NTNU – Trondheim
Norwegian University of
Science and Technology

Seasonal hydropower planning using linear decision rules

Simen V. Braaten
Ola Gjønnes
Knut Jonny Hjertvik

Industrial Economics and Technology Management
Submission date: June 2015
Supervisor: Stein-Erik Fleten, IØT

Norwegian University of Science and Technology
Department of Industrial Economics and Technology Management

Problem formulation

To implement an optimization model for seasonal planning under uncertainty in suitable software, by using linear decision rules. Additionally, the applicability of linear decision rules in this context will be analyzed.

Abstract

Hydroelectric reservoirs need to be managed carefully in order to maximize their benefits. Existing stochastic optimization approaches for making seasonal reservoir schedules often suffer from very high computing efforts, that grow exponentially with the problem size. This thesis investigates a relatively new approach to solve this problem, known as a linear decision rules (LDR) approximation. Here, the decision variables are replaced with affine functions of the realizations of the uncertain parameters. The presented LDR optimization model takes uncertainty in price and multiple inflow series into account. Important considerations in hydropower scheduling, such as head corrections and finding water values, are also incorporated. The model is tested to find the applicability of LDR in the seasonal hydropower planning problem, and theoretical upper bounds are obtained to evaluate the quality of the solutions. The main result is that accuracy is traded expensively for tractability, as the solutions from the LDR-formulation are far from the upper bound. The solution values are also clearly beaten by a rolling intrinsic heuristic already in use in similar industries. However, reasonable policies and water values are generated, and head corrections are satisfactorily included. Additionally, by reducing the number of realizations the decision rules are dependent on, very short computational times are achieved. Therefore, LDR are only a good fit if short running times are important.

Sammendrag

Magasin brukt til vannkraftproduksjon må styres med omhu for å levere sin maksimale nytte. Eksisterende stokastiske optimeringsmetoder for å sesongplanlegge produksjon krever veldig lang kjøretid, og den vokser eksponensielt med problemstørrelsen. I denne oppgaven undersøkes en relativt ny måte å løse dette problemet på, kjent som tilnærming ved lineære beslutningsregler (LDR). Her blir beslutningsvariablene erstattet av affine funksjoner, som avhenger av realiseringen av de usikre parameterne elektrisitetspris og tilsig. Vannkraftspesifikke konsepter som fallhøydeeffekter og vannverdier blir også hensyntatt. Modellen testes, med formål om å finne ut hvor anvendbare LDR er til å løse sesongplanleggingsproblemet innen vannkraft. Til dette utvikles teoretiske øvre grenser for objektivverdien, og det viser seg at korte kjøretider koster dyrt i form av optimalitetstap. Løsningsverdiene er langt fra de øvre grensene, og blir også slått av en heuristikk med rullende horisont, som allerede er i bruk i bransjen. Beslutningsreglene fra modellen er imidlertid fornuftige, og det er enkelt å inkludere fallhøydeeffektene. I tillegg kan reduksjoner i antall parametere hver beslutningsregel avhenger av føre til veldig korte kjøretider. Konklusjonen er at LDR bare er en god løsning på dette problemet dersom kort løsningstid er viktig.

Preface

This master thesis is written within Applied Economics and Operations Management at the Department of Industrial Economics and Technology Management, Norwegian University of Science and Technology (NTNU).

Most of all, we would like to thank our supervisor Professor Stein-Erik Fleten for inspiration and invaluable guidance, and our co-supervisor, Powel employee Gro Klæboe, PhD., for swift and helpful answers to our countless questions. We also owe a thanks to Postdoc Marco Zugno at DTU, who patiently answered our inquiries, and Powel AS for access to their benchmark watercourse Fantasi. Additionally, Daniel Risberg at KTH has provided valuable input on the stochastic data, and without the help from Stefan and Uffe at Powel, our model would hardly exist. Last, we want to thank our fellow student Håkon Sveinssønn Mork, whose input on both content and typesetting has been highly appreciated.

Trondheim, June 11th, 2015.

Simen V. Braaten | Ola Gjønnes | Knut Hjertvik

Contents

1 Introduction

For at least 2000 years, people have harnessed falling and running water for useful purposes [32], and in the late 19th century, hydropower became a source for generating electricity. In the 20th century, technology matured, networks for transmitting power were developed, and markets for trading were established.

Today, hydropower producers make short-term bidding decisions every day to maximize profit. These decisions are, among other factors, based on estimates of how much the water in their reservoirs will be worth at the end of the short term planning period. To generate these estimates, a seasonal planning optimization model that accounts for effects further into the future is used. Additionally, a seasonal planning model outputs schedules that can be used in risk management [44], and within budgeting, a seasonal plan is used to forecast the revenue stream for the year ahead.

The seasonal planning model can be implemented as either a deterministic or a stochastic model. A deterministic model uses future inflows and electricity prices as if they are fixed, known values.

However, due to the large inherent uncertainty in inflows, it is essential to take stochasticity into account. This reduces the risk of spilling water or emptying the reservoir. Considering the price as stochastic is preferred for all deregulated markets. Especially so in the Nordic market, where the high share of hydropower leads to electricity prices that are strongly influenced by the hydrological situation [56].

Several methods can be used when taking stochasticity into account. Historically, the most promising ones have been based on either dynamic programming (DP) or nested Benders' decomposition. One alternative is deterministically solving for the entire planning horizon, and locking the first stage variables. Then updating with the actual realizations of inflow and price, the remaining stages can be iteratively solved with a receding horizon. This is known as a rolling instrisic (RI) heuristic [43].

A relatively new way of looking at the seasonal hydropower problem is called linear decision rules (LDR) approximation. This solution method treats the decision variables as affine functions of the uncertain parameters. The result is an interesting trade-off, where tractability is achieved at the expense of loss of optimality due to the linearization.

After Ben-Tal resurrected LDR in [4] a decade ago, several papers have been published describing this new method applied to hydropower planning. In [37], the method of iterative LDR is introduced as the approach where an LDR formulation is solved iteratively, updating problem parameters, while in [22, 41] LDR approximations are applied on both the primal and dual of the seasonal hydropower problem, thereby also presenting values for the optimality gap caused by the approximation.

It is in the wake of these publications our thesis align. More specifically, the research question adressed in this thesis is:

How suitable are *linear decision rules* for seasonal hydropower scheduling?

Answering this breaks down to a series of subquestions:

1. Can hydropower specific implementation details, such as head effects and water values, be included in the solution process?
2. Will LDR be able to give quicker results or higher objective values than RI?
3. What bounds can be found on the suboptimality of the LDR approximation?

Our main contribution is a thorough computational test of LDR in seasonal hydropower planning. It is shown that the method can not be conclusively endorsed when compared with RI, and with a new, state of the art, theoretical upper bound. Additionally, we present LDR with explicit focus on several uncertainty types, in an introductory guide to the topic. To the research on LDR, we contribute with novel combinations of elements from earlier works. This includes normalized, zero mean uncertainties [63] and iterative recalculations of energy coefficients. Our work in general, and these contributions in particular, are presented in two papers, "Linear

Decision Rules for seasonal hydropower planning: Modelling Considerations" and "Linear Decision Rules for seasonal hydropower planning: Implementation and computational performance". Both articles are planned submitted to international academic journals.

