



A stochastic policy algorithm for seasonal hydropower planning

Håkon S. Grini¹ · Anders S. Danielsen¹ · Stein-Erik Fleten¹ · Andreas Kleiven¹

Received: 2 December 2022 / Accepted: 12 July 2023
© The Author(s) 2023

Abstract

Hydropower producers need to plan several months or years ahead to estimate the opportunity value of water stored in their reservoirs. The resulting large-scale optimization problem is computationally intensive, and model simplifications are often needed to allow for efficient solving. Alternatively, one can look for near-optimal policies using heuristics that can tackle non-convexities in the production function and a wide range of modelling approaches for the price- and inflow dynamics. We undertake an extensive numerical comparison between the state-of-the-art algorithm stochastic dual dynamic programming (SDDP) and rolling forecast-based algorithms, including a novel algorithm that we develop in this paper. We name it Scenario-based Two-stage ReOptimization abbreviated as STRO. The numerical experiments are based on convex stochastic dynamic programs with discretized exogenous state space, which makes the SDDP algorithm applicable for comparisons. We demonstrate that our algorithm can handle inflow risk better than traditional forecast-based algorithms, by reducing the optimality gap from 2.5 to 1.3% compared to the SDDP bound.

Keywords Stochastic programming · Stochastic policy · Hydropower reservoir management

1 Introduction

Sequential decision-making in the presence of uncertainty is of relevance in many diverse applications, including inventory control problems, production scheduling, and system modelling [3]. Typically, the goal is to minimize cost or maximize profits by dynamically adapting decisions to exogenous information. These problems are often computationally intractable because of exponential growth in the number of scenarios needed to evaluate future expected costs or profits. Therefore, approximate

✉ Håkon S. Grini
hkong@gmail.com

¹ Norwegian University of Science and Technology, 7491 Trondheim, Norway

dynamic programming (ADP) algorithms are often used to obtain near-optimal policies [26].

In some applications, one can utilize the problem structure in developing computationally tractable algorithms, such as decomposition-based algorithms [18, 33]. The stochastic dual dynamic programming algorithm (SDDP), is an example [23]. This algorithm is state of the art for solving the seasonal hydropower production planning problem [12, 31]. The algorithm requires the problem formulation to be convex and uncertainty to be stage-wise independent [24]. Therefore, certain aspects of the real problem need to be approximated, to adhere to the convexity requirement. To handle non-convexities, some recent papers have proposed to extend the algorithm [11, 17], or to approximate the problem using McCormick envelopes [6]. To adhere to the requirement of stage-wise independence, serial correlation is typically accounted for by either imposing a linear dependence structure and increasing the state vector, or discretizing the state space before applying the algorithm [20]. The first approach greatly restricts the dynamics of exogenous factors, while the latter approach induces a discretization error.

An alternative class of algorithms that often provide good policies is forecast-based reoptimization heuristics, e.g. the rolling intrinsic (RI) heuristic [5]. The heuristic repeatedly solve deterministic problems based on the expectation of exogenous factors. When averaging revenues from sufficiently many sample paths, it provides a feasible policy and an estimated lower bound on maximization problems. Reoptimization heuristics are widely used by practitioners in various fields. In gas storage management, numerical experiments indicate that the rolling intrinsic algorithm provides near-optimal policies for price- and storage-dependent injections and withdrawals [19, 35]. Despite the ability to handle complex relationships, few works have assessed the performance of reoptimization-based heuristics in hydropower scheduling. An exception, albeit for short-term operations, is [21].

In this work, our main contributions are to provide an extensive comparison of the rolling intrinsic algorithm, the state-of-the-art algorithm SDDP, and a novel, heuristic-based, algorithm that does not impose any restrictions on the modelling of exogenous factors. We call this algorithm Scenario-based Two-stage ReOptimization (STRO). The main innovative aspect is that there is inherent stochasticity in the algorithm, and that there is no requirement that the policy yielded by the algorithm is deterministic. Still, the algorithm provides an estimate of the lower bound on the optimal value, from which marginal water values can be derived.

STRO samples possible future outcomes and solves two-stage programs repeatedly to make decisions. The idea is that this will better capture the range of possible realizations of exogenous factors, compared to the rolling intrinsic heuristic which uses a deterministic forecast based on expectations. Capturing the range of possible realizations is crucial in applications where all risk factors cannot be hedged, e.g. a hydropower producer which may risk spillage. In the setting of hydropower scheduling using SDDP, our algorithm allows for a more realistic representation of inflow dynamics, e.g. by using dynamic artificial neural networks [32], as no restrictions are imposed on the modelling of the exogenous factors.

Decision-making under uncertainty is also highly relevant in power system analyses [27, 34]. Our method is similar to the approach in [15, 16], where two-stage problems

are sequentially reoptimized along historical values of uncertain variables. We extend this method in several ways. First, we suggest to estimate the value function without requiring a deterministic policy. By doing so, we can sample a small subset of scenarios from which the two-stage problem is based on, and as such, control the degree of computational complexity. This makes our method capable of solving non-convex problems as well, as opposed to [15, 16], which is limited to the linear problems. Second, contrary to [15, 16], we evaluate the performance against the state-of-the-art algorithm SDDP, which is known to be optimal under certain restrictions. This gives us an upper bound on the profit maximization problem, which allows us to analyse different configurations of our suggested approach and provide useful insights about the representation of uncertainty in hydropower scheduling.

Another benefit of STRO is the opportunity for highly scalable parallel processing. Because scenarios are generated independently, their respective policies can be computed individually. The use of parallelization allows for a drastic reduction in the processing time required. As seen in previous works, traditional parallelization schemes do not scale to the same extent for SDDP [1, 14]. We provide an extensive numerical comparison of optimality gaps and computation time of the proposed algorithm, the state-of-the-art algorithm SDDP, and the rolling intrinsic algorithm.

The most important findings are that both the RI and STRO methods achieve revenues that are less than 3% from the estimate of the upper bound on optimal revenue based on several trial runs of the SDDP algorithm. Furthermore, by increasing the number of samples per decision for the STRO method, the performance of the algorithm increases substantially. When using 2 or more samples per decision, STRO outperforms RI by a decent margin and achieves revenues less than 2% from the estimated upper bound. The benefit of increasing the number of samples for the STRO heuristic is largest when moving from 1 to 2 samples, and the benefit from increasing the number of samples further was found to diminish quickly.

The paper is structured as follows. First, a detailed mathematical formulation of the problem is given, along with an in-depth problem description, in Sect. 2. Then the novel heuristic is presented in Sect. 3. We then give an illustrative example of how this heuristic works in Sect. 4. In Sect. 5 we give an overview of the specific case the algorithms were tested on, as well as the software and hardware used. Finally, the results of our experiments are outlined in Sect. 6 before we give our final remarks in Sect. 7.

2 Mathematical formulation

In this section, we provide a problem description and describe the stochastic processes and a mathematical model for the seasonal hydropower planning problem.

2.1 Problem description

We consider the medium-term reservoir management problem, where the objective is to value-maximize production over 1–2 years. The production decisions for medium-term reservoir management are typically of a weekly granularity.

Modern hydropower plants commonly consist of multiple interconnected reservoirs, that allow for coordinated water release that maximizes profit for the system as a whole. The state of the system at a given time is given by the amount of water that is stored in the different reservoirs. Furthermore, as regulations can put constraints on the water stored on a seasonal basis, these constraints can vary for different times t .

