016), a buffer of 2 years is considered in this study. This assumption is motivated by the physical characteristics of the system and the fairly detailed data set, which prevent the presence of multiple near-optimal solutions (Rougé & Tilmant, 2016). It means that once the SDDP algorithm has converged, the cuts from January to December of the third year are kept and the others are discarded. The

results are analyzed after reoptimizing the policies along the 107 years of historical river discharges. Reoptimizing the policies is similar to the procedure proposed by Tejada-Guibert et al. (1993) with SDP except that the benefit-to-go function F_{t+1} is now piecewise linear (Tilmant et al., 2020). The reoptimization procedure was implemented here because we wanted to see how the proposed extension would handle the multiyear drought of the 80ies, which had catastrophic consequences on riverine communities and led to mass migration (Dia, 2007).

Table 3
Transition Probability Matrix $P_i^{j,j}$ From i th Climate State to the j th Derived From Fitted Hidden Markov Model at Bakel in SRB When $\Phi = 3$

	Dry	Normal	Wet
Dry	0.921	0.067	0.012
Normal	0.078	0.858	0.064
Wet	0.010	0.110	0.880

Table 4
Characteristics of Two Different Built-In Hydro Model of Stochastic Dual Dynamic Programming

Formulation	p	Φ	Time series of inflow	Time series of climate states
SDDP-MPAR(p)	Up to 6	–	✓	×
SDDP-HMM(Φ)	1	3	✓	✓
DDDP	–	–	✓	×

Table 5 lists the average and standard deviation of annual hydropower generation for the three dual DP formulations. As we can see, with SDDP-MPAR, the mean annual hydropower generation is roughly 5.7% less than the deterministic case. When the hidden climate states are included in the algorithm, this difference is reduced to 4%. In other words, incorporating the climate states can fill 30% ($1.7/5.7 \times 100$) of the gap in terms of energy. The standard deviation reflects the variability of hydropower generation in the basin. For example, with optimum operating policies provided by SDDP-HMM, the variability of power production during dry, normal, and wet years is 2.3% less than the traditional SDDP-MPAR formulation. The impact of these decisions on the amount of water spilled is also presented in Table 5 where we can see with SDDP-HMM the average annual spillage is reduced by 3.7% compared to the traditional SDDP formulation.

The stationary probabilities associated with dry, normal, and wet years are 28%, 53%, and 19% respectively. These stationary probabilities can be used to analyze the differences between model formulations within the dry, normal, and wet periods. For example, Figure 9 traces out the cumulative distribution of total annual hydropower generation for the three dual DP formulations. As we can see, the DDDP clearly provides an upper bound on the performance in terms of energy. During dry years, with a non-exceedance probability between zero and 28%, considering the long-term persistence is beneficial because the results with SDDP-HMM are almost as good as DDDP. During normal years, however, this gain is less pronounced and the difference between the cumulative distribution associated with SDDP-MPAR and SDDP-HMM becomes marginal; both SDDP formulations (with and without climate states) generate similar policies meaning that from a decision-making point-of-view, the climate states become informationless during “normal” years. The added-value comes back again during wet years, which are associated with non-exceedance probabilities beyond 81%. These results clearly indicate that the incorporation of hidden climate states yields better policies during extreme hydrological conditions (dry and wet years) compared to the traditional SDDP formulation.

Taking the DDDP model as a benchmark, Table 6 lists, for both SDDP formulations, the mean annual reduction in energy output and the mean annual increase in spillage losses that one can expect during dry, normal, and wet years. During dry years, compared to the deterministic case, the energy output is reduced by 10.7% when allocation policies only consider the short-term persistence. This difference reaches 6.51% when allocation policies are tailored to climate states. During normal years, as indicated above, SDDP-MPAR and SDDP-HMM propose similar allocation policies since the differences with respect to the deterministic case are roughly the same: 4.41% versus 4.11%. Allocation policies derived from SDDP-MPAR and SDDP-HMM are again distinct during wet years where the difference, in terms of energy output, with respect to the deterministic case are now: 4.22% versus 1.82%.

The impact of hidden climate states on storage trajectories is analyzed in Figure 10 where we can see the time series of aggregated storages over the simulation period for the three dual DP formulations. At the onset of the multiyear drought centered around the 80ies, SDDP-MPAR is almost depleting the reservoirs too quickly. This can be explained by the fact that the autoregressive model is overestimating the inflows during dry years. However, when the hidden climate states are included in the algorithm, allocation policies better hedge against such adverse conditions, and storage levels are kept higher and closer to that of DDDP.

Table 5
Average Annual Results in Terms of Hydropower Generation and Spillage: Mean and Standard Deviation

Energy	SDDP-MPAR	SDDP-HMM	DDDP
Mean (GWh)	3493	3554	3704
Std (GWh)	920	898	881
Spillage	SDDP-MPAR	SDDP-HMM	DDDP
Mean (hm ³)	10923	10512	8966
Std (hm ³)	12582	12077	9976