This thesis is split in three main parts. Chapters 2-5 serve as background and additional notes to the articles. Here, theory on the hydropower industry, alternative solution methods, LDR, and theoretical upper bounds is presented. The purpose is to provide some context to readers who are not already familiar with the topics discussed in the articles.

In Article 1 [10], an LDR model is derived from a deterministic seasonal hydropower scheduling model. The intention is to introduce the topic for practitioners of hydropower planning. In the scope of this thesis, the paper functions as an explanation of the transformations done to obtain an LDR model for the seasonal hydropower scheduling problem.

Article 2 [9] presents further analysis on the model. Here, the LDR model is benchmarked against an implementation of RI, and bounds for the optimality gap are presented. Additionally, implementation details regarding the important hydropower specific details are outlined and explained, before the findings are discussed.

Last, in Section 6, our conclusions from Article 1 and Article 2 are summarized, and the reasearch questions are answered.

It should be noted that several stretches in this thesis are similar, as the articles are written to be able to be read independently. This means that some theory and introductory material needs to be presented multiple times.

2 Hydropower scheduling

This chapter serves as a brief introduction to the hydropower industry. Key figures are given, and main characteristics of hydropower scheduling are presented.

2.1 Overview

By building large, engineered dams, humans have added over half a million artificial reservoirs to the world's surface [15]. These reservoirs serve many purposes such as flood control, drinking water supply, and low flow augmentation [48]. A fair share is also used for electricity generation. Globally, more than 50 000 hydropower installations have been built, although most of them are rather small. Worldwide, there are around 1000 power plants with a capacity of more than 100 MW, and a comprehensive list can be found in [19]. The utilization of hydropower resources is not evenly divided between countries, for instance one third of the power stations larger than 2000 MW are situated in China. An overview of the relative shares of hydropower generation by country is given in Figure 1.

Since 1990, global hydropower generation has increased by about 50 %, to 3288 TW h (2008) [28], see Figure 2. Hydropower is the leading renewable energy source, with new capacity additions since 2005 generating more electricity than all other renewables combined. This growth is expected to continue as demand for clean, economically viable energy increases [27], also increasing the demand for decision support in the industry.

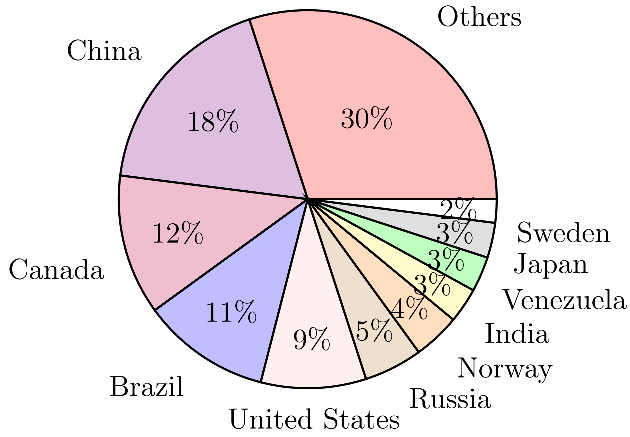

Figure 1: Shares in hydropower generation 2008 [28]

.

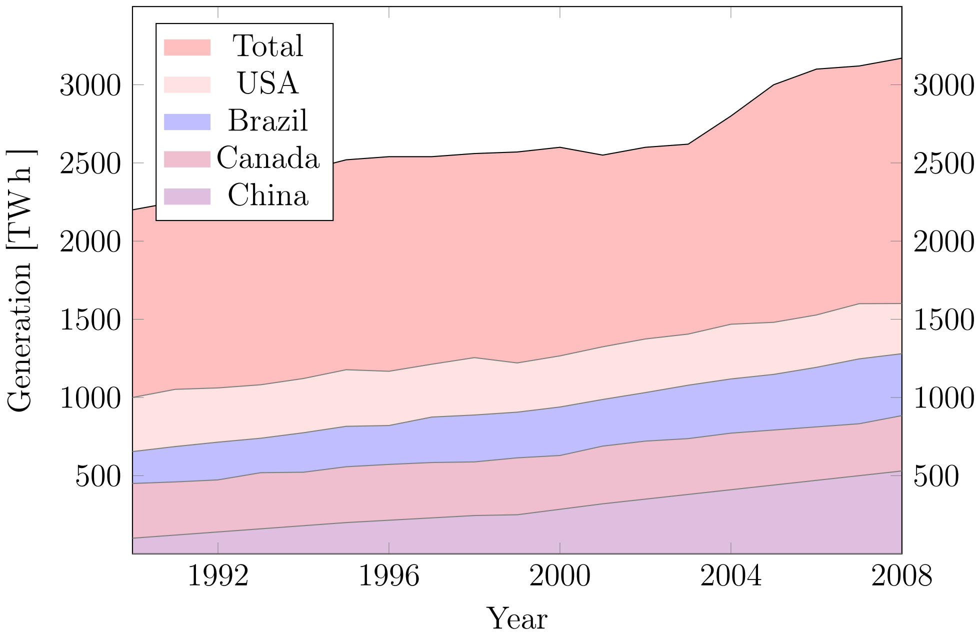

Figure 2: Development of global hydropower generation, 1990-2008, [28].

Although the global generation of hydropower grows each year due to the building of new capacity, there can be significant variations in a single water system's annual power output. This can also be observed on a national level. For instance, the annual hydropower generation in Norway has had a standard deviation of 11.7 TW h from a mean of 126.8 TW h since 2000 [50]. New types of reservoir management may help reduce the deviation between years, if the risk of spill and emptying is handled and mitigated in a better way.

2.2 Hydropower planning problem

A typical installation is a system with three parts: an electric plant where the electricity is generated, a dam that can be opened or closed to control water flow, and a reservoir where water can be stored. The water behind the dam flows through an intake and pushes against blades in a turbine, causing them to turn. The turbine spins a generator to create electricity. Then water exits through a draft tube to a lower-elevated reservoir.

This report considers the problem a price-taking hydroelectric power producer faces when deciding a seasonal production plan in a deregulated market setting. The aim is, given a forecast of future spot prices and inflow, to decide the generation strategy in order to maximize profits. More specifically, the strategy consists of deciding how much water to save, discharge and pump in every reservoir for all time periods in the planning horizon, all relevant constraints taken into account.

2.3 Scheduling hierarchy

Seasonal planning is just one of many challenges for a hydropower producer. Examples include investment decisions, reservoir management, risk management, and short term bidding decisions. The ideal solution to face all of these issues, would be one single optimization process that took all future effects of a decision into account [14]. However, due to the complexity and size of such a model, hydropower decision making is usually split into several parts with different degree of detail and different scheduling horizon; long term planning, and seasonal and short term scheduling [57]. Long term planning models focus on investments and typically have a 15–20 year horizon, while seasonal term planning deals with reservoir management with a horizon of up to five years [51]. Short term planning typically deals with decisions with horizons of one week or shorter, such as making generation plans to meet obligations, known as the *unit commitment problem* [57].

To utilize results across different scheduling horizons, it is necessary to couple the long-, medium-, and short term models. Several principles can be used, for example coupling through price or volume, or by penalty functions [14]. Volume coupling means that the longer term model specifies a target reservoir volume, and the shorter term model schedules optimally to achieve this. Penalty functions can be included to allow some deviation from the specified volume. Price coupling is most common due to its flexibility, given that it is possible to obtain a correct valuation of the water stored in a reservoir. Using this principle, the longer term model supplies a valuation function for keeping water in the reservoirs, leaving it to the shorter term model to decide whether it is more valuable to discharge it now. Commercial seasonal models such as Vansimtap and Prodrisk use this type of coupling [17], by outputting *water values* to the short term optimization.