Several simplifications and assumptions are often made when modelling the reservoir management problem. This includes ignoring operating costs, as these variable costs are usually negligible compared to the revenue. This makes the problem a revenue-maximizing one. Another common simplification is to make the energy coefficient constant, to keep the problem convex. This coefficient is the factor that gives how much energy a production facility can get from a single unit of water. In reality, this factor is dependent on the head (the height difference between the turbine and the variable water level), as well as the intensity of the flow through the turbine. The head can vary substantially with the amount of water in the reservoir. In the sample system, the variations represented between 4.0% and 6.5% of the overall head for the different reservoirs, but these variations can be substantially higher or lower depending on the system. Since potential energy increases linearly with the head, this incentivizes producers to keep their reservoir levels high, thereby generating more electricity per unit of water. This poses an interesting dynamic, as the risk of spillage increases with higher water levels. To capture this dynamic, head variations are implemented in the model with a variable energy coefficient. This leads to bi-linearity which is accounted for using McCormick envelopes, similarly to [6].

2.2 Marginal water values

The main goal of the medium-term reservoir management problem is to estimate marginal water values. These values naturally depend on the amount of water in the system. To illustrate this, one can compare the additional value of an extra unit of water in a full reservoir and an empty reservoir. In the case of a full reservoir, one would have to discharge the unit immediately, independently of price, to avoid spillage. If the full reservoir is already producing at full capacity, the additional unit would naturally be spilled, and the marginal value would be 0. For the other extreme case where the reservoir is empty, the marginal value would be much higher as the producer could wait until the prices are high to discharge the unit, without having to worry about spillage. In practice, that means that the marginal value of water decreases, as the water level in the reservoir increases.

The marginal water value represents the current alternative cost of discharging a unit and is therefore used to make short-term production decisions. The way the medium-term reservoir management problem relates to this is that one can calculate the expected discounted revenue over a 1–2 year horizon starting with different water levels. If one plots the expected revenue against the starting water level, the slope of the curve would represent the current marginal water value for different water levels.

2.3 Nomenclature

Before formally defining the optimization model, we provide the notation used in the following section. Some of the notation is only used for the additional constraints seen in the appendix.

Sets and indices

R	Reservoirs, indexed by i and j
R^P	Reservoirs with electricity production
R_i^{Dis}	Upstream reservoirs releasing water into reservoir i
R_i^{Spill}	Upstream reservoirs spilling water into reservoir i
M_i	Set of points on the piecewise linear curve representing the relation between head and water level for reservoir with electricity production i , indexed by k
Π	Set of feasible policies
S	Set of level 1 scenarios, indexed by s
\mathcal{N}	Set of level 2 scenarios, indexed by n
T_t	Set of stages from stage t to and including the end of the horizon, indexed by τ

Decision variables

$x_{i,t}$	Discharge from reservoir i in period t	(m ³)
$r_{i,t}$	Slack variable for water spillage from reservoir i in period t	(m ³)
$l_{i,t}$	Water level in reservoir i at the beginning of time period t	(m ³)
$l_{i,t}^{avg}$	Average water level in reservoir i through time period t	(m ³)
$h_{i,t}$	Water head of reservoir i in period t	(m)
$w_{i,t}$	Substitution variable for the bi-linear term $h_{i,t}x_{i,t}$	(m)
$\lambda_{i,t,k}$	Weight of point k on the piece-wise linear graph describing the relation between water level and head in reservoir i at time t	

Parameters

T	Stages (time periods) in planning horizon	
S	Number of level 1 (outer) scenarios	
N	Number of level 2 (inner) scenarios	
β	Discount factor	
η_i	Constant factor describing the efficiency of the turbine of reservoir i	(MWh/m ³)
g	Gravitational acceleration constant	(m/s ²)
ρ	Density of water	(Kg/m ³)
P_t	Power price in stage t	(€/MWh)
$I_{i,t}$	Water inflow intensity for reservoir i in period t	(m ³)

2.4 Stochastic processes

Restructured energy markets have uncertain prices, denoted P_t , which means that operators have to optimize over stochastic prices. Furthermore, as weather cannot be forecasted perfectly, the inflow to reservoirs is also uncertain. The set of all reservoirs is denoted by $R = \{1, 2, \dots, |R|\}$. We denote inflow to reservoir $i \in R$ at time t by $I_{i,t}$.

For prices, we apply the Schwartz-Smith two-factor approach [28]. This assumes that the logarithm of the price can be cast as the sum of a short-term factor, a long-term factor, and finally a seasonal factor, as the energy price typically follows a seasonal pattern. The parameters are estimated using Kalman filtering and maximum likelihood estimation [13]. Data points are obtained from synthetic futures curves [9] using the approach in [2]; see also [9].

The inflow specification is primarily based on the one seen in [12], where the inflow is modelled according to AR-1 processes. To reduce dimensionality and capture the covariance of inflow to different reservoirs, a PCA transformation of the inflow data is performed to project it to a lower dimensionality, as seen in [8]. The original inflow data is normalized to follow a standard normal distribution, before being transformed using PCA. After fitting AR-1 models to the lower dimension data, one can produce scenarios using the model for every dimension k in the projected space. After generating the low-dimensional scenarios, we can transform the generated scenario back to the original dimensionality, using the PCA matrix. We then denormalize this data, to ensure that the generated scenarios follow the distribution of the historical observations. To circumvent the issue of negative inflows generated by the AR model, negative values are converted to 0 according to Eq. (1).

$$I_{i,t} = \max(I_{i,t}, 0) \quad (1)$$

2.5 Actions and transition function

We present here stage- t actions and transition functions. As only a subset of these reservoirs have connected turbines allowing for energy production, we denote this by subset $R^P = \{1, 2, \dots, |R^P|\}$. These sets are both indexed by i and j , for the modelling of interconnections between reservoirs. Every reservoir i , has two sets of reservoirs, R_i^{Dis} and R_i^{Spill} , that represent the reservoirs whose discharged and spilled water flows directly into reservoir i , respectively.

The vector of decisions in stage t includes discharge $x_{i,t}$ and spillage $r_{i,t}$ for each station and reservoir, where $\Pi_t(I_t, I_t)$ is the feasible stage- t action set and $(x_{i,t}, r_{i,t}, i \in R) \in \Pi_t(I_t, I_t)$. The discharge variable has both upper and lower bounds depending on physical limitations and regulations that can be seasonal. The spillage and the discharge can naturally not be negative.

The transition function consists of three components (I_t, P_t, I_t) . The exogenous components, includes price P_t and the vector of inflows, $I_t = (I_{i,t}, i \in R)$. These factors get updated according to the stochastic processes discussed in the previous

section, independently of the stage- t action. The endogenous component is the vector of reservoir volumes in the watercourse, $l_t = (l_{i,t}, i \in R)$. Executing an action (x_t, r_t) at stage t and state (l_t, P_t, I_t) leads to the following update of the endogenous state,

$$l_{i,t+1} = l_{i,t} + I_{i,t} - r_{i,t} - x_{i,t} + \sum_{j \in R_i^{Spill}} r_{j,t} + \sum_{j \in R_i^{Dis}} x_{j,t}, \tag{2}$$

$$i \in R \quad t = 0, \dots, T - 1$$

This function ensures water balance in time and the topology of the watercourse. It gives the relationship between the inflow, production decisions, spillage, and the water level of a given reservoir i . It ensures that spillage happens when the net inflow is higher than the reservoir capacity. Due to the longer time steps taken with medium-term reservoir management, water delay is not taken into account, and it is assumed that water can flow through the entire system within a single time step.