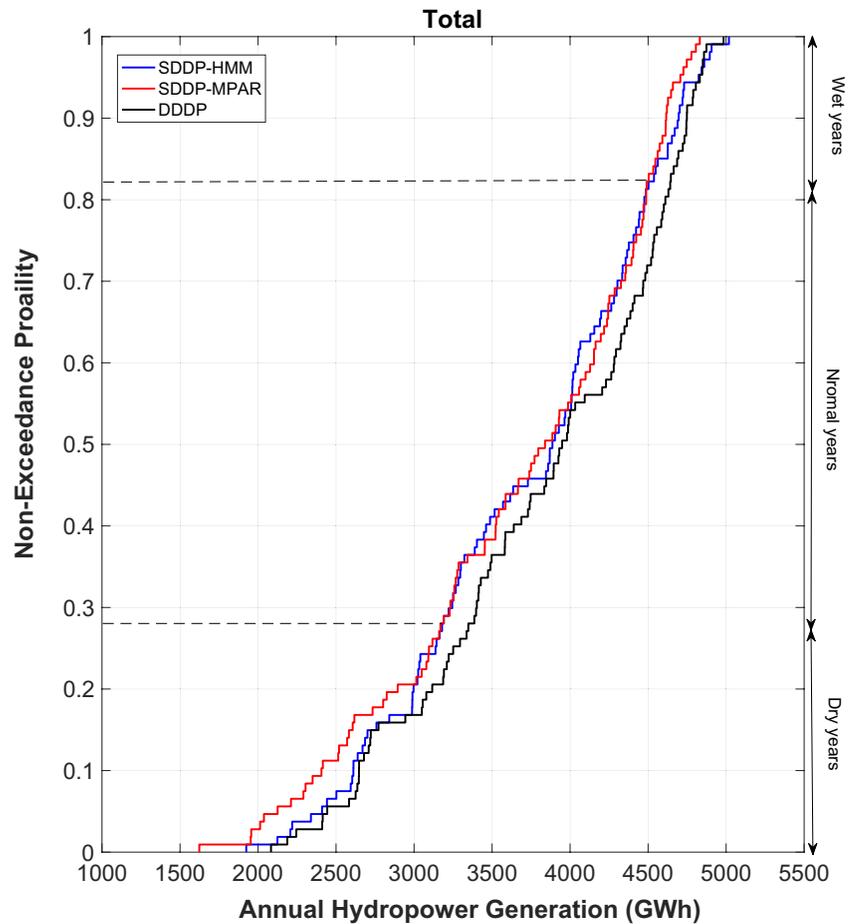


Figure 9. Cumulative distribution of total annual hydropower generation for three Stochastic Dual Dynamic Programming formulation in the SRB.

4.3. Evaluating the Performance of the SDDP Built-In Hydro Models

As mentioned earlier in Section 2.4, the SDDP algorithm requires that, at stage t , K vectors of inflows \mathbf{q}_t be generated by the built-in hydrological model so that K OSOP can be solved to capture the stochasticity of the hydrologic conditions and to derive the expected values of the parameters of the cuts approximating the benefit-to-go function. In other words, for each node n , an ensemble of K streamflow forecasts is available at the beginning of each stage. Generally speaking, the better the forecasts, the better the approximation of the benefit-to-go function F_t since the cuts' parameters directly depend on the primal and dual information associated with the K OSOP solved at stage t , one for each forecast. See Appendix A for more details on the procedure to derive the cuts' parameters.

In hydrology, the overall quality of ensemble streamflow forecasts is assessed with the continuous ranked probability score (CRPS) (Matheson & Winkler, 1976). This common verification metric measures the quadratic distance between the cumulative distribution of the forecasts and the cumulative distribution of the observations, with a value

Hydro condition	Annual loss in terms of energy (%)		Annual increase in spillage losses (%)	
	SDDP-MPAR(p)	SDDP-HMM(3)	SDDP-MPAR(p)	SDDP-HMM
Dry years	10.77	6.51	35.73	8.35
Normal years	4.41	4.11	21.66	21.32
Wet years	4.22	1.82	22.34	13.89

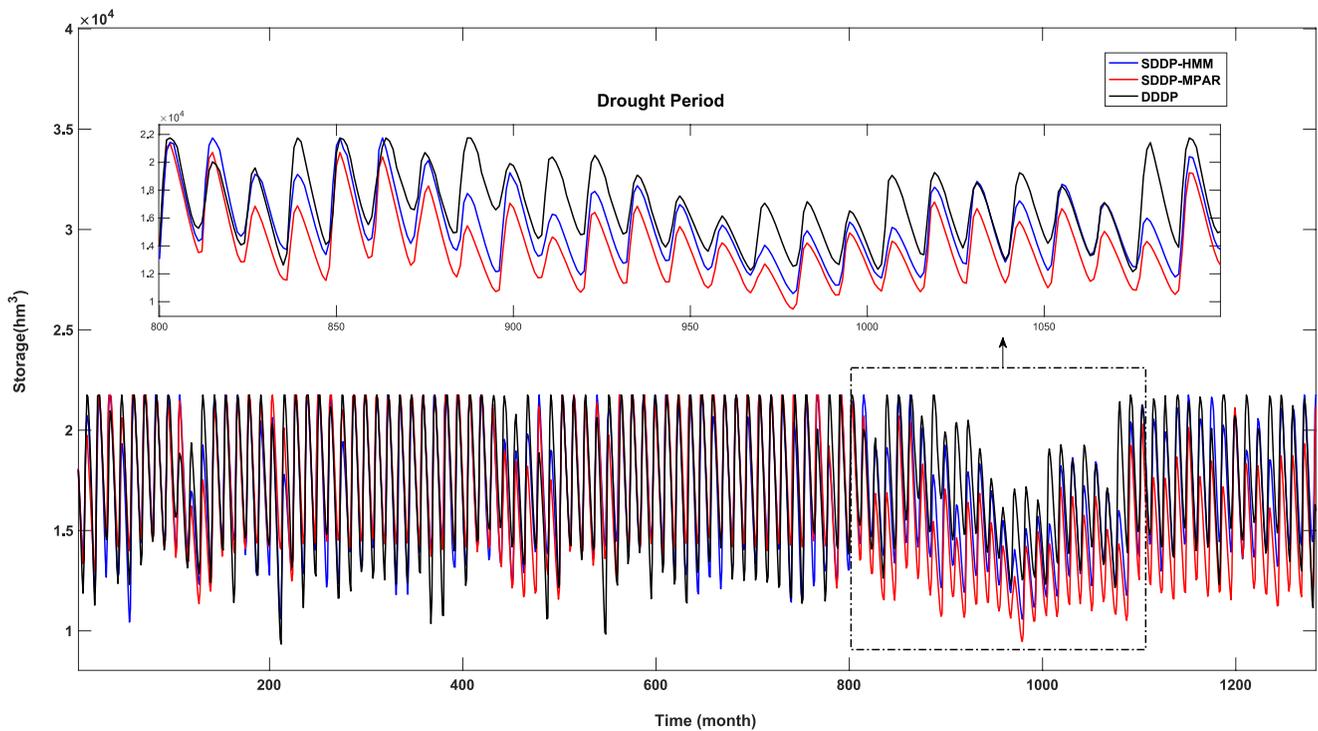


Figure 10. Aggregated reservoir trajectories for the three Stochastic Dual Dynamic Programming formulations (1904–2011).