2.4 Water Values

The water value is defined as the expected future marginal value of stored water in a reservoir [58]. The water value will increase with decreasing reservoir level, and vice versa. This is intuitive, as extra water is increasingly important when the water gets more scarce, reducing the flexibility of the planner.

2.5 Energy coefficient

The energy coefficient measures the amount of energy generated in a power plant per unit water. In hydropower generation it is defined as:

$$E(h) = \frac{1}{3.6 \cdot 10^6} \eta \rho g h \qquad [kWh/m^3].$$

Here, η is the power plant's efficiency, ρ is the water density $[\frac{kg}{m^3}]$, and g is the gravitational constant $[\frac{m}{s^2}]$. Importantly, h is the fall height of the water, known as the plant's head $[m]$ (see Section 2.6).

The efficiency of a plant's turbine and generator is typically dependent on the throughput of water. At optimal conditions, a plant can usually make use of over 90% of the potential energy in the water [54]

In a typical power plant, the energy coefficent is considered to have a non-negative value. For pumping stations, however, energy is consumed in order to move water to another reservoir, hence the pumping coefficients at these stations are strictly negative. For pumped-storage facilities, logically, the pumping coefficient is of greater magnitude than the corresponding energy coefficient.

2.6 Head Effects

Head, h_r, is considered to be the difference in elevation between the preceding reservoir's level and where the water ends up after passing through the turbine. This is either the succeeding reservoir level, or where r's power station releases the water to a river or the ocean. The latter is called r's tailwater level. Mathematically, the head level is given by

$$h_r = z_r - max\{z_{S_r}, z_r^{TW}\}$$

Here z_r and z_{S_r} are the reservoir levels in meters above sea level, while z_r^{TW} is the tailwater level.

This means that deciding between discharging or keeping water affects the head levels that can be utilized in power generation. Taking head effects into account might alter the optimal plan, since it gives the model an incentive to keep the reservoir level high. A higher head increases the energy coefficient, and makes it possible to generate more energy from the same amount of water.

Consider the simplified situation first presented in [52]. Here, the reservoir is initially full and should be emptied within a given period, during which inflow is known and low relative to the generation limit, and electricity prices are constant[1]. The optimal strategy will then be to only discharge the inflow, and thus keep the energy coefficient as high as possible, until the point where the producer has to drain more water to empty the reservoir in time. Looking at Figure 3b, we can conclude that the production plan to the right will generate more power than the left one, because the average energy coefficient is higher. A model that does not consider head effects, however, will treat the two plans as equivalent, and potentially deliver suboptimal results.

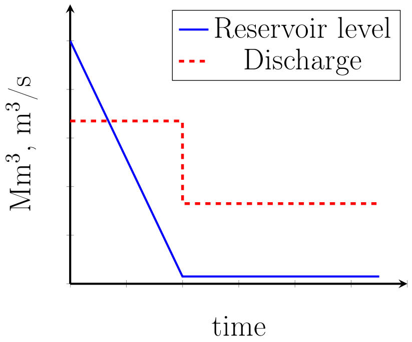

(a) Generation plan and reservoir trajectories without head effects

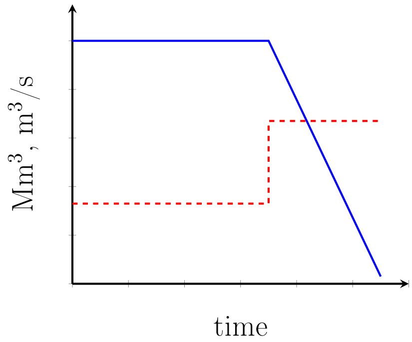

(b) Generation plan and reservoir trajectories with head effects

Figure 3: How including head effects can alter optimal production, as seen in [52]

2.7 Degree of flexibility

The term flexibility in this context refers to the ability to allocate water resources to certain time periods in order to maximize profit [13]. The producer's flexibility in choosing the optimal strategy depends on the generation capacity of the power stations, the size of the reservoirs, the amount and distribution of inflow, and judicial regulations on reservoir level and waterway flow.

Hydropower installations are normally categorized within two different types; "run of river" and "reservoir" [42]. A "Run of river"-installation has a small or nonexis-

[1] A scenario that is not necessarily very different from a regular winter season for many reservoirs.

tent forebay, and its operation is solely determined by the rate of instream flow, hence discharge for generation cannot be adjusted to changes in price and inflow [29]. It is characterized by relatively high flow volumes and low degree of control on power output.

On the other end we have reservoir-installations which can save water for later use. Pumped-storage hydroelectricity can further increase flexibility by moving water from a reservoir at lower elevation, to a reservoir at high elevation during low-cost, off-peak hours. Restrictions on maximum and minimum generation and bypass, as well as time varying reservoir limits periodic maintenance, will on the other hand reduce the flexibility of a hydropower plant.

Various factors can be used to measure an installation's flexibility, among them the *degree of regulation* (DoR).

DoR is defined as the reservoir volume divided by the average annual inflow. It describes the number of years it takes to fill the reservoir given average annual inflow and no generation, and can be used as a tool when determining the schedule horizon. High DoR means the reservoir takes a long time to fill up, and a longer scheduling horizon should be considered.

The largest reservoirs have DoRs in the range 2-3 [32]. Seasonal variations in inflow will then have limited effect on the discharge decision, making a seasonal model less useful. The same conclusion can be drawn for run-of-rivers installations. Since water cannot be stored, there are no decisions to be made. Between these two extremes, hydropower installations face both uncertainty and flexibility, so reservoir management is important.

3 Competing methods

Operations research has been used in hydropower planning for half a century [49]. During this time, many solution methods have been developed, some of them very successful. In this chapter we present these alternatives to LDR, to give the reader a sense of each method's benefits and drawbacks. The method we test our computational results against, RI, is presented in detail in Article 2.

3.1 Stochastic dynamic programming

Dynamic programming (DP) is a solution method where the initial problem is broken down into simpler subproblems [2]. Each of them is solved, before they are combined to reach a final solution. In the context of hydropower generation planning, an optimal seasonal plan can, at least in theory, be constructed by solving weekly scheduling subproblems. This is done by iterating over possible states in each stage. In this context, a state should include information about reservoir levels, prices and inflow. In each state, a number of actions are possible, each associated with a payoff in the current stage, and a transition to a new state in the next stage. A typical example would be a weekly production plan. The plan includes discharging water for generation of energy, which has an immediate payoff. At the same time, this discharge of water changes the reservoir levels, and thus defines a transition to a different state.

Then, the objective is to find the action that maximizes the sum of earnings in this stage, and future earnings in the resulting state. By starting at the last stage, working backwards, and recording the value of each state in each stage, a recursive solution can be found. To find the best action in a given state and stage, the future value of each potential resulting state is simply looked up in a table, as the next stage is already solved. Then this value is added to the payoff of the action in this stage. From this sum, the best action is found, and the best result is stored, so the next iteration can look it up.