2.6 Immediate reward and policy

At each stage, the immediate reward is given by electricity production and electricity price (3)

$$P_t \sum_{i \in R^p} g(l_{i,t}, x_{i,t}), \tag{3}$$

where $g(l_{i,t}, x_{i,t})$ is the energy output from discharge $x_{i,t}$ at reservoir level $l_{i,t}$. We denote the set of feasible policies by Π . A policy π is a collection of stage-dependent actions, mapping states at time t to feasible actions. We let $l_{i,t}^\pi$ and $x_{i,t}^\pi$ respectively denote the endogenous state reached at, and action made in, stage t for reservoir i , when following policy π . We aim to find a policy that maximizes the expected accumulated discounted reward,

$$V_0(l_0, I_0, P_0) = \max_{\pi \in \Pi} \mathbb{E} \left[\sum_{t=0}^{T-1} \beta^t P_t \sum_{i \in R^p} g(l_{i,t}^\pi, x_{i,t}^\pi) + \beta^T \sum_{i \in R} \Phi_i(l_{i,T}^\pi) \middle| I_0, P_0 \right], \tag{4}$$

where β is the discount factor and $\Phi(l_{i,T})$ is the end of horizon value for each reservoir, given by how much water is remaining.

2.7 Relaxation and approximations

The reward function contains the term $g(l_{i,t}, x_{i,t})$. This function is typically concave in $l_{i,t}$ and $x_{i,t}$. The higher the reservoir, the higher the generation output. Similarly, as a function of discharge, the efficiency of a turbine is first increasing and then possibly decreasing. Piece-wise linear approximations of the discharge function can handle the latter feature. In our model, we assume it to be a constant factor for each reservoir i denoted $\eta_i \in [0, 1]$, i.e. $g(l_{i,t}, x_{i,t}) = g(l_{i,t})x_{i,t}\eta_i$. The dependency of reservoir volume, i.e. head variations, is complicated since it leads

to a bi-linear objective and thereby non-convexity. This makes it challenging to optimize with conventional methods such as SDDP, which we use as a baseline. To deal with this issue, we relax the problem using McCormick envelopes. We start by modifying the reward in (3) to

$$P_t \sum_{i \in R^P} h_i(l_{i,t})x_{i,t}g\rho\eta_i \tag{5}$$

The functions $h_i(\cdot)$ represent the head in reservoir i , while g is the gravitational acceleration constant and ρ is the density of water. The $h_i(\cdot)$ functions depend on $l_{i,t}$ and are calculated using piece-wise linear approximations for the individual reservoirs. These piece-wise linear graphs are represented as sets of points, where the x-coordinate represents the water level and the y-coordinate represents the head. These sets are denoted M_i , for every reservoir i . To find the head $h_{i,t}$ given any feasible water level $l_{i,t}$, we introduce the variables $\lambda_{i,t,k}$ representing the weight of every point $k \in M_i$ at time t . By ensuring that all these weights are non-negative and sum to 1, and that at most two points adjacent to each other can be non-zero, we can find the exact point on the y-axis on the piece-wise linear graph, given the value on the x-axis.

The restriction requiring that at most two weights can be non-zero, and that these must represent points adjacent to each other, is called a special ordered set of type 2 (SOS2) restriction. This restriction can be ignored in cases where the piece-wise linear function is concave, as the head $h_{i,t}$ would be the same with or without the SOS2 constraint. The approximated graphs for the energy-producing reservoirs in the Matre Haugsdal system can be seen in Fig. 1. It can be seen that these are all concave, ensuring that the SOS2 constraints are unnecessary.

The piece-wise linear approximation of the head function $h_i(l_{i,t})$ is explained in further detail in Appendix A. This approximation leads to bi-linear terms in (5). To approximate these terms, we introduce the substitution variable $w_{i,t} = h_i(l_{i,t})x_{i,t}$ and use McCormick envelope constraints. This gives us the modified linear reward

$$P_t \sum_{i \in R^P} w_{i,t}g\rho\eta_i, \tag{6}$$

which leads to a convex and linear multistage optimization problem, satisfying the requirements of SDDP. For additional constraints associated with the McCormick relaxation, see Appendix B.




Fig. 1 Volume-head relations for reservoirs with production

3 Scenario-based two-stage re-optimization heuristic

One of the main ideas behind rolling horizon methods, such as RI, is that a finite number of sampled scenarios can be used to approximate the sample space of a continuous-state problem. The RI heuristic has shown promising performance in various applications and is simple and intuitive. We utilize a similar concept as the RI heuristic in developing a novel Scenario-based Two-stage Re-Optimization heuristic (STRO). Instead of using expectations of exogenous variables as forecasts from any given state, we propose to generate possible scenarios from the stochastic model and repeatedly solve two-stage problems based on these scenarios.

To distinguish between the outer scenarios that are the realizations of the environment that are unknown to the heuristic, and the inner scenarios that are generated to build the two-stage stochastic programs that the heuristic solves to find the first-stage decision. We refer to the outer scenarios as level 1 scenarios and the inner ones as level 2 scenarios.

We introduce the following notation for formulating the two-stage problem used by the STRO scheme. Let (I_t^s, P_t^s) denote the realization of stochastic variables in stage t of level 1 scenario s . Let I_t^s denote the endogenous state reached at time t by following the STRO scheme along level 1 scenario s . We denote the decision variables for discharge and spillage from reservoir i in stage t and level 1 scenario s as $x_{i,t}^s$ and $r_{i,t}^s$ respectively. Then, let $(\bar{I}_t^n, \bar{P}_t^n) = \{(\bar{I}_t^n, \bar{P}_t^n), (\bar{I}_{t+1}^n, \bar{P}_{t+1}^n), \dots, (\bar{I}_T^n, \bar{P}_T^n)\}$ denote a realization of level 2 scenario n generated according to the stochastic processes, from time $t + 1$ to end of horizon T , conditional on the level 1 realization at time t . Further, let the variables $\bar{I}_{i,t}^n, \bar{x}_{i,t}^n, \bar{r}_{i,t}^n$ represent the decision variables for the endogenous state, water discharge and spill for reservoir i at time t in level 2 scenario n . Let $\mathcal{T}_t = \{t, t + 1, \dots, T\}$ be the set of remaining periods from time t to time T . For each stage t and each level 1 scenario s , the two-stage problem that is being solved can be formulated as

$$\begin{aligned}
 & \max_{\{\bar{I}_t^n, \bar{x}_t^n, \bar{r}_t^n\}, n \in \mathcal{N}, \tau \in \mathcal{T}_t\}} \frac{1}{N} \sum_{n \in \mathcal{N}} \left(\sum_{\tau \in \mathcal{T}_t \setminus \{t\}} \beta^\tau \bar{P}_\tau^n \sum_{i \in R^p} g(\bar{I}_{i,\tau}^n, \bar{x}_{i,\tau}^n) + \beta^T \sum_{i \in R} \Phi_i(\bar{I}_{i,T}^n) \right), \\
 & \text{s.t. } (I_{i,t}^s, P_t^s, I_t^s) = (\bar{I}_{i,t}^n, \bar{P}_t^n, \bar{I}_t^n), \quad n \in \mathcal{N}, i \in R \\
 & \bar{I}_{i,\tau+1}^n = \bar{I}_{i,\tau}^n + \bar{I}_{i,\tau}^n - \bar{r}_{i,\tau}^n - \bar{x}_{i,\tau}^n + \sum_{j \in R_i^{spill}} \bar{r}_{j,\tau}^n + \sum_{j \in R_i^{Dis}} \bar{x}_{j,\tau}^n, \tau \in \mathcal{T}_t \setminus \{t\}, i \in R, n \in \mathcal{N} \\
 & \quad x_{i,t}^s = \bar{x}_{i,t}^n, \quad n \in \mathcal{N}, i \in R \\
 & \quad r_{i,t}^s = \bar{r}_{i,t}^n, \quad n \in \mathcal{N}, i \in R
 \end{aligned} \tag{7}$$

The objective of the two-stage problem is to maximize expected future profits in the discrete set of (level 2) scenarios \mathcal{N} . In this formulation, every scenario in \mathcal{N} is considered equally likely, but it is also possible to introduce a separate weighting for each scenario if they have different likelihoods. The first constraint ensures that the initial condition is set to the realized time t level 1 scenario s and the endogenous state reached in that scenario by following the STRO scheme. The next constraint is the reservoir balance in level 2 scenario n , which is equivalent to (2). The last 2 constraints ensure that the first-stage decisions are the same in all level 2 scenarios $n \in \mathcal{N}$ and that the production decisions that will be used in level 1 scenario s are the same as the first-stage decisions in the LP.