of 0 indicating a perfect forecast. As the CRPS assesses the forecast for a single time step, the MCRPS is defined as the average CRPS over the entire planning period T (Thiboult et al., 2016). To allow for the comparison among different nodes, we normalized the MCRPS with the standard deviation of the observations (SD_o) (Cassagnole et al., 2021). The normalized mean continuous ranked probability score is estimated using the following equation:

$$NMCRPS(P_{cum}, o_t) = \frac{\left(\frac{1}{T} \sum_{t=1}^T \int_{-\infty}^{\infty} [P_{cum_t}(q_t) - I_{o_t}(q_t)]^2 dq\right)}{SD_o} \quad (14)$$

where o_t is the observed inflow, P_{cum_t} is the cumulative probability distribution of the forecasts at time t , q_t is the predicted inflow, and I_{o_t} is the Heaviside function defined as:

$$I_{o_t}(q_t) = \begin{cases} 0 & q_t < o_t \\ 1 & q_t \geq o_t \end{cases} \quad (15)$$

Table 7
NMCRPS Scores for the Two Hydrological Models: Multisite Periodic Autoregressive Model (MPAR) and MPAR-HMM

Name	Node ID	MPAR	MPAR-HMM
Koukoutamba	2	0.103	0.094
Boureya	4	0.103	0.097
Manantali	6	0.104	0.097
Gouina	12	0.108	0.107
Felou	13	0.108	0.109
Gourbassi	16	0.126	0.109

Table 7 essentially compares the overall quality of the forecasts of the incremental flows generated by the two built-in hydrological models: MPAR or MPAR-HMM. The values highlighted in bold show the best model. We can see that HMM-based forecasts are more accurate than those derived from a model that only captures the short-term persistence (MPAR). The difference is larger for the nodes draining portions of the Fouta Djallon where precipitations are abundant but where the contrast between dry and wet years is also more pronounced. Further downstream, as the river runs north-west and the climate gets drier, the relative contribution of tributaries draining the northern part of the basin becomes not only less important but also less variable year-to-year. Consequently, forecasting the incremental flows at the two run-of-river power plants, Félou and Gouina, is achieved with a similar accuracy regardless of the model.

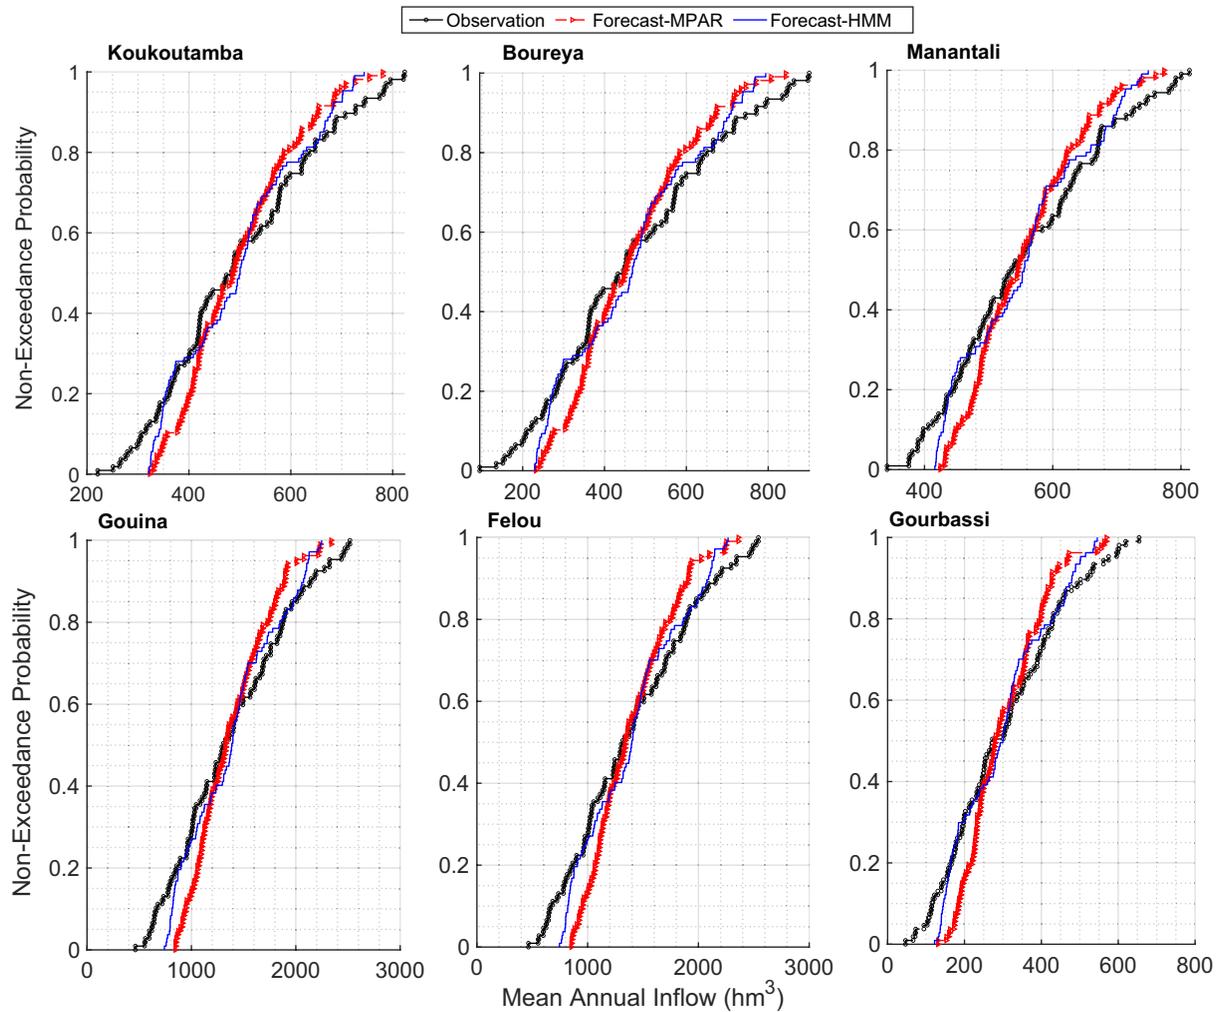


Figure 11. Cumulative distribution of Annual inflows for each hydropower plant in the SRB.

As mentioned above when analyzing both the time series of aggregated storages and the results for the energy sector, the difference between both SDDP formulations is more obvious during dry and wet years. Such a difference comes from the fact that the corresponding built-in hydrological models perform differently as soon as the hydrologic conditions are departing from normal. Figure 11 shows the cumulative distribution of the accumulated annual inflows at each power plant compared to the historical observations for both hydrological models: MPAR and MPAR-HMM. As we can see, none of the models is able to fully capture the variability of the observations. In other words, the spread of the annual streamflows is underestimated. However, as soon as the hidden climate states are incorporated, the situation improves and the cumulative distribution of annual streamflows derived from the MPAR-HMM model moves closer to the CDF of the observations, especially during dry and wet years.