Stochastic dynamic programming (SDP) differs from traditional DP in that the

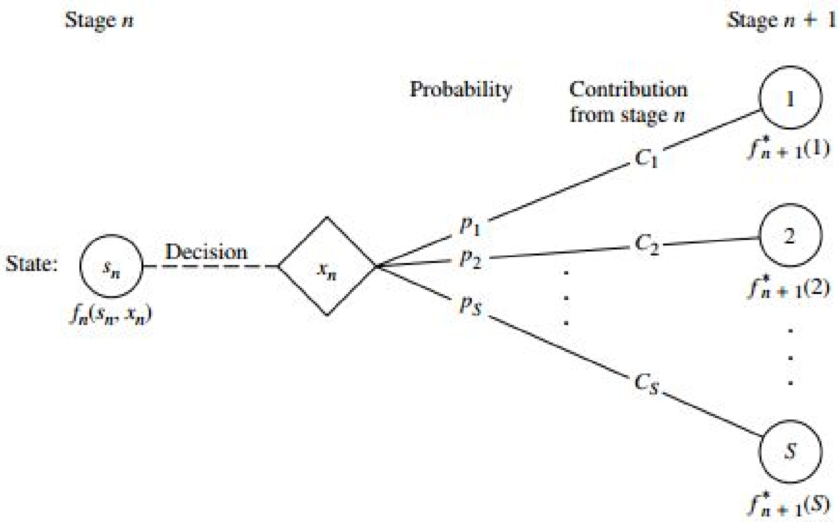

Figure 4: Basic structure of stochastic DP, figure collected from [26]

outcome of each stage is not completely determined by the state and action of the current stage. Rather, the contribution C_i and resulting stage i' is also dependent on a known probability distribution [62], as seen in Figure 4. For example, the weekly plan discussed above does not deterministically lead to a certain state in the next stage. To tell the reservoir levels with certainty, the realization of the uncertain inflow is also necessary.

Formally, let s_n be the current state where decision x_n is made. The probability of the system then going to state i is p_i, and the contribution to the objective in that case is C_i. The goal of each subproblem is to make decisions that maximize the expected sum of contributions from stage n and onwards, yielding the profit to go function [26]:

$$f_n(s_n, x_n) = \sum_{i=1}^{S} p_i[C_i + f^*_{n+1}(i)],$$

where S is the number of possible states, and with

$$f^*_{n+1}(i) = \max_{x_{n+1}} f_{n+1}(i, x_{n+1}),$$

for all feasible values of the next decision, x_{n+1}. The main drawback of SDP is the computational intractability that occurs when the number of dimensions grows. As one might imagine, the number of possible states grows rapidly in this case, often called the *curse of dimensionality.* An SDP that checks every possible state, quickly becomes too complex to solve. In resolving this issue, a series of good techniques have been discovered, such as state aggregation [1], and approximate dynamic programming [31]. These make SDP a feasible alternative, that still thrives in the field of reservoir management [53]. Additionally, this research has resulted in an alternative approach, *stochastic dual dynamic programming* (SDDP).

3.2 Stochastic dual dynamic programming

The main contribution of SDDP is to remove the need to explore all possible states [38]. Instead of looking up solutions in every possible next stage, the algorithm collects dual variables from the solutions at some possible states in the next stage. Which states to solve is decided by sampling scenarios [29], and each set of dual variables contribute to an approximation function of the future value. This way, the number of states to solve in each stage can be drastically reduced, effectively counteracting the curse of dimensionality. For each state i_k that is solved in stage $n+1$, the dual variables are used to generate a cut $\kappa_k(i)$ (see Figure 5) in a piecewise linear approximation $\hat{f}^*_{n+1}(i)$ of the future value. With this approximation, there is no need to discretize the state space, which results in reduced computational complexity.

Since the total future earnings as a function of reservoir levels can be closely approximated by a piecewise linear function, SDDP has delivered good results and become very popular within reservoir management [39]. Here, the cuts have a fairly natural interpretation. In each solved state i_k, let α_{rk} be the dual variable associated with the reservoir balance equation for reservoir r. Then, α_{rk} can be seen as the additional value of one unit of water in r. Now, if R_k is the total future value at state i_k, and $\hat{m}_{rk}$ the optimal reservoir levels, total future value, as a function of state i, can be approximated by the cut (remember that the state i

contains the reservoir levels m_r):

$$\hat{f}^*_{n+1}(i) \leq \kappa_k(i) = R_k + \sum_r \alpha_{rk}(m_r - \hat{m}_{rk}), \quad \forall i_k.$$

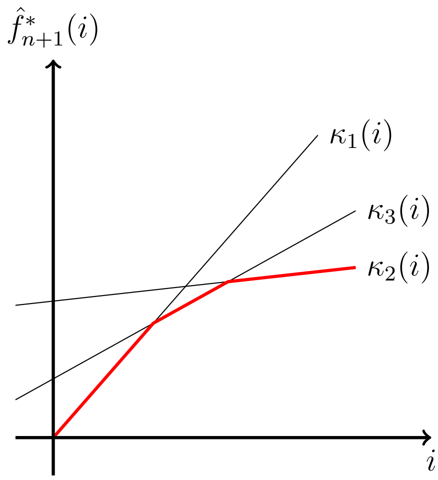

Figure 5: In SDDP, the future value is approximated by cuts calculated at sampled states.

These types of cuts are very commonly in use when describing the water value we discussed earlier.

The main critique against SDDP, and the reason our comparison is against RI, is that SDDP struggles with jointly stochastic price and inflow. This is because letting both of these parameters be stochastic may destroy the convexity (or concavity) of the objective functions, making the cutting plane approach infeasible. Some authors have circumvented this problem by assuming price and inflow to be independent parameters. An example with discrete price levels in a combined SDP/SDDP model is presented in [18], but a significant layer of complexity is still added.

3.3 Scenario-based methods

Even though the SDP and SDDP formulations have been the most successful methods for solving the hydropower scheduling commercially, other approaches have also been studied in the literature [8, 29]. The most important of these approaches is perhaps traditional multistage stochastic programming based on scenario generation [33]. Here, given a set of scenarios with associated probabilities, a deterministic equivalent model can be built [45]. In this, each scenario has its own *recourse variables*, representing decisions that would be made, given that the scenario was realized [25]. For each scenario, all constraints must be restated in terms of the associated recourse variables, to ensure they hold no matter the realized scenario [46]. Then, each scenario will give rise to a certain objective value, and the overall objective function can be defined as the average of these. Variations exist, where the objective function is the worst case [3] earnings, or the constraints only need to hold with a certain probability [5], but the core idea remains.

The primary advantage of scenario-based stochastic programming over the other mentioned approaches, is the flexibility it offers in modelling and defining scenarios, particularly if the state dimension is high [60]. Further, it is natural to base stochastic programming on linear or mixed integer programming, which means that many well-studied and standard techniques are available, like Lagrange relaxation [36], Benders decomposition [47], and Dantzig-Wolfe reformulation [35]. Still, the formulation often results in very large models, requiring special solution algorithms.

4 Linear decision rules

Linear decision rules are the main focus of this thesis. They provide a tractable approximation of a stochastic optimization problem [4]. Contrasted to scenario based models, LDR do not require the distributions of the uncertain parameters or state discretization. However, it reduces flexibility, and this may result in a larger optimality gap. This chapter sets the stage for the transformation in Article 1 by describing robust optimization and briefly introducing LDR as affine reactions to uncertainty.

4.1 Robust Optimization

In robust optimization the decision-maker is seeking to construct a solution that is feasible for any realization of the uncertainty. It is a field of growing interest and importance, and has mainly been developed in the last 15 years [20].