Each time this problem is being solved, the time t decisions, i.e. first-stage decisions, $(x_{i,t}^s, r_{i,t}^s)$ are used to update $l_{i,t}^s$ along the level 1 scenario s according to Eq. (2). This updated state then serves as an initial condition for the next two-stage problem.

By generating S level 1 scenarios, the total expected profits from the STRO scheme can be computed as the sample average

$$\frac{1}{S} \sum_{s \in S} \left(\sum_{t=0}^{T-1} \beta^t P_t^s \sum_{i \in R^p} g(l_{i,t}^s, x_{i,t}^s) + \beta^T \sum_{i \in R} \Phi_i(l_{i,T}^s) \right) \tag{8}$$

A visualization of how a two-stage stochastic program is built from generated level 2 scenarios to explore the underlying sample space can be seen in Fig. 2. Starting in a state in stage τ , one can generate N possible realizations of the future, starting in the current state. Then, these scenarios can be combined into a two-stage stochastic program, which can be solved to find the first-stage decision.

Given that stochastic programs are built for every state in every level 1 scenario, the idea is to approximate the continuous sample space of the exogenous variable from each state. The higher number of level 2 scenarios N the more accurate this approximation will be. This approximation is illustrated in Fig. 3. The main difference compared to RI is that STRO makes multiple explorations (level 2 scenarios) from every state in every underlying level 1 scenario and uses these to approximate the sample space from that point. The red nodes in the figure represent the level 2 scenarios. New information about the underlying level 1 scenario is revealed with each step of the horizon. Thus, new two-stage stochastic programs must be built for every state by generating new level 2 scenarios.

None of the decisions made by STRO's heuristic are based on any knowledge of the future in the underlying price-inflow scenario. Following this, the decisions made using the heuristic are implementable and the value of the feasible policies found for every individual level 1 scenario yields lower bounds for the optimal revenue for that given scenario. Same as for RI, this means that




Fig. 2 Two-stage stochastic program built from N level 2 scenarios. These are solved to find an implementable decision for stage t in a level 1 scenario




Fig. 3 Sample space approximation for STRO with N level 2 scenarios per state

with enough sampled scenarios, one will get an estimated lower bound of the expected revenue, subject to sampling error.

As every generated level 2 scenario represents possible realizations of the sample space, one would gradually get a better approximation of all possible realizations of the future in all possible states by generating more scenarios per state.

As mentioned, the STRO heuristic finds the first-stage decision by considering these realizations as a two-stage program. Using two stages is an approximation, and a more accurate method would be to solve a multistage problem instead. However, building and solving multistage stochastic programs is more computationally costly, and approximation through two stages has been shown to be a good solution for hydropower problems in the literature [29].

Due to the growth in complexity when increasing the number of level 2 scenarios per state, one can in practice only build programs based on a small number of scenarios. Therefore, the performance of STRO hinges on the assumption that most of the value from the approach can be achieved using a small number N of level 2 scenarios, e.g. $N = 2$ or $N = 3$, which from early experiments seems to be the case.

4 Complexity

Considering that the STRO heuristic is based on solving a two-stage stochastic program for every production decision, the amount of programs solved is quite large. Given S level 1 scenarios and T stages per level 1 scenario, one will have to solve a total of ST LPs to get an estimate of the marginal water value.

Note that the number of LPs solved is not dependent on the amount of level 2 scenarios N . However, this instead affects the complexity of each LP, meaning that an increase in level 2 scenarios makes each individual LP harder to solve.

5 Main innovations and benefits

As mentioned in Sect. 2.1, medium-term hydropower planning is mainly concerned with calculating the expected discounted revenue, given different initial water levels, to calculate the marginal value of the water in the reservoir. Consequently, there is no requirement to generate a policy that deterministically yields the same production decision when faced with the same state, such as RI and SDDP. Instead, we forego this requirement and allow the STRO heuristic to yield stochastic production decisions based on the level 2 scenarios drawn from the sample space. We refer to this as it being a stochastic policy. The benefit is that by not having to aggregate the possible outcomes, for instance, by taking the expectation such as RI, we can explore more of the sample space and consider more varied outcomes. As shown in the illustrative example in Sect. 4, this can lead to more accurate revenue estimations.

Another significant benefit of STRO and other rolling horizon methods is that they pose few constraints on the problem class. Unlike SDDP, which either requires inflow and price to be modelled linearly or with a Markov decision process over discretized states, this allows price and inflow to be modelled using non-linear methods such as deep learning-based methods. There is also no strict constraint that requires the problem to be modelled convexly, meaning that one for instance, could model head variations more accurately without having to relax the problem with McCormick envelopes. However, a non-convex model would require a longer processing time.

Other significant benefits of rolling horizon methods such as STRO and RI is that they are significantly more straightforward to implement than SDDP and are more suited for parallelization since the level 1 scenarios are all solved independently. Although SDDP is also parallelizable, it is harder to implement and does not lead to as linear of a speedup, as seen in previous works. One example is [1] which reports speedup factors of 7–11 for 10 CPUs, and 9–14 for 20 CPUs, for a case study of the Brazilian hydrothermal system. For this case, the speedup is measured using 10 forward pass scenarios for 10 CPUs vs a single forward scenario for one CPU. Using a more detailed representation of the Brazilian power system, [22] report speedup factors of 5–6.5 for 10 CPUs and 8–11 for 20 CPUs. Here the speedup factor is measured using the same number of scenarios for the parallel and serial runs.

We leave to further work to explore the practical merits of these benefits and concentrate our experiments on a case where SDDP can be used as a benchmark.

6 Illustrative example

This section illustrates the proposed heuristic using a simple three-stage example. Consider a plant operator with a current reservoir volume of 8, maximum reservoir capacity of 10, and maximum generation capacity of 10. In this setting,

Fig. 4 Illustrative example inflow tree




Table 1 Solutions and objective values

	Optimal	Rolling intrinsic	STRO(1)	STRO(2)	STRO(3)	STRO(4)
x_0	1	0	{0, 0, 1, 1}	{0, 1, 1, 1, 1, 1}	{1, 1, 1, 1}	1
$x_{1,1}$	3	2	–	–	3	3
$x_{1,3}$	0	0	–	–	3	0
$x_{2,1}$	10	10	–	–	10	10
$x_{2,2}$	8	9	–	–	8	8
$x_{2,3}$	9	10	–	–	9	9
$x_{2,4}$	8	9	–	–	8	8
Mean spill	0	$\frac{3}{4}$	$\frac{1}{2}$	$\frac{1}{12}$	0	0
Mean profit	131.5	125.0	127.0	130.83	131.5	131.5

Note that STRO does not yield a deterministic policy since there are several decisions at the first stage

the chronology is as follows: First, inflow realizes, then the reservoir volume gets updated. Spillage happens if the current reservoir level plus inflowing water exceeds the reservoir capacity. Then the generation decision is taken before the new inflow arrives.