Regarding the agricultural sector, the performances associated with the three dual DP formulations are identical: all irrigation schemes are supplied with a reliability above 95%. Even though the analyzed system corresponds to a future development scenario with almost three times as many irrigated hectares as today, the agricultural water demand is still lower than the available resources, especially when the five reservoirs are operational.

5. Conclusions

As low frequency signals are affecting the hydrology worldwide, several river basins are experiencing cycles of prolonged dry and wet periods. Properly capturing such a regime-like behavior is essential to enhance the performance of our water resources systems. This is especially the case for the multireservoir systems with large storage

capacity because they can move water over extended periods of time. However, determining regime-tailored allocation policies for large-scale water resources systems is challenging as it often requires trading-off hydrological complexity for system complexity.

In this study, we address this trade-off by combining Hidden Markov Modeling with SDDP. The former can be used to capture the long-term persistence that characterizes the flow regime of many rivers, while the latter can be used to determine optimal allocation policies in large multireservoir systems.

The proposed extension of SDDP, called SDDP-HMM, is tested and compared to two other dual DP formulations using a future development scenario of the Senegal River basin as an example. The analysis of simulation results shows: (a) Considering the hidden climate states is particularly relevant during hydrological extremes; reservoir operating policies provided by SDDP-HMM better hedge against adverse hydrologic conditions; (b) The hydrologic forecasts computed by a model that captures both the short-term and the long-term persistences are more accurate than those obtained by the traditional MPAR model found in many SDDP formulations.

Our results are consistent with previous studies such as Turner and Galelli (2016) and show that the incorporation of relevant hydrologic information can improve the performance of a water resources system. The CPU time required for SDDP-MPAR, and SDDP-HMM, was about 8 and 25 min, respectively, on MacBook Pro with the eight-core Apple M1 chip with 8 GB of RAM. Despite the fact that the new SDDP formulation is three times slower, the computational time remains fairly low, making it possible to embed SDDP-HMM in decision support systems.

The proposed approach to address the trade-off between system complexity and hydrological complexity relies on an extension of SDDP, an optimization algorithm which requires that the one-stage optimization problems be linear programs. This condition has several implications. The most important ones are: (a) that non-linear relationships must be made piecewise linear such as Equation 8 for the hydropower production function, (b) the built-in hydrological model Equation 10 must be linear. As shown by Rougé and Tilmant (2016), a side-effect associated with these approximations is that the algorithm is prone to the presence of multiple near-optimal solutions when inputs data are limited and/or when the system has significant carryover storage capacity. When these two conditions are met, we recommend implementing the year-periodic reoptimization procedure proposed by Rougé and Tilmant (2016) instead of directly using the benefit-to-go functions extracted after the buffer period. It must also be stressed that the proposed optimization model assumes that the river basin is managed by a single entity, which is the case in the Senegal. In transboundary river basins where countries do not cooperate, the model could be implemented sequentially, starting with the upstream country and imposing the resulting outflow to the downstream country and so on (Jeuland & Whittington, 2014). In this study, we also assume that the climate state is the same for the whole basin. Extending the model to river basins with asynchronous climate states is straightforward: sub-basin specific climate states must be included in the state vector. Note that this will automatically increase the computation time as all combinations of climate states must be investigated. It must also be stressed that the proposed optimization model assumes that the operator is able to assess the climate state at the beginning of each month. This 1-month perfect foresight is considered here as a reasonable assumption since it can be adjusted (together with the corresponding allocation decisions) as time goes by and the operators know more about the actual climate state during that month. Moreover, even though the climate state is hidden, in some cases, it can be partially revealed through the use of other hydroclimatic information such as sea surface temperature (SST) or soil moisture (Gelati et al., 2010). In the Senegal River basin, for example, there is a relationship between seasonal flows and SST in the Gulf of Guinea (Gu & Adler, 2004).

Appendix A: Estimating Cuts' Parameters in SDDP-HMM

A key step of the Stochastic Dual Dynamic Programming (SDDP) algorithm is the derivation of the cuts' parameters. In the SDDP version with hidden climate states, the l th cut approximating the benefit-to-go $F_{i,j}$ when the climate is in state j has the following form:

$$F_{i,j} - \phi_{i,j}^l s_t \leq \gamma_{i,j}^l q_{t-1,i} + \beta_{i,j}^l \quad (\text{A1})$$