Given that the objective or one or more restrictions depend on a random vector δ drawn from an set $\mathcal{U}$ of possible values, a robust problem can be expressed as a stochastic program [4]

$$\left\{ \min_{\mathbf{x}\in\mathbb{R}} \mathbf{c}(\delta)^\top \mathbf{x} \; s.t. \; A(\delta)\mathbf{x} \geq \mathbf{b}(\delta) \right\}_{\delta\in\mathcal{U}}. \tag{$\mathcal{SP}$}$$

Here, $\mathcal{U}$ can be thought of as the set of possible realizations of δ, and a certain realization of δ as a scenario. In the *decision rule* approach, the decision variables are functions $\mathbf{x}(\delta)$ of the uncertainty, often called decision rules or *reactions.* Sets of such reaction functions are called *policies*, as they allow a coherent mapping from any realization of uncertainty, to appropriate decisions. Now, the problem is transformed to a search through the space $\mathbb{F}$ of possible functions, in order to find the set of decision rules that perform best under the uncertainty of the problem:

$$\left\{ \min_{\mathbf{x}(\cdot)\in\mathbb{F}} \mathbb{E}_{\mathcal{U}} \left[\mathbf{c}^\top \mathbf{x}(\delta)\right] \; s.t. \; A(\delta)\mathbf{x}(\delta) \geq \mathbf{b}(\delta), \; \delta \in \mathcal{U} \right\}. \tag{$\mathcal{DR}$}$$

As noted in [16], such a search can be very difficult, as the formulation is #P-hard, i.e. harder than any NP-hard problem [55]. This leaves the problem severely computationally intractable for realistic sizes, so in order to be able to efficiently find solutions, some simplifications or restrictions are needed.

4.2 Affine Reaction Functions

Only looking at *linear* reactions is one such simplification, suggested by multiple authors [4, 12]. This results in affine reaction functions, consisting of a fixed intercept, and a linear reaction to the uncertainty. We call such functions linear decision rules, and they greatly simplify our search space. Letting $\hat{\mathbf{x}}$ denote the intercepts, and K^x be the matrix of slopes on δ, we can now write

$$\mathbf{x}(\delta) = \hat{\mathbf{x}} + K^x \delta,$$

or equivalently

$$x_j(\delta) = \hat{x}_j + \sum_i K^x_{ij} \delta_i, \quad \forall j.$$

Here, i is the index of an element in δ, and j the index of a variable. With this reformulation, $(\mathcal{DR})$ can be made computationally tractable by a transformation procedure. This procedure is closely described in Article 1, and results in a purely linear program if some conditions are met. First, the uncertainty set should be polyhedral, i.e. bounded by linear constraints and thus on the form $\mathcal{U} = \{ \delta : H\delta \leq h \}$. Second, the problem should have a fixed recourse, which means that the constraint matrix A should not depend on δ. Third, the right hand side should be affine in δ, so $b(\delta) = \hat{b} + M\delta$.

Available online at www.sciencedirect.com
SciVerse ScienceDirect
Energy Procedia 00 (2015) 000–000
Energy Procedia
www.elsevier.com/locate/procedia

2015 Workshop on Hydro Scheduling in Competitive Electricity Markets

Linear Decision Rules for seasonal hydropower planning: Modelling Considerations

Simen V. Braaten[a], Ola Gjønnes[a], Knut Hjertvik[a,*]

[a] *Norwegian University of Science and Technology, NO-7491 Trondheim, Norway*

Abstract

The seasonal hydropower planning problem is presented as a robust optimization model using linear decision rules (LDR). Uncertainty is present in multiple inflow series and the price of electricity. The objective is to generate weekly policies that maximize profit. The LDR approximation is effective at reducing computational complexity, and is well-suited to multistage problems. By restricting the decision variables to be affine reactions of the realizations of the uncertain parameters, the original intractable problem is transformed into a tractable one with shorter computational time. This article is intended as an introduction to LDR for practitioners within the hydropower industry. The basic results will be presented, along with explanations of the ideas behind the method. An example case is given to illustrate the solution method.

Peer-review under responsibility of SINTEF Energi AS.

Keywords: OR in energy; Robust optimization; Seasonal planning problem; Hydropower optimization; Linear decision rules

1. Introduction

This paper uses linear decision rules (LDR) to develop a robust seasonal planning model for hydropower scheduling. Hydropower is the world's leading renewable energy source with nearly 3500 TW h generated in 2010, thereby accounting for 16 % of the electricity generation [1].

* Corresponding author. Tel.: +47-908-44-997
E-mail address: knutjonn@stud.ntnu.no

Peer-review under responsibility of SINTEF Energi AS.

Nomenclature

Indices

t, τ	time
r, ρ	reservoir
d	decision type
j	uncertainty polytope constraint
v	inflow series
u	general uncertainty series

Sets

$\mathcal{T}$	time periods
$\mathcal{R}$	reservoirs
$\mathbb{D}$	decision types
$\mathbb{C}_r^d$	connected reservoirs in reservoir r for decision type d
$\mathbb{U}$	uncertainty series
$\mathcal{M}_t^d$	time periods whose uncertain parameters influence decision d at time t

Parameters

F_{rt}	natural inflow to reservoir r at time t [Mm^3]
π_t	power price at time t [€/MW h]
T^H, T^S	number of hours in week and million seconds in time period, respectively
β_t	discount factor at time t
$E_{rt}^d, (\bar{E}_r^d)$	energy coefficient (at reference head) for decision d in reservoir r at time t [MJ/m^3]
U_{rt}^d, L_{rt}^d	upper/lower bound for decision d for reservoir r at time t
$\overline{M}_{rt}, \underline{M}_{rt}, M_r^0$	upper limit/lower limit/initial volume for reservoir r at time t [Mm^3]
$\overline{D}_{rt}, \underline{D}_{rt}$	upper/lower discharge limit for reservoir r at time t [m^3/s]
C_{rt}^B	minimum bypass from reservoir r at time t [m^3/s]
$F_{rt\tau}^V, \Pi_{t\tau}$	mapping from the uncertainty in τ into the realization of inflow/price at time t in reservoir r [Mm^3]

Variables

x_{rt}^d	decision of type d for reservoir r at time t [m^3/s]
m_{rt}	level in reservoir r at the end of time t [Mm^3]
$\hat{x}_{rt}^d$	LDR intercept of decision d for reservoir r at time t [m^3/s]
$K_{rtu\tau}^d$	LDR slope of decision d, for reservoir r at time t for the realization of uncertainty u at τ [m^3/s]
μ_{rtj}^{bd}	dual variables associated with the j-th constraint of the uncertainty polytope of decision d for r at t, b being either upper (U) or lower (L) primal bound

It is expected to increase by 3.1 % annually for the next 25 years [2]. In order to meet this demand and maximize benefit, developing well-grounded generation strategies is highly relevant for production planners.

A hydropower producer is typically faced with the problem of deciding *when* to discharge water from their reservoirs. Discharging water for generation leads to profits right now, while saving water in the reservoir enables you to get higher head levels and sell more electricity later on, at potentially higher prices. This scheduling problem is usually adressed with a hierarchy of short-term and longer-term models [3–5]. A longer-term model, for instance a seasonal planning model, can be used to produce estimated water values for the short-term model.