The price is deterministic and given by $P_0 = 10$, $P_1 = 11$, $P_2 = 12$, which incentivizes later production. Inflow, on the other hand, is stochastic and the different inflow scenarios can be seen in Fig. 4. From every state in this figure, there is an equal chance that the subsequent inflow is high or low compared to the previous inflow. This results in a binary tree and 4 different scenarios since inflow is deterministic in the first stage.

As an example: Since the current reservoir volume is 8 and 1 unit of inflow arrives deterministically in the first stage, no spillage happens at this point, as the reservoir can hold 10 units of water. Suppose then that the producer decides to generate nothing at $t = 0$, and that the high inflow is realized at $t = 1$ (scenario 1 or 2), meaning that 2 units arrive in addition to the 9 already in the reservoir. This would result in 1 unit of spillage.

To see how RI and STRO with different numbers of level 2 scenarios perform for this example, we compare their different decisions and revenues in Table 1.

We also compare them to the optimal strategy. Here we use $x_{i,j}$ to denote the production decision at time i and state j , numbered according to the scenarios in Fig. 4. The scenario index for stage 0 is omitted, as all these decisions need to be identical. Similarly, in stage 1, we only use scenario indices $\{1, 3\}$ since there only are 2 possible inflow realizations at $t = 1$. Index 3 is used to stay consistent with the indices in stage 2.

We observe that it is optimal to produce 1 unit in stage 0 to avoid spillage in the future. Rolling intrinsic, which bases decisions on expectations of the stochastic variables, ignores that there is a 50% chance of 2 units of inflow getting realized in the next stage and only considers the expected inflow, which is $\frac{0+2}{2} = 1$, thereby deciding to produce 0 at $t = 0$.

The notation STRO(N) is used in the table to indicate that N level 2 scenarios are used per state for the STRO heuristic. Since there are only 4 possible realizations of the future, we use a maximum N of 4. For STRO(1), STRO(2) and STRO(3) the production decision naturally depends on which of the 4 scenarios are sampled. Since we sample the scenarios without replacement, the number of combinations of level 2 scenarios is $\binom{4}{N}$, which results in 4, 6, 4 and 1 combinations for $N = 1, 2, 3$ and 4 respectively. The production decisions in the first stage are given in the table for all different combinations of level 2 scenarios in curly brackets. To deal with uncertain production decisions, we list the production decisions in the first stage for each combination of level 2 scenarios for every STRO(N) in curly brackets. Since the production decision at $t = 0$ is uncertain for $N = 1$ and $N = 2$, the next state is also uncertain (STRO(3) yields the same production decision at $t = 0$ independent of the combination of level 2 scenarios). For this reason, the production decisions at $t = 1$ and $t = 2$ are not listed for these, as these decisions naturally will depend on the previous decision, which is not deterministic. This relates to the discussed idea of the policy being stochastic.

The table shows that STRO(4) and STRO(3) both find the optimal strategy. This is unsurprising for STRO(4) as it considers all 4 possible scenarios when making decisions. STRO(3) does not have this privilege, but since it only omits a single possible level 2 scenario at $t = 0$, it will always sample at least one of the 2 scenarios with high inflow at $t = 1$. This allows it to see that there is a chance of spillage occurring and thereby plan its production accordingly.

It can further be seen that STRO(1) and STRO(2) both outperform RI, despite not necessarily seeing that there is a chance of high inflow at $t = 1$. As STRO(2) will sample 2 level 2 scenarios per state it will not consider the chance of high inflow only in the case that the 2 level 2 scenarios are both low inflow scenarios. There is only a $\frac{1}{6}$ chance that this happens, making it unlikely that the wrong production decision is made. When making a decision at $t = 1$ there are only 2 possible future inflow realizations, depending on the previously observed inflow. As a consequence of this, STRO(2) will make the optimal decisions at $t = 1$ since it can consider both possible future realizations from that point. STRO(1) only samples a single level 2 scenario per state. At $t = 0$ it is only a 50% chance that it samples a high inflow scenario and adjusts production to avoid risking spillage. However, it still outperforms RI, because RI will only consider the expected inflow of 1, meaning that it will always make the wrong production decision at $t = 0$. The same effect would

be seen if a high inflow scenario was realized at $t = 1$. Then RI would not take the fact that a high inflow of 3 could occur at $t = 2$ into account and only consider the expected inflow of 2. This results in a higher chance of spillage.

Although STRO(1) outperforms RI in this illustrative example, it must be stressed that this is not the case in general, and under different circumstances RI has been found to outperform the simplest version of STRO, as we will demonstrate in the next section.

7 Hydropower case and implementation

7.1 Implementation environments

7.1.1 Software

All code is implemented and run in Python 3.8, except the parameter estimation of the price model, which is done in MATLAB2020b [13]. In Python, the `gurobipy` package is used for the optimization models. A previous Python implementation of SDDP was also used [10].

7.1.2 Hardware

The computational tests are conducted on a system that allows for running up to 40 processes in parallel per node. In Table 2, the relevant technical aspects of the hardware are given.

7.2 Matre Haugsdal and Vemundsbotn

This paper is written in cooperation with a major Norwegian hydropower producer. They have provided data for one of their hydropower systems. Matre Haugsdal and Vemundsbotn are power plants in western Norway. The plants are parts of the more extensive system of interconnected reservoirs built around the Matre watercourse. Together they produce enough electricity to cover the needs of more than 50k households. An overview of the system's layout as modelled can be seen in Fig. 5, and the publicly available technical specifications for the two production plants are given in Table 3.

Table 2 Hardware specifications

Specification	Details
Node system	Dell PowerEdge R640
Node CPU	2x 2.4 GHz Intel Xeon Gold 5115 CPU - 10 core
Node RAM	96Gb