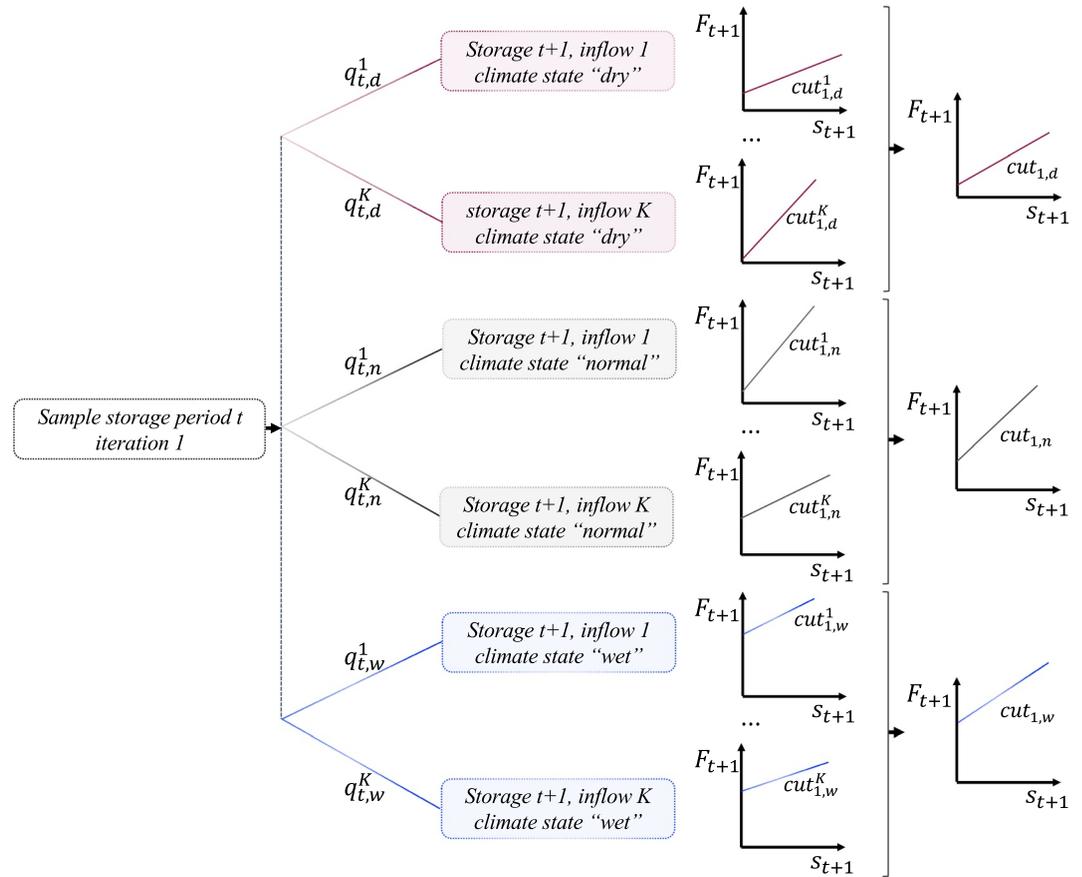


Figure A1. Illustration of the backwards openings and the aggregation of the cuts when the climate is switching between 3 different states (e.g., dry (d), normal (n), and wet (w)).

where $\varphi_{i,j}^l$, $\gamma_{i,j}^l$, and $\beta_{i,j}^l$ are cuts' parameters when system is in climate state j . These parameters can be analytically derived from the primal and dual information that becomes available as the algorithm progresses backward. The parameters of the cuts can be derived from the Karush-Kuhn-Tucker (KKT) conditions (Kemp & Kimura, 1978):

$$\nabla F(x^*) - \sum_i \lambda_i^* g_i(x^*) = 0 \quad (A2)$$

where λ is the dual information of the optimization problem, x denotes a primal variable, and g_i represents the i th linear constraint. Rewriting KKT for our reservoir optimization problem yields:

$$\begin{aligned} \frac{\partial F_{t,j}^k}{\partial x_j} &= \lambda_{w,t,j} \frac{\partial}{\partial x_j} (s_{t+1} + r_t + i_t + l_t - s_t - q_{t,j}) \\ &+ \sum_{l=1}^L \lambda_{c,t,j}^{l,k} \frac{\partial}{\partial x_j} (F_{t+1,j} - \varphi_{t+1,j}^l s_{t+1} - \gamma_{t+1,j}^l q_{t,j} - \beta_{t+1,j}^l) \\ &+ \sum_{l=1}^L \sum_{h=1}^H \lambda_{hp,t,j}^{l,k,h} \frac{\partial}{\partial x_j} (\hat{P}_{t,j} - \psi^h s_{t+1}/2 - \omega^h r_t - \delta^h - \psi^h s_t/2) \end{aligned} \quad (A3)$$

where λ_w is dual variable associated with the mass balance Equation 3, λ_c is dual variable associated with the cuts Equation 9, and λ_{hp} is dual variable associated with convex hull approximation of the hydropower production function Equation 8. Cut parameters can be calculated by taking the partial derivative with respect to the corresponding state variable.

The vector of slopes $\varphi_{t,j}^{l,k}$ with respect to the storage variable s_t is given by:

$$\phi_{t,j}^{l,k} = \frac{\partial F_{t,j}^{k,h,l}}{\partial s_t} = \lambda_{w,t,j}^{l,k} + \sum_{h=1}^H \lambda_{hp,t,j}^{l,k,h} \psi^h / 2 \quad (\text{A4})$$

Since the one-stage SDDP optimization problem is evaluated for K inflows scenarios (branches, backward openings) to capture the hydrologic uncertainty, the individually-calculated cut parameters must be aggregated so that the *expected* benefit-to-go is approximated (Figure A1). For the j th climate state, the vector of expected slopes $\phi_{t,j}^l$ is given by

$$\phi_{t,j}^l = \frac{1}{K} \sum_{k=1}^K \phi_{t,j}^{l,k} \quad (\text{A5})$$

where K is the number of backward openings, that is, inflow branches generated by the built-in hydrologic model as the algorithm progresses backward.

Similarly, the vector of slopes $\gamma_{t,j}^{l,k}$ can be obtained from the dual and primal information associated with mass balance equalities and cuts inequalities:

$$\begin{aligned} \gamma_{t,j}^{l,k} &= \frac{\partial F_{t,j}^{k,h}}{\partial q_{t-1,i}} \\ &= \frac{\partial F_{t,j}^{k,h}}{\partial q_{t,j}} \frac{\partial q_{t,j}}{\partial q_{t-1,i}} \\ &= \left(\lambda_{w,t,j}^{l,k} + \sum_{h=1}^L \lambda_{c,t,j}^{l,k,h} \psi^h / 2 \right) \times \left(\frac{\sigma_{t,j}}{\sigma_{t-1,i}} \phi_t \right) \end{aligned} \quad (\text{A6})$$

Again, the aggregation is carried out by taking the expectation over the K artificially generated flows:

$$\gamma_{t,j}^l = \frac{1}{K} \sum_{k=1}^K \gamma_{t,j}^{l,k} \quad (\text{A7})$$

Finally, the constant term of the cut can be calculated with:

$$\beta_{t,j}^l = \frac{1}{K} \sum_{k=1}^K F_{t,j}^k - \phi_{t,j}^l s_t^o + \gamma_{t,j}^l q_{t-1,i}^o \quad (\text{A8})$$

where s_t^o and $q_{t-1,i}^o$ are the sampled storages and inflows respectively.

Acknowledgments

This project was supported by the Natural Sciences and Engineering Research Council of Canada Discovery Grant of the second author and the GoNEXUS project, which is funded by the European Union Horizon Programme call H2020-LC-CLA-2018-2019-2020 - Grant Agreement Number 101003722. The participation of Canadian researchers to GoNEXUS is made possible by the New Frontiers Research Fund program (Canada) (grant NFRFG-2020-00430) and by the Fonds Nouvelles Frontières en Recherche (Quebec) (grant 2022-FNFR-310131). Figure 6 is reprinted from Journal of Hydrology, 587, Tilmant et al., Probabilistic trade-off assessment between competing and vulnerable water users - The case of the Senegal River basin, Copyright (2020), with permission from Elsevier.