Additionally, a seasonal planning model could create policies that serve several purposes. Within accounting, the policies can be used for budgeting future revenues, while in risk management the policies are valuable when measuring the overall exposure, for instance the amount of possible future generation that is price-dependent for different seasons [6]. For risk managers who want to evaluate future generation in their own price scenarios, price-dependent policies are preferable to state-contingent plans that are made according to other, potentially incompatible, scenarios.

It might be challenging to find policies for optimal reservoir operations, due to tight operational constraints, a complex dynamic system of interconnected reservoirs, and power output being a non-linear and non-convex function of head and water discharge [7]. Uncertain future price and inflow levels complicate matters further. Such a problem translates well into robust mathematical optimization. The main complications lie in ensuring sufficient water levels without having to unnecessarily spill water in high inflow scenarios.

Many existing approaches for solving the seasonal planning problem rely on stochastic dynamic programming (SDP) or stochastic dual dynamic programming (SDDP). Both of these methods have been very successful in this industry, but neither is without fault. SDP requires discretization of the state space, leading to an exponential increase of the computational effort with the number of state variables [8]. SDDP, on the other hand, cannot easily handle jointly stochastic inflow and price, since it can lead to objective functions that are neither convex nor concave [9]. Still, these are the standard solution approaches today, and routinely deliver good results.

LDR offer a way to take several types of uncertainties into account in recourse decisions, without resorting to discretization of possible outcomes. This is achieved by letting the recourse actions be functions of the uncertain parameters. A set of such reaction functions can be seen as a policy, i.e. a coherent mapping from any realization of uncertainty to a proper recourse action. In order to efficiently find good solutions, the search is restricted to affine functions. As will be seen, this approximation reduces to an, admittedly large, purely linear program.

In recent years, several other papers regarding LDR in hydropower planning have been published. An iterative procedure that can handle non-linear value functions is presented in [10], while [11,12] do analysis on the optimality gap, applying LDR to both the primal and the dual of the seasonal hydropower planning problem.

This paper is intended as an introductory guide to understanding, modelling and implementing LDR in practice, specifically in the context of hydropower scheduling. It presents a new approach to LDR in reservoir management, with zero-mean forecast errors as uncertainty parameters, and slopes on multiple uncertainty types. Additionally, it discusses issues related to limited dependencies on uncertainty parameters more thouroughly than earlier literature. These contributions to the field will be shown to give compact models that return sensible schedules, and intuitively understandable policies.

The rest of this paper is organized as follows. First, the hydropower planning problem is introduced and formulated as a deterministic linear program in Section 2. A description of the uncertainty is given in Section 3, before the reformulation into an LDR model is done in Section 4. An illustrative example is presented in Section 5. Finally, concluding remarks are given in Section 6.

2. Problem formulation

The problem at hand involves optimally deciding a seasonal production plan. Given a forecast for weekly prices and inflow, the aim is to maximize profit. Specifically, discharge for generation, x^q_{rt}, pumping, x^p_{rt}, and spill, x^s_{rt} should be decided for all reservoirs r at all times t. For simplicity of notation, the three types of decisions constitute the *decision set* $\mathbb{D} = \{q, p, s\}$. From these decisions, the reservoir levels m_{rt} are deduced. All the variables need to be within certain bounds, that may be time dependent. Pumping, reservoir level and generation have upper and lower limits, while spill only needs to be non-negative. Additionally, reservoirs have limits on *total discharge*, i.e. water for generation together with water that bypass the connected power station.

$$
\begin{aligned}
\max \quad & T^H \sum_{t\in\mathcal{T}} \beta_t \pi_t \sum_{r\in\mathcal{R}} \sum_{d\in\mathbb{D}} E^d_{rt} x^d_{rt} & & (1)\\
\text{s.t.} \quad & m_{r0} = M^0_r & r\in\mathcal{R} \quad & (2)\\
& m_{rt} = m_{r(t\text{-}1)} + F_{rt} + T^S \sum_{d\in\mathbb{D}} \left(\sum_{\rho\in\mathbb{C}^d_r} x^d_{\rho t} - x^d_{rt} \right) & r\in\mathcal{R}, t\in\mathcal{T} \quad & (3)\\
& \underline{M}_{rt} \le m_{rt} \le \overline{M}_{rt} & r\in\mathcal{R}, t\in\mathcal{T} \quad & (4)\\
& L^q_{rt} \le E^q_{rt} x^q_{rt} \le U^q_{rt} & r\in\mathcal{R}, t\in\mathcal{T} \quad & (5)\\
& L^p_{rt} \le E^p_{rt} x^p_{rt} \le U^p_{rt} & r\in\mathcal{R}, t\in\mathcal{T} \quad & (6)\\
& \underline{D}_{rt} \le x^q_{rt} + C^B_{rt} \le \overline{D}_{rt} & r\in\mathcal{R}, t\in\mathcal{T} \quad & (7)\\
& x^s_{rt} \ge 0 & r\in\mathcal{R}, t\in\mathcal{T} \quad & (8)
\end{aligned}
$$

All reservoirs are considered to have a power station immediately downstream, either real, or a dummy station without generation capabilities if the reservoir discharges directly to another reservoir. Power stations can therefore be omitted from the model formulation, as seen in the simplified topology in Fig. 1. The objective function (1) maximizes the present value of the difference between revenues and costs for all decision types. Here, prices are given per MW h, while decisions are in m^3/s. Thus, energy coefficients E^d are needed to convert volumetric sizes to power levels, and the number of hours in a time period, T^H, is used to get the total energy level. The energy coefficients are defined such that generation produces, pumping costs, and spill neither produces nor costs electricity. Equalities (2) and (3) ensure that the reservoir balance starts at its initial level and is updated between time periods. The reservoir level in one period depends on the previous level and the natural inflow, as well as the decisions made both in connected upstream reservoirs, $\mathbb{C}^d_r$, and in the current reservoir. Constraints

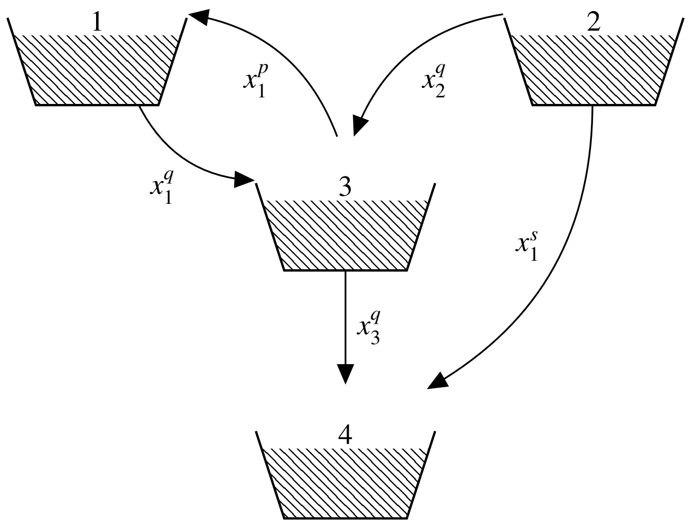

Fig. 1: An illustration of some of the variables used in the formulation.