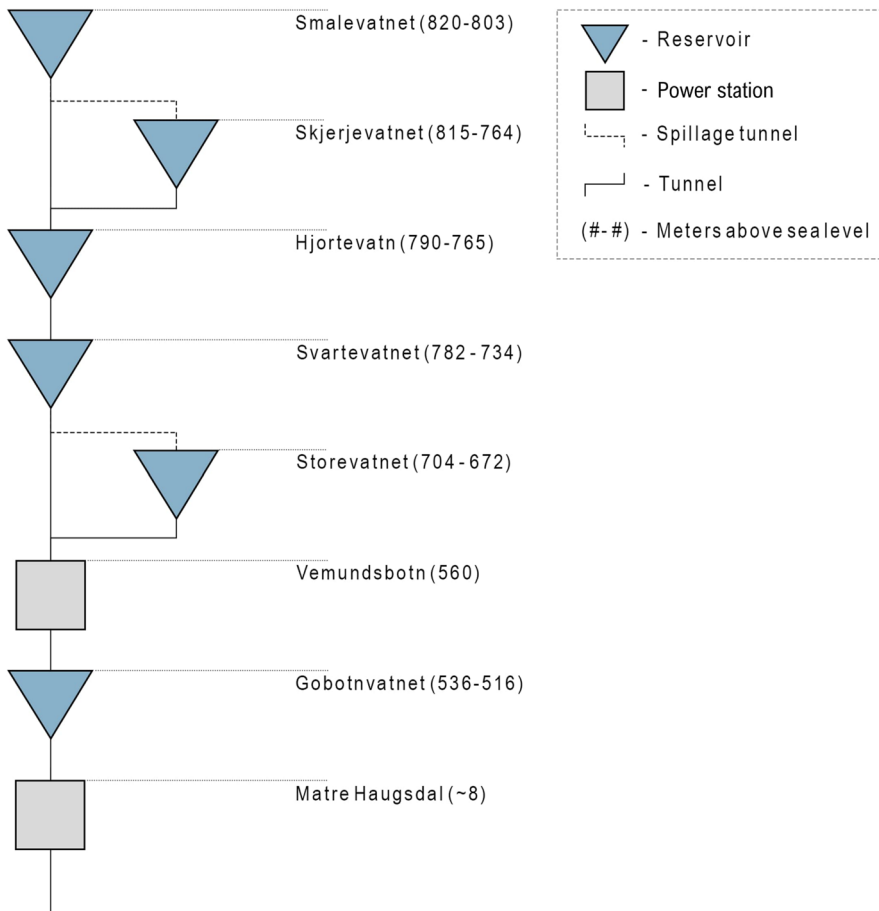


Fig. 5 Matre Haugsdal and Vemundsbotn, with connected reservoirs

Table 3 Specifications for plants in the considered system

	Matre Haugsdal	Vemundsbotn
Build year	2016	1976
Turbines	2	1
Production capacity (MW)	180	45
Elevation (m)	537	(164/242) ^a
Head difference (m)	20.5	(32/48) ^a

^aTwo connected reservoirs (Storevatnet/Svartevatnet)

Table 4 shows the average annual inflows for each reservoir, along with their respective volumes and discharge capacities. A glance at the data shows that the reservoirs are of a “fast” nature, meaning that the inflow to the reservoirs is relatively high compared to the storage capacity. We were not given any specified minimum

Table 4 Reservoir specifications

Reservoir	Volume (Mm ³)	Inflow (Mm ³ /year)	Max. discharge (m ³ /s)
Gobotvatn	22.0	134.6	37.6
Storevatn	23.3	57.6	18.0
Svartevatn	60.8	66.8	23.6
Skjerjevatn	70.5	61.3	22.0
Hjortevatn	10.7	60.9	10.0
Smalevatn	13.1	96.5	19.8

discharge for the reservoirs in question, and thus, no lower bound is set for the model. Furthermore, the water level regulations for the reservoirs do not vary with season, but are constant throughout the year.

7.3 Price and inflow data

Our price model uses futures data, and to generate the smoothed future curves, we use futures price data for the trading period 2006–2018, retrieved from Montel. The price simulations start 1 January 2019. The data is given per trading day, with closing prices for monthly, quarterly, and annual delivery for up to 3 years ahead. To fit the inflow model, we use time series inflow data for the reservoirs. The inflow data series provided are of hourly resolution in the period 2009–2019.

8 Computational study

To see how the STRO method with different number of level 2 scenarios per decision performs, the results are benchmarked to the SDDP method, as well as the most well-known rolling horizon method RI. To avoid disclosing any non-public information about the Matre Haugsdal hydropower system, the results of the models are represented as percentages relative to the baseline given by SDDP.

For these studies, revenue has been used as the primary metric for performance comparisons between different algorithms. Seeded level 1 scenarios were used for the rolling horizon methods, to make the revenues as comparable as possible. This means that all rolling horizon methods receive the same realizations of inflow and price. This ensures that the objective values of the rolling horizon methods are comparable, as they solve the exact same scenarios. However, this was not achievable for SDDP due to its discretization. We address this potential source of bias by using a relatively large number of samples for both methods over 10 runs. The planning horizon for each method is 2 years with weekly stages resulting in 104 stages in total.

The results presented are average values from all these runs, along with relative standard deviations of the objective values achieved over different trials using the

different methods. This approach should give a reasonable estimate of the difference in performance.

Note that since SDDP yields upper bounds while RI and STRO yield estimated lower bounds, there will always be some discrepancy between their results.

The discretizations are done on a sample base of 200k samples of each exogenous variable, which were discretized to a Markov chain with 125 discrete states per stage. Each of these stages contains a combination of price and inflow. No further scenario reductions were applied for the SDDP method.

The results from the trial runs are presented in Table 5. It presents the expected revenue for the different methods, and the standard deviation over the different runs for each method. Every run of the Rolling Horizon methods was made using 1000 level 1 scenarios, and reported revenue is the average over these runs.

From these results there are multiple takeaways. First, it is apparent that increasing the number of samples per decision in the heuristic leads to a substantial improvement in the results achieved. When comparing the results of STRO for $N = 1$ to $N = 7$ it can be seen that the gap to the SDDP algorithm is approximately halved. Furthermore, it can be seen that the improvement in the results is diminishing, and the most significant improvement can be seen when increasing N from 1 to 2. However, when comparing the results for $N = 6$ to $N = 7$, the gain has tapered out.

Another takeaway is that the STRO algorithm clearly outperforms the traditional RI algorithm for $N \geq 2$. The results of the RI algorithm were found to be within 2.5% of the estimated optimum from SDDP, in line with previous findings from [4]. However, the results from the STRO(7) algorithm, is within 1.4% of the results of SDDP. Given that the estimated revenues from the rolling horizon methods, STRO and RI, are lower bounds, whereas SDDP gives an upper bound, this gap is quite close, and it is apparent that the STRO method is well suited for the hydropower planning problem.

As seen in the table, the standard deviations of the results achieved using different methods are quite high compared to the differences in the results achieved with

Table 5 Average objective values of the compared methods

Solution method	Revenue	
	Relative value (%)	Standard deviation (%)
SDDP	100.000	0.487
RI	97.546	0.489
STRO(1)	97.249	0.484
STRO(2)	98.115	0.487
STRO(3)	98.411	0.489
STRO(4)	98.548	0.486
STRO(5)	98.611	0.485
STRO(6)	98.654	0.482
STRO(7)	98.674	0.483

the different algorithms. This implies that the true differences in revenues between SDDP and the other methods can be slightly higher or lower than these experiments indicate. Since seeded scenarios were used for RI and STRO we can be more confident that the comparisons between these methods are more accurate. In future experiments, we will make efforts to reduce the standard deviations by using more level 1 scenarios to ensure more stable revenue estimates over the different trials.

8.1 Runtime analysis

Rolling Horizon methods are well suited for parallelization and multiprocessing, as each level 1 scenario is processed independently. To minimize the processing time, our implementation runs 40 processes in parallel, which was the maximum allowed by our hardware. With improved hardware, a further speedup is naturally achievable. There exist parallel implementations of SDDP [1, 7, 25]. However, as [1] points out, parallel implementation of SDDP is far from straightforward. Therefore, we opt for a non-parallel version for benchmarking. With such parallel methods implemented, we would naturally have a substantially lower processing time with SDDP.

In Table 6, the run times of SDDP and the STRO methods are presented. RI solves a problem of the same size as STRO(1), the computational times for these methods are, therefore, approximately the same. The total run time is given in hours (h), and the average time used to solve each level 1 scenario is given in seconds (s). Note that the run time of SDDP is given without the additional time required for sampling and discretization.