Data Availability Statement

All the data used in this study are from the Senegal River Basin Authority (OMVS) and were collected during a project funded by UN-FAO (TCP/INT/3602). Because the model contains sensitive information on existing and planned hydropower plants that is protected by a nondisclosure agreement with UN-FAO, it cannot be made public. These data can be requested by contacting OMVS (www.omvs.org, T: +221 338598182, E: omvssphc@omvs.org).

References

- Akaike, H. (1974). A new look at the statistical model identification. *IEEE Transactions on Automatic Control*, 19(6), 716–723. <https://doi.org/10.1109/TAC.1974.1100705>
- Akintug, B., & Rasmussen, P. (2005). A Markov switching model for annual hydrologic time series. *Water Resources Research*, 41(9), W09424. <https://doi.org/10.1029/2004wr003605>
- Bader, J.-C., Cauchy, S., Duffar, L., & Saura, P. (2014). *Monographie hydrologique du fleuve Sénégal. De l'origine des mesures jusqu'en 2011*, In IRD Editio (Ed.), IRD. Retrieved from <https://www.documentation.ird.fr/hor/fdi:010065190>
- Bayazit, M. (2015). Nonstationarity of hydrological records and recent trends in trend analysis: A state-of-the-art review. *Environmental Processes*, 2(3), 527–542. <https://doi.org/10.1007/s40710-015-0081-7>

- Bilmes, J. A. (1998). *A gentle tutorial of the EM algorithm and its application to parameter estimation for Gaussian mixture and hidden Markov models*. In I. C. S. Institute (Ed.), (Vol. 4). (No. 510). <https://doi.org/10.1080/0042098032000136147>
- Bracken, C., Rajagopalan, B., & Zagona, E. (2014). A hidden Markov model combined with climate indices for multidecadal streamflow simulation. *Water Resources Research*, *50*(10), 7836–7846. <https://doi.org/10.1002/2014WR015567>
- Bruckmann, L., & Beltrando, G. (2016). Transboundary water resources management and livelihoods: Interactions in the Senegal river. In *EGU general assembly conference abstracts* (pp. EPSC2016–7355).
- Cassagnole, M., Ramos, M.-H., Zalachori, I., Thirel, G., Garçon, R., Gailhard, J., & Ouillon, T. (2021). Impact of the quality of hydrological forecasts on the management and revenue of hydroelectric reservoirs—a conceptual approach. *Hydrology and Earth System Sciences*, *25*(2), 1033–1052. <https://doi.org/10.5194/hess-25-1033-2021>
- Castelletti, A., Pianosi, F., & Restelli, M. (2013). A multiobjective reinforcement learning approach to water resources systems operation: Pareto Frontier approximation in a single run. *Water Resources Research*, *49*(6), 3476–3486. <https://doi.org/10.1002/wrcr.20295>
- Coulibaly, P., & Burn, D. H. (2004). Wavelet analysis of variability in annual Canadian streamflows. *Water Resources Research*, *40*(3), W03105. <https://doi.org/10.1029/2003wr002667>
- de Matos, V. L., & Finardi, E. C. (2012). A computational study of a stochastic optimization model for long term hydrothermal scheduling. *International Journal of Electrical Power & Energy Systems*, *43*(1), 1443–1452. <https://doi.org/10.1016/j.ijepes.2012.06.021>
- Dia, H. (2007). Migrant investments in the Senegal River valley: Trust and conflicts of interest. *European Journal of International Migration*, *23*(3), 29–49.
- Faye, C., Diop, E. H. S., & Mbaye, I. (2015). Impacts des changements de climat et des aménagements sur les ressources en eau du fleuve Sénégal: Caractérisation et évolution des régimes hydrologiques de sous-bassins versants naturels et aménagés. *Belgeo. Revue belge de géographie*, *4*. <https://doi.org/10.4000/belgeo.17626>
- Fortin, V., Perreault, L., & Salas, J. (2004). Retrospective analysis and forecasting of streamflows using a shifting level model. *Journal of Hydrology*, *296*(1–4), 135–163. <https://doi.org/10.1016/j.jhydrol.2004.03.016>
- Gelati, E., Madsen, H., & Rosbjerg, D. (2010). Markov-switching model for nonstationary runoff conditioned on El Niño information. *Water Resources Research*, *46*(2), W02517. <https://doi.org/10.1029/2009wr007736>
- Goor, Q., Kelman, R., & Tilmant, A. (2010). Optimal multipurpose-multireservoir operation model with variable productivity of hydropower plants. *Journal of Water Resources Planning and Management*, *137*(3), 258–267. [https://doi.org/10.1061/\(ASCE\)WR.1943-5452.0000117](https://doi.org/10.1061/(ASCE)WR.1943-5452.0000117)
- Gu, G., & Adler, R. F. (2004). Seasonal evolution and variability associated with the West African monsoon system. *Journal of Climate*, *17*(17), 3364–3377. [https://doi.org/10.1175/1520-0442\(2004\)017<3364:SEAWAW>2.0.CO;2](https://doi.org/10.1175/1520-0442(2004)017<3364:SEAWAW>2.0.CO;2)
- Hurst, H. E. (1951). Long-term storage capacity of reservoirs. *Transactions of the American Society of Civil Engineers*, *116*, 770–799. <https://doi.org/10.1061/taceat.0006518>