(4)-(8) impose upper and lower bounds on reservoir level, generation, pumping, discharge, and spill, respectively. Pumping and generation limits are given in MW, so energy coefficients are also needed in the generation constraints (5) and pumping constraints (6). Notably, pumping is considered to be decided at the upstream reservoir, and can simply be seen as negative generation at a different energy coefficient. Thus, inequalities (6) typically have an upper bound of zero and always a non-positive lower bound. Since reservoir levels and natural inflows are given in Mm^3, the decisions that are given in m^3/s need to be multiplied by the length of a time period. Constraints (7) are the total discharge limits, so here bypass (C^B_{rt}) is considered in addition to the water for generation.

Some details are omitted due to the weekly time resolution. Omitted restrictions include ramping constraints, variable efficiency, costs related to start and shutdown of operations, and time delay for water travelling between reservoirs. Furthermore, the discount rate is assumed deterministic. Water may be spilled at unlimited rates, and all spill is assumed to end up at its intended target, be that another reservoir, or the ocean. Lastly, bypass is assumed fixed at its lower bound.

3. Introducing uncertainty

The deterministic model assumes future inflow and price to be known with certainty. In reality, however, it is rare that forecasts are accurate beyond the very close future. Regulations and capacities, on the other hand, are certain, and need to be respected regardless of the realized inflow or price.

Therefore, (1)-(8) should be defined to hold for any possible realization of inflow and price. This can be done by considering some discrete set of possible scenarios, and for each of them ensure that every constraint holds [4]. Scenarios can be predefined like in classical stochastic programming or SDP, but in the context of hydropower scheduling, in-run simulations, e.g. in SDDP, are more common. However, such scenarios only cover a subset of the actual continuous possibility space, and extrapolation is needed when faced with other scenarios.

Based on the research in [13,14], the LDR approach has emerged as an alternative. Here, the proposition is that each recourse decision in a stage needs only one, function valued, solution.

To understand this, consider representing scenarios as points in a cartesian space. For instance, in one scenario the price could be $\pi = 15$ €/MW h, and the inflow $F = 12\,\mathrm{Mm}^3$, represented as the point $(15, 12)$ in the plane. Now, all possible scenarios together form the *uncertainty set*, see Fig. 2. In principle, this subspace can take any shape or form, but one natural modelling choice is a polytope [15]. This general shape has the properties of being simply connected, convex, and easily described by linear constraints on the parameters.

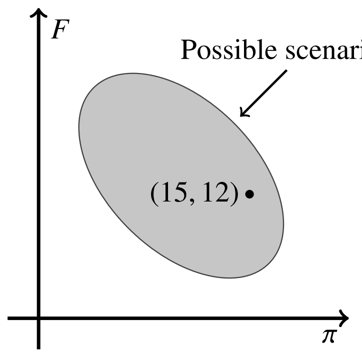

Fig. 2: Realization of possible scenario (15,12) in uncertainty set

Then, the idea of decision rules can be seen as replacing decision variables $x_{\delta'}$ for a number of selected scenarios $\delta' \in \mathcal{U}$, with a function $X : \mathcal{U} \to \mathbb{R}$. An illustration of this is given in Fig. 3. Here, X can be seen as a policy, while the actual decision in any given scenario is found by evaluating the function at the corresponding point $X(\delta')$.

(a) Scenario based model

(b) LDR approximation

Fig. 3: Illustrations, adapted from [14], of feasible regions and decisions in scenario based models and LDR approximation. LDR restrict the decisions to lie on a hyperplane imposed by the uncertain parameters.

Allowing X to take any functional form makes it very difficult to find *the* optimal functions [16]. However, by restricting the search space to functions that are affine with respect to δ, significant performance gains can be achieved [13,14].

To bring this back to the context of the model presented in Section 2, consider how the uncertainty parameters δ and the uncertainty set $\mathcal{U}$ are defined. The realized inflow and price could be described by

$$F_{rt} = \hat{F}_{rt} + \sum_{\tau \in \mathcal{T}} F^V_{rt\tau} \delta_{V(r)\tau} \qquad r \in \mathcal{R}, t \in \mathcal{T} \tag{9}$$

$$\pi_t = \hat{\pi}_t + \sum_{\tau \in \mathcal{T}} \Pi_{t\tau} \delta_{P\tau} \qquad r \in \mathcal{R}, t \in \mathcal{T}. \tag{10}$$

In these definitions, F denote inflow, and π price, while $\hat{F}$ and $\hat{\pi}$ denote their respective forecasted values. Every reservoir r belongs to an inflow series $V(r) = v \in \mathcal{V}$, where reservoirs in the same series face similar uncertainty, because as seen in (9), the same parameters δ_v decide their deviations. Additionally, there is the series, P, of parameters, δ_P, corresponding to the price. In total, these give rise to the set of uncertainty series $\mathbb{U} = \mathcal{V} \cup \{P\}$.

Each uncertainty series $u \in \mathbb{U}$ has its corresponding vector δ_u of uncertain parameters, with one element for each time period. The matrices containing $F^V_{rt\tau}$ and $\Pi_{t\tau}$ represent the mapping from the uncertainty parameters δ to the actual deviation from forecasted price and inflow. While complex relationships between periods can be modelled through these matrices, letting them be diagonal gives rise to some attractive simplifications. That is, for each time period t, $F^V_{rt\tau}$ is only non-zero when $t = \tau$, so natural inflow in period t only depends on the δ of period t. This gives δ_{vt} a natural interpretation, as the normalized deviation from the forecasted inflow during period t at reservoirs in inflow series v. With non-diagonal mapping matrices, it might be difficult to find the proper dependencies between time periods in historical data, and estimate historical values of δ.

The uncertainty set considered in this paper is a polytope, described by constraints on the form

$$-1 \leq \quad \delta_{ut} \quad \leq \quad 1 \qquad u \in \mathbb{U},\ t \in \mathcal{T} \tag{11}$$

$$-\Gamma^L_i \leq \sum_{t \in \mathcal{T}} \delta_{ut} \leq \Gamma^U_i \qquad u \in \mathbb{U}, \tag{12}$$

as illustrated in Fig. 4. The uncertain parameters $\delta_{u\tau}$ are normalized deviations from the forecast, so Eqs. (11) and (12) define how much any scenario can deviate from the expected case represented by (0,0). Individual parameters are bounded by [-1,1], and the cumulative normalized deviation is limited by the parameter Γ. In the remainder of the paper, the uncertainty set will be denoted by the shorthand $\mathcal{U} = \{\delta \mid H\delta \leq h\}$.

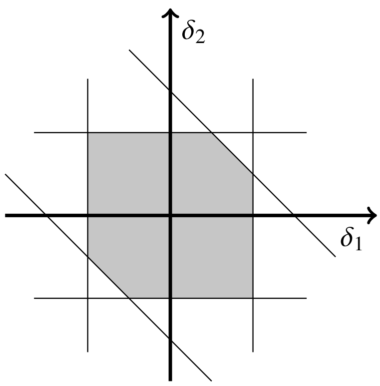

Fig. 4: Uncertainty polytope as described by (11) and (12)

4. Linear decision rules and reformulation

Understanding the parameters δ, decisions of how much to discharge, pump and spill at reservoir r at time t can now be defined as affine functions. As an example, the amount of discharge for generation becomes

$$x_{rt}^{q} = \hat{x}_{rt}^{q} + \sum_{u\in\mathbb{U}} \sum_{\tau\in\mathcal{M}_t^u} K_{rtu\tau}^{q} \delta_{u\tau}. \tag{13}$$

Here, the intercept $\hat{x}_{rt}^{q}$ and slopes $K_{rt\tau\upsilon}^{q}$ uniquely define the function, as long as the uncertainty parameters are linearly independent [17]. The intercepts represent scheduled values, while the slopes represent the real-time adjustments. Note that these adjustments depend on the realizations of the uncertain parameter $\delta_{u\tau}$ for all time periods τ in the memory set $\mathcal{M}_t^u$.