The table shows that the most significant increase in run-time comes with the increase from 2 to 3 level 2 scenarios per decision in STRO, followed by a nearly linear increase in computational time. Once the scenarios are sampled and given to the two-stage stochastic program, the program becomes deterministic [30]. By doubling the number of scenarios for the two-stage program, the program size approximately doubles.

We stress that 2000 iterations were significantly more than what the algorithm needed to converge. In most cases, the upper bound only decreased by 0.2% after the

Table 6 Run time for 1000 level 1 scenarios of STRO(M), and 2000 iterations of SDDP

Solution method	Time	
	Total (h)	Per sampled scenario (s)
SDDP	11.38	
STRO(1) / RI	0.26	0.94
STRO(2)	0.73	2.63
STRO(3)	3.47	12.48
STRO(4)	5.94	21.38
STRO(5)	8.19	29.49
STRO(6)	10.58	38.10
STRO(7)	11.88	42.77

initial 200 iterations and by 0.1% after the first 500 iterations. Admittedly, it would have been reasonable to terminate the algorithm after 1000 iterations as the final 1000 iterations only yielded a $\sim 0.03\%$ decrease of the upper bound. Our reasoning behind using so many iterations was that we wanted the upper bound of SDDP to get as low as possible to make the results as comparable to the lower bound provided by the rolling horizon methods as possible. However, this makes the provided runtime for SDDP higher than what it could have been.

It should also be noted that 125 discrete states per stage for the SDDP method is a relatively small amount, and with an increase in states there would have been an increase in processing time.

9 Concluding remarks

We propose a new approach to solving the medium-term reservoir management problem. The aim is to approximate the value function as it depends on current price levels and current reservoir levels. When medium-term hydropower planning is used to support the decision-making for short-term planning, only the value function is passed to the short-term planning problem, not necessarily a deterministic policy. Thus, we do not require the algorithm to generate deterministic policies, and rather allow them to be stochastic while taking care not to relax information constraints. This ensures that the algorithm yields estimated lower bounds on the objective function value, as our method relies on possible scenarios for the future development of prices and inflow. The new approach, called STRO, is compared to SDDP and RI.

Although STRO can be applied to broader classes of problems compared to those solved by SDDP, we designed the experiments to suit SDDP well. We leave for future work to explore the performance on non-convex problems.

Sources of bias potentially impacting the comparisons of the implemented models include both discretization and sampling error subjected to the different methods. Efforts were made to mitigate this by performing multiple trial runs, averaging the results, and performing experiments that, for instance, aimed to bypass the discretization error of SDDP. We have also used seeding to ensure that the exact sampling error is imposed on the rolling horizon methods.

Another factor to consider is that the methods discussed and compared were only tested on a single hydropower system, using only one type of inflow and price model. Specific properties of the case system can make the method more suitable, which can mean that similar results will not be achievable for other systems. These concerns were not explored in the interest of scope, but they should be further examined in the future.

With these points made, we emphasize that the RI algorithm has been shown to produce promising results when applied to the reservoir management problem by other authors with different stochastic models for the exogenous processes. This strengthens our confidence in the results.

The idea of using a sampling-based heuristic for rolling horizon methods can also be explored further. The STRO method only considers building two-stage stochastic

convex programs, but in principle, multi-stage non-convex stochastic programs could be built and solved instead.

A. Piece-wise linear approximations

As the head $h_{i,t}$ is a function of the amount of water in a reservoir $l_{i,t}$ this function needs to be approximated in a piece-wise manner. To do this, the relation between head and water level for a specific reservoir i can be represented by a set M_i which contains sample points $((l_{i,1}, h_{i,1}), \dots, (l_{i,K}, h_{i,K}))$, representing the mapping between water level and head. Increasing the number of points naturally leads to a better approximation, but comes at the consequence of increased complexity.

To find the head from the water level using the piece-wise linear graph, a set of variables representing the weight of point k for reservoir i at time t denoted $\lambda_{i,t,k}$. By ensuring that no weights can be negative (9), the weights sum to 1 (10), and that only two adjacent points are non-zero using a special ordered set of type 2 constraint (SOS2) (11), one can find the head $h_{i,t}$ based on the water level $l_{i,t}$. This is done using Eqs. (12)–(13). Equation (12) forces the weights to correspond with the water level, and Eq. (13) finds the head given the weights. If the function describing the mapping from water level to head is concave, the SOS2 constraint (11) is redundant.

$$\lambda_{i,t,k} \geq 0, \quad k \in M_i \quad i \in R^P \quad t = 0, \dots, T - 1 \tag{9}$$

$$\sum_{k \in M_i} \lambda_{i,t,k} = 1, \quad i \in R^P \quad t = 0, \dots, T - 1 \tag{10}$$

$$\lambda_{i,t,k} \in \text{SOS2}, \quad k \in M_i \quad i \in R^P \quad t = 0, \dots, T - 1 \tag{11}$$

$$\sum_{k \in M_i} \lambda_{i,t,k} l_{i,k} = l_{i,t}^{avg}, \quad i \in R^P \quad t = 0, \dots, T - 1 \tag{12}$$

$$\sum_{k \in M_i} \lambda_{i,t,k} h_{i,k} = h_{i,t}, \quad i \in R^P \quad t = 0, \dots, T - 1 \tag{13}$$

As the incoming and outgoing water levels can be quite different, the average water level $l_{i,t}^{avg}$ is used as the basis for the head for the period. This variable is given according to constraint (14).

$$l_{i,t}^{avg} = \frac{l_{i,t} + l_{i,t+1}}{2}, \quad i \in R^P \quad t = 0, \dots, T - 1 \tag{14}$$

B. McCormick envelope constraints

In the following the constraints comprising the McCormick envelopes are shown. These are used to ensure that the substitution variable $w_{i,t}$ closely approximates the bi-linear term $h_{i,t}x_{i,t}$. These constraints depend on the variables $h_{i,t}$ and $x_{i,t}$, as well as the maximum and minimum values that these can take on. These maximum and minimum values are denoted $H_{i,t}^{max}, H_{i,t}^{min}, X_{i,t}^{max}$ and $X_{i,t}^{min}$.

$$w_{i,t} \geq X_{i,t}^{min} h_{i,t} + x_{i,t} H_{i,t}^{min} - X_{i,t}^{min} H_{i,t}^{min}, \quad i \in R^P \quad t = 0, \dots, T-1 \quad (15)$$

$$w_{i,t} \geq X_{i,t}^{max} h_{i,t} + x_{i,t} H_{i,t}^{max} - X_{i,t}^{max} H_{i,t}^{max}, \quad i \in R^P \quad t = 0, \dots, T-1 \quad (16)$$

$$w_{i,t} \leq X_{i,t}^{max} h_{i,t} + x_{i,t} H_{i,t}^{min} - X_{i,t}^{max} H_{i,t}^{min}, \quad i \in R^P \quad t = 0, \dots, T-1 \quad (17)$$

$$w_{i,t} \leq X_{i,t}^{min} h_{i,t} + x_{i,t} H_{i,t}^{max} - X_{i,t}^{min} H_{i,t}^{max}, \quad i \in R^P \quad t = 0, \dots, T-1 \quad (18)$$

$$0 \leq w_{i,t} \leq X_{i,t}^{max} H_{i,t}^{max}, \quad i \in R^P \quad t = 0, \dots, T-1 \quad (19)$$

Acknowledgements We gratefully acknowledge support from the research center HydroCen, RCN No. 257588. We are thankful to Montel for supplying price data, and to Eviny for supplying data regarding their hydropower plants.

Funding Open access funding provided by NTNU Norwegian University of Science and Technology (incl St. Olavs Hospital - Trondheim University Hospital).