- Jasson, P., Amaury, T., & Pascal, C. (2017). Optimizing multireservoir system operating policies using exogenous hydrologic variables. *Water Resources Research*, *53*(11), 9845–9859. <https://doi.org/10.1002/2017WR021701>
- Jeuland, M., & Whittington, D. (2014). Water resources planning under climate change: Assessing the robustness of real options for the Blue Nile. *Water Resources Research*, *50*(3), 2086–2107. <https://doi.org/10.1002/2013WR013705>
- Jha, S., Das, J., & Goyal, M. K. (2021). Low frequency global-scale modes and its influence on rainfall extremes over India: Nonstationary and uncertainty analysis. *International Journal of Climatology*, *41*(3), 1873–1888. <https://doi.org/10.1002/joc.6935>
- Karamouz, M., Szidarovszky, F., & Zahraie, B. (2003). *Water resources systems analysis*. CRC Press.
- Kelman, J., Stedinger, J. R., Cooper, L. A., Hsu, E., & Yuan, S.-Q. (1990). Sampling stochastic dynamic programming applied to reservoir operation. *Water Resources Research*, *26*(3), 447–454. <https://doi.org/10.1029/WR026i03p00447>
- Kemp, M. C., & Kimura, Y. (1978). *Introduction to mathematical economics*. [Book; Book/illustrated]. Springer-Verlag. Retrieved from <http://digitool.hbz-nrw.de:1801/webclient/DeliveryManager?pid=2363886&custom%5Fatt%5F2=simple%5Fviewer>
- Kim, Y.-O., & Palmer, R. N. (1997). Value of seasonal flow forecasts in Bayesian stochastic programming. *Journal of Water Resources Planning and Management*, *123*(6), 327–335.
- Kliot, N., Shmueli, D., & Shamir, U. (2001). Institutions for management of transboundary water resources: Their nature, characteristics and shortcomings. *Water Policy*, *3*(3), 229–255. [https://doi.org/10.1016/S1366-7017\(01\)00008-3](https://doi.org/10.1016/S1366-7017(01)00008-3)
- Koutsoyiannis, D. (2002). The Hurst phenomenon and fractional Gaussian noise made easy. *Hydrological Sciences Journal*, *47*(4), 573–595. <https://doi.org/10.1080/02626660209492961>
- Labadie, J. W. (2004). Optimal operation of multireservoir systems: State-of-the-art review. *Journal of Water Resources Planning and Management*, *130*(2), 93–111.
- Libisch-Lehner, C., Nguyen, H., Taormina, R., Nachtnebel, H., & Galelli, S. (2019). On the value of ENSO state for urban water supply system operators: Opportunities, trade-offs, and challenges. *Water Resources Research*, *55*(4), 2856–2875. <https://doi.org/10.1029/2018WR023622>
- Liu, H., Sun, Y., Yin, X., Zhao, Y., Cai, Y., & Yang, W. (2020). A reservoir operation method that accounts for different inflow forecast uncertainties in different hydrological periods. *Journal of Cleaner Production*, *256*, 120471. <https://doi.org/10.1016/j.jclepro.2020.120471>
- Lohmann, T., Hering, A. S., & Rebennack, S. (2016). Spatio-temporal hydro forecasting of multireservoir inflows for hydro-thermal scheduling. *European Journal of Operational Research*, *255*(1), 243–258. <https://doi.org/10.1016/j.ejor.2016.05.011>
- Loucks, D. P., & Van Beek, E. (2017a). An introduction to probability, statistics, and uncertainty. In *Water resource systems planning and management* (pp. 213–300). Springer.
- Loucks, D. P., & Van Beek, E. (2017b). *Water resource systems planning and management: An introduction to methods, models, and applications*. Springer.
- Matheson, J. E., & Winkler, R. L. (1976). Scoring rules for continuous probability distributions. *Management Science*, *22*(10), 1087–1096. <https://doi.org/10.1287/mnsc.22.10.1087>
- Mbeutcha, Y., Gendreau, M., & Emiel, G. (2021). Benefit of PARMA modeling for long-term hydroelectric scheduling using stochastic dual dynamic programming. *Journal of Water Resources Planning and Management*, *147*(3), 05021002. [https://doi.org/10.1061/\(asce\)wr.1943-5452.0001333](https://doi.org/10.1061/(asce)wr.1943-5452.0001333)
- Mujumdar, P., & Nirmala, B. (2007). A Bayesian stochastic optimization model for a multi-reservoir hydropower system. *Water Resources Management*, *21*(9), 1465–1485. <https://doi.org/10.1007/s11269-006-9094-3>
- Olsen, J. R., Stedinger, J. R., Matalas, N. C., & Stakhiv, E. Z. (1999). Climate variability and flood frequency estimation for the upper Mississippi and lower Missouri rivers I. *JAWRA Journal of the American Water Resources Association*, *35*(6), 1509–1523.
- Pereira, M., & Pinto, L. (1983). Application of decomposition techniques to the mid-and short-term scheduling of hydrothermal systems. *IEEE Transactions on Power Apparatus and Systems*, *PAS-102*(11), 3611–3618. <https://doi.org/10.1109/tpas.1983.317709>
- Pereira, M. V., & Pinto, L. M. (1991). Multi-stage stochastic optimization applied to energy planning. *Mathematical Programming*, *52*(1–3), 359–375. <https://doi.org/10.1111/j.1752-1688.1999.tb04234.x>