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

1. Ávila, D., Papavasiliou, A., Löhdorf, N.: Parallel and distributed computing for stochastic dual dynamic programming. *Comput. Manage. Sci.* **19**, 199–226 (2022). <https://doi.org/10.1007/s10287-021-00411-x>
2. Benth, F.E., Koekkebakker, S., Ollmar, F.: Extracting and applying smooth forward curves from average-based commodity contracts with seasonal variation. *J. Deriv.* **15**(1), 52–66 (2007)
3. Bertsekas, D.: *Dynamic Programming and Optimal Control*, vol. I. Athena Scientific, Belmont (2005)
4. Braaten, S.V., Gjønnnes, O., Hjertvik, K.J., Fleten, S.E.: Linear decision rules for seasonal hydro-power planning: modelling considerations. *Energy Procedia* **87**, 28–35 (2016)

5. Breslin, J., Clewlow, L., Elbert, T., Kwok, C., Strickland, C., van der Zee, D.: Gas storage: rolling intrinsic valuation. *Energy Risk* **January**, 61–65 (2009)
6. Cerisola, S., Latorre, J.M., Ramos, A.: Stochastic dual dynamic programming applied to nonconvex hydrothermal models. *Eur. J. Oper. Res.* **218**(3), 687–697 (2012)
7. da Silva, E., Finardi, E.: Parallel processing applied to the planning of hydrothermal systems. *IEEE Trans. Parallel Distrib. Syst.* **14**(8), 721–729 (2003). <https://doi.org/10.1109/TPDS.2003.1225052>
8. da Costa, J.P., de Oliveira, G.C., Legey, L.V.L.: Reduced scenario tree generation for mid-term hydrothermal operation planning. In: 2006 International Conference on Probabilistic Methods Applied to Power Systems (2006)
9. Dietze, M., Chávarry, I., Freire, A.C., Valladão, D., Street, A., Fleten, S.E.: A novel semiparametric structural model for electricity forward curves. *IEEE Trans. Power Syst.* **38**(4), 3268–3278 (2023). <https://doi.org/10.1109/TPWRS.2022.3197982>
10. Ding, L., Ahmed, S., Shapiro, A.: A Python package for multi-stage stochastic programming. *Optimiz. Online* (2019)
11. Downward, A., Dowson, O., Baucke, R.: Stochastic dual dynamic programming with stagewise-dependent objective uncertainty. *Oper. Res. Lett.* **48**(1), 33–39 (2020)
12. Gjelsvik, A., Mo, B., Haugstad, A.: 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*, pp. 33–55. Springer, Berlin (2010)
13. Goodwin, D.: Schwartz-Smith 2-factor model: parameter estimation (2020). <https://www.mathworks.com/matlabcentral/fileexchange/43352-schwartz-smith-2-factor-model-parameter-estimation>. Accessed 9 Nov 2021
14. Helseth, A., Braaten, H.: Efficient parallelization of the stochastic dual dynamic programming algorithm applied to hydropower scheduling. *Energies* **8**, 14287–14297 (2015). <https://doi.org/10.3390/en81212431>
15. Helseth, A., Mo, B., Warland, G.: Long-term scheduling of hydro-thermal power systems using scenario fans. *Energy Syst.* **1**(4), 377–391 (2010)
16. Helseth, A., Mo, B., Lote Henden, A., Warland, G.: Detailed long-term hydro-thermal scheduling for expansion planning in the Nordic power system. *IET Gen. Transmiss. Distrib.* **12**(2), 441–447 (2018)
17. Hjelmeland, M.N., Zou, J., Helseth, A., Ahmed, S.: Nonconvex medium-term hydropower scheduling by stochastic dual dynamic integer programming. *IEEE Trans. Sustain. Energy* **10**(1), 481–490 (2018)
18. Infanger, G., Morton, D.P.: Cut sharing for multistage stochastic linear programs with interstage dependency. *Math. Program.* **75**(2), 241–256 (1996)
19. Lai, G., Margot, F., Secomandi, N.: An approximate dynamic programming approach to benchmark practice-based heuristics for natural gas storage valuation. *Oper. Res.* **58**(3), 564–582 (2010)
20. Löhdorf, N., Shapiro, A.: Modeling time-dependent randomness in stochastic dual dynamic programming. *Eur. J. Oper. Res.* **273**, 650–661 (2019)
21. Löhdorf, N., Wozabal, D.: The value of coordination in multimarket bidding of grid energy storage. *Oper. Res.* (2022). <https://doi.org/10.1287/opre.2021.2247>
22. Machado, F.D., Diniz, A.L., Borges, C.L., Brandão, L.C.: Asynchronous parallel stochastic dual dynamic programming applied to hydrothermal generation planning. *Electric Power Syst. Res.* **191**, 106907 (2021)
23. Pereira, M., Pinto, L.: Multi-stage stochastic optimization applied to energy planning. *Math. Program.* **52**(1–3), 359–375 (1991)
24. Philpott, A., Guan, Z.: On the convergence of stochastic dual dynamic programming and related methods. *Oper. Res. Lett.* **36**(4), 450–455 (2008)
25. Pinto, R.J., Borges, C.T., Maceira, M.E.P.: An efficient parallel algorithm for large scale hydrothermal system operation planning. *IEEE Trans. Power Syst.* **28**(4), 4888–4896 (2013). <https://doi.org/10.1109/TPWRS.2012.2236654>
26. Powell, W.: *Approximate Dynamic Programming: Solving the Curses of Dimensionality*. Wiley, Oxford (2007)
27. Roald, L.A., Pozo, D., Papavasiliou, A., Molzahn, D.K., Kazempour, J., Conejo, A.: Power systems optimization under uncertainty: a review of methods and applications. *Electr. Power Syst. Res.* **214**, 108725 (2023)
28. Schwartz, E., Smith, J.E.: Short-term variations and long-term dynamics in commodity prices. *Manage. Sci.* **46**(7), 893–911 (2000)

29. Séguin, S., Audet, C., Côté, P.: Scenario-tree modeling for stochastic short-term hydropower operations planning. *J. Water Resour. Plan. Manag.* **143**(12), 04017073 (2017)
30. Shapiro, A., Nemirovski, A.: On complexity of stochastic programming problems. In: Jeyakumar, V., Rubinov, A. (eds.) *Continuous Optimization. Applied Optimization*, vol. 99, pp. 111–146. Springer, Boston (2005)
31. Shapiro, A., Tekaya, W., da Costa, J.P., Soares, M.P.: Risk neutral and risk averse stochastic dual dynamic programming method. *Eur. J. Oper. Res.* **13**, 375–391 (2013)
32. Valipour, M., Banihabib, M.E., da Costa, J.P., Behbahani, S.M.R.: Comparison of the ARMA, ARIMA, and the autoregressive artificial neural network models in forecasting the monthly inflow of Dez dam reservoir. *J. Hydrol.* **13**, 433–441 (2012)
33. Van Slyke, R.M., Wets, R.: L-shaped linear programs with applications to optimal control and stochastic programming. *SIAM J. Appl. Math.* **17**(4), 638–663 (1969)
34. Wallace, S.W., Fleten, S.E.: Stochastic programming models in energy. *Handb. Oper. Res. Manag. Sci.* **10**, 637–677 (2003)
35. Wu, O., Wang, D., Qin, Z.: Seasonal energy storage operations with limited flexibility: the price-adjusted rolling intrinsic policy. *Manuf. Serv. Oper. Manag.* **14**(3), 455–471 (2012)

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.