- Pina, J., Tilmant, A., & Côté, P. (2017). Optimizing multireservoir system operating policies using exogenous hydrologic variables. *Water Resources Research*, 53(11), 9845–9859. <https://doi.org/10.1002/2017WR021701>
- Poorepahy-Samian, H., Espanmanesh, V., & Zahraie, B. (2016). Improved inflow modeling in stochastic dual dynamic programming. *Journal of Water Resources Planning and Management*, 142(12), 04016065. [https://doi.org/10.1061/\(asce\)wr.1943-5452.0000713](https://doi.org/10.1061/(asce)wr.1943-5452.0000713)
- Potter, K. W. (1991). Hydrological impacts of changing land management practices in a moderate-sized agricultural catchment. *Water Resources Research*, 27(5), 845–855. <https://doi.org/10.1029/91WR00076>
- Powell, W. B. (2007). *Approximate dynamic programming: Solving the curses of dimensionality* (Vol. 703). John Wiley & Sons.
- Pritchard, G. (2015). Stochastic inflow modeling for hydropower scheduling problems. *European Journal of Operational Research*, 246(2), 496–504. <https://doi.org/10.1016/j.ejor.2015.05.022>
- Rabiner, L. R. (1989). A tutorial on hidden Markov models and selected applications in speech recognition. *Proceedings of the IEEE*, 77(2), 257–286. <https://doi.org/10.1109/5.18626>
- Rani, D., & Moreira, M. M. (2010). Simulation–optimization modeling: A survey and potential application in reservoir systems operation. *Water Resources Management*, 24(6), 1107–1138. <https://doi.org/10.1007/s11269-009-9488-0>
- Raso, L., Malaterre, P.-O., & Bader, J.-C. (2017). Effective streamflow process modeling for optimal reservoir operation using stochastic dual dynamic programming. *Journal of Water Resources Planning and Management*, 143(4), 04017003. [https://doi.org/10.1061/\(asce\)wr.1943-5452.0000746](https://doi.org/10.1061/(asce)wr.1943-5452.0000746)
- Rougé, C., & Tilmant, A. (2016). Using stochastic dual dynamic programming in problems with multiple near-optimal solutions. *Water Resources Research*, 52(5), 4151–4163. <https://doi.org/10.1002/2016WR018608>
- Sagna, P., Dipama, J., Vissin, E., Diomandé, B., Diop, C., Chabi, P., et al. (2021). Climate change and water resources in West Africa: A case study of ivory coast, Benin, Burkina Faso, and Senegal. In *Climate change and water resources in Africa* (pp. 55–86). Springer.
- Schwarz, G. (1978). Estimating the dimension of a model. *Annals of Statistics*, 6(2), 461–464. Retrieved from <http://www.jstor.org/stable/2958889>
- Shoelson, B. (2021). Deleteoutliers. Retrieved from <https://www.mathworks.com/matlabcentral/fileexchange/3961-deleteoutliers>
- Stedinger, J. R., & Griffis, V. W. (2011). Getting from here to where? Flood frequency analysis and climate 1. *JAWRA Journal of the American Water Resources Association*, 47(3), 506–513. <https://doi.org/10.1111/j.1752-1688.2011.00545.x>
- Stedinger, J. R., Sule, B. F., & Loucks, D. P. (1984). Stochastic dynamic programming models for reservoir operation optimization. *Water Resources Research*, 20(11), 1499–1505. <https://doi.org/10.1029/WR020i01p01499>
- Tan, W. L., Yusof, F., & Yusop, Z. (2017). Subseasonal to multidecadal variability of northeast monsoon daily rainfall over Peninsular Malaysia using a hidden Markov model. *Theoretical and Applied Climatology*, 129(1–2), 577–586. <https://doi.org/10.1007/s00704-016-1795-9>
- Tang, G., Zhou, H., Li, N., Wang, F., Wang, Y., & Jian, D. (2010). Value of medium-range precipitation forecasts in inflow prediction and hydropower optimization. *Water Resources Management*, 24(11), 2721–2742. <https://doi.org/10.1007/s11269-010-9576-1>
- Tejada-Guibert, J. A., Johnson, S. A., & Stedinger, J. R. (1993). Comparison of two approaches for implementing multireservoir operating policies derived using stochastic dynamic programming. *Water Resources Research*, 29(12), 3969–3980. <https://doi.org/10.1029/93WR02277>
- Thibault, A., Antil, F., & Boucher, M.-A. (2016). Accounting for three sources of uncertainty in ensemble hydrological forecasting.
- Tilmant, A., Goor, Q., & Pinte, D. (2009). Agricultural-to-hydropower water transfers: Sharing water and benefits in hydropower-irrigation systems. *Hydrology and Earth System Sciences*, 13(7), 1091–1101. <https://doi.org/10.5194/hess-13-1091-2009>
- Tilmant, A., Pina, J., Salman, M., Casarotto, C., Ledbi, F., & Pek, E. (2020). Probabilistic trade-off assessment between competing and vulnerable water users—the case of the Senegal River basin. *Journal of Hydrology*, 587, 124915. <https://doi.org/10.1016/j.jhydrol.2020.124915>
- Tilmant, A., Pinte, D., & Goor, Q. (2008). Assessing marginal water values in multipurpose multireservoir systems via stochastic programming. *Water Resources Research*, 44(12), W12431. <https://doi.org/10.1029/2008wr007024>
- Treisman, F., Maceira, M. E. P., Penna, D. D. J., Damázio, J. M., & Rotunno Filho, O. C. (2020). Synthetic scenario generation of monthly streamflows conditioned to the El Niño–Southern Oscillation: Application to operation planning of hydrothermal systems. *Stochastic Environmental Research and Risk Assessment*, 34(2), 1–23. <https://doi.org/10.1007/s00477-019-01763-2>
- Turner, S., & Galelli, S. (2016). Regime-shifting streamflow processes: Implications for water supply reservoir operations. *Water Resources Research*, 52(5), 3984–4002. <https://doi.org/10.1002/2015WR017913>
- United Nations Report. (2003). *World water assessment programme (united nations), water for people water for life*. United Nations Educational, Scientific and Cultural Organization (UNESCO) and Berghahn Books.
- Viterbi, A. J. (1967). Error bounds for convolutional codes and an Asymptotically optimum decoding algorithm. *IEEE Transactions on Information Theory*, 13(2), 260–269. <https://doi.org/10.1109/TIT.1967.1054010>
- Welch, L. R. (2003). Hidden Markov models and the Baum-Welch algorithm. *IEEE Information Theory Society Newsletter*, 53(4), 9–13.
- Whitcher, B., Byers, S. D., Guttorp, P., & Percival, D. B. (2002). Testing for homogeneity of variance in time series: Long memory, wavelets, and the Nile River. *Water Resources Research*, 38(5), 12. <https://doi.org/10.1029/2001WR000509>
- Whiting, J. P., Lambert, M. F., Metcalfe, A. V., Adamson, P. T., Franks, S. W., & Kuczera, G. (2004). Relationships between the El-Niño Southern Oscillation and spate flows in southern Africa and Australia. *Hydrology and Earth System Sciences Discussions*, 8(6), 1118–1128. <https://doi.org/10.5194/hess-8-1118-2004>
- You, J.-Y., & Cai, X. (2008). Determining forecast and decision horizons for reservoir operations under hedging policies. *Water Resources Research*, 44(11), W11430. <https://doi.org/10.1029/2008wr006978>
- Yu, J., Park, Y. J., Kwon, H.-H., & Kim, T.-W. (2018). Probabilistic assessment of meteorological drought over South Korea under RCP scenarios using a hidden Markov model. *KSCSE Journal of Civil Engineering*, 22(1), 365–372. <https://doi.org/10.1007/s12205-017-0788-2>
- Zhao, T., Yang, D., Cai, X., Zhao, J., & Wang, H. (2012). Identifying effective forecast horizon for real-time reservoir operation under a limited inflow forecast. *Water Resources Research*, 48(1), W01540. <https://doi.org/10.1029/2011wr010623>
- Zucchini, W., & Guttorp, P. (1991). A hidden Markov model for space-time precipitation. *Water Resources Research*, 27(8), 1917–1923. <https://doi.org/10.1029/91WR01403>
- Zucchini, W., MacDonald, I. L., & Langrock, R. (2017). *Hidden markov models for time series: An introduction using r*. Chapman and Hall/CRC.