The memory set contains the periods a decision is dependent on, and can include all or some of the previous time periods, but never future ones. Allowing the LDR to depend on only a subset of the uncertain parameters will speed up computation, but reduce the flexibility of decisions. Some ways to reduce the size of the memory set include limiting the number of periods the decision can look backwards, and neglecting the uncertainty in time periods where the reservoir level typically is well within its upper and lower limits. Similarly, one could restrict what uncertainty series a variable is allowed to depend on, but in the following, it is assumed that every decision depends on every uncertainty series.

Now, a reformulation of the deterministic model to a robust optimization problem with LDR is described, following the setup in [18]. The aim is to show how a purely linear model can be used to find the intercepts and slopes as decision variables, and thus identify the optimal affine reactions. First, the inequalities for discharge, pumping and spill are converted, before the reservoir volume inequalities and the objective are transformed.

4.1. General inequalities

Consider inequality (5) rewritten with the function valued variable for generation,

$$\hat{x}_{rt}^{q} + \sum_{u\in\mathbb{U}} \sum_{\tau\in\mathcal{M}_t^u} K_{rt\tau u}^{q} \delta_{u\tau} \leq \frac{U_{rt}^{q}}{E_{rt}^{q}}, \quad r \in \mathcal{R}, t \in \mathcal{T}, \delta \in \mathcal{U}. \tag{14}$$

Rewriting (1)-(8) this way would yield an LP where the intercepts and slopes are the decision variables, so the function valued search is reduced to an LP. However, since there should be one such constraint for each point in $\mathcal{U}$, and there are uncountably many points in a polytope, the problem will still be impossible to solve. As noted in [18], this problem can be circumvented, because only constraints describing the worst case is really needed. It is not trivial, though, to tell which constraint will be the strictest. To this end, the shape of the uncertainty set can be utilized. As a polytope defined by $H\delta \leq h$, it can be used as the feasible region for a linear program determining the parameter values maximizing the LHS.

$$\max_{\delta} \left\{ \begin{array}{l} \displaystyle\sum_{u\in\mathbb{U}} \sum_{\tau\in\mathcal{M}_t^u} K_{rt\tau\upsilon}^{q} \delta_{u\tau} \\ \displaystyle s.t. \sum_{u\in\mathbb{U}} \sum_{\tau\in\mathcal{T}} H_{ju\tau} \delta_{u\tau} \leq h_j, \quad \forall j \quad : \mu_{rtj}^{Uq} \end{array} \right\} \leq \frac{U_{rt}^{q}}{E_{rt}^{q}} - \hat{x}_{rt}^{q} \qquad r \in \mathcal{R}, t \in \mathcal{T} \tag{15}$$

This formulation of the constraints is equivalent with (14), but has finite cardinality. However, it is not linear as both the slopes and the uncertainties are decision variables in the objective. To regain a completely linear structure, LP duality theory [19] can be exploited to unlink the different variables K and δ (Note the dual variables in gray with their corresponding constraint). An equivalent formulation to (15) is

$$\min_{\mu^{Uq}} \left\{ \begin{array}{lll} \sum_j h_j \mu^{Uq}_{rtj} & & \\ s.t. \sum_j H_{ju\tau}\mu^{Uq}_{rtj} = K^q_{rt\tau\upsilon}, & u \in \mathbb{U}, \tau \in \mathcal{M}^u_t & : \delta_{u\tau} \\ \sum_j H_{ju\tau}\mu^{Uq}_{rtj} = 0, & u \in \mathbb{U}, \tau \notin \mathcal{M}^u_t & : \delta_{u\tau} \\ \mu^{Uq}_{rtj} \geq 0, \quad \forall j & & \end{array} \right\} \leq \frac{U^q_{rt}}{E^q_{rt}} - \hat{x}^q_{rt} \qquad r \in \mathcal{R}, t \in \mathcal{T}. \tag{16}$$

Here, the equalities come from the fact that the δ-variables are free. Now, this formulation only requires that $\sum_j h_j \mu^{Uq}_{rtj} \leq \frac{U^q_{rt}}{E^q_{rt}} - \hat{x}^q_{rt}$ for *some* set of μ^{Uq}. That is, the set of constraints that needed to hold true *for all* realizations is replaced with a set of constraints that only need to be satisfied *once*. Thus, the original (14) holds whenever the following set of constraints have a solution:

$$\sum_j h_j \mu^{Uq}_{rtj} \leq \frac{U^q_{rt}}{E^q_{rt}} - \hat{x}^q_{rt}, \quad r \in \mathcal{R}, t \in \mathcal{T} \tag{17}$$

$$\sum_j H_{ju\tau}\mu^{Uq}_{rtj} = K^q_{rt\tau\upsilon}, \quad r \in \mathcal{R}, t \in \mathcal{T}, u \in \mathbb{U}, \tau \in \mathcal{M}^u_t \tag{18a}$$

$$\sum_j H_{ju\tau}\mu^{Uq}_{rtj} = 0, \quad r \in \mathcal{R}, t \in \mathcal{T}, u \in \mathbb{U}, \tau \notin \mathcal{M}^u_t \tag{18b}$$

$$\mu^{Uq}_{rtj} \geq 0, \quad r \in \mathcal{R}, t \in \mathcal{T}, j. \tag{19}$$

Similar transformations can be carried out for all the original constraints, to formulate a robust linear program in which (4)-(8) holds for any realization of the uncertainty. Further examples are omitted here for brevity, as the transformations are very similar to the one shown above. However, there are some subtleties to be noted.

First, constraints (18) are defined for *every* combination (u, τ), even if the decision in question does not depend on $\delta_{u\tau}$. This simply means that $\delta_{u\tau}$ has objective coefficient 0 in the primal in (15), but the variable is still present in the constraints below. Thus, the corresponding dual constraint (18b) should have a right hand side of 0. Second, completely analogous transformations could be done to $\geq$-restrictions. But here, the left hand side becomes a minimization problem in the primal space, so the restrictions $H\delta \leq h$ are non-standard and give rise to non-positive dual variables [19]. Third, equality constraints can be handled in a simpler way, see [18].

4.2. Reservoir volumes

The reservoir volumes m_{rt} are not decisions themselves, but merely consequenses of all previous discharge, pumping and spill decisions, which again depend on all previous realizations of uncertainty. This makes limited memory ($|\mathcal{M}^u_t| < t$) impossible[1]. The idea is that the reservoir balance is simply

[1] Similarly, if restrictions are imposed on which uncertainty series decisions are allowed to depend on, these restrictions cannot apply as strictly to the reservoir level. This is because the reservoir level is a consequence of decisions at upstream reservoirs as well, and these can depend on other series.

the initial level plus the sum of all natural inflows, adjusted for all decisions in the current and connected reservoirs. A more detailed deduction is given in the appendix. Also note that the purpose of stating Eqs.(2)-(3) in the deterministic model is for readability, and they can be substituted into (4).

If the expressions of the linear decision rules for discharge, pumping and spill are inserted, then (2)-(4) can be desribed by the following inequalities[2],

$$\underline{M}_{rt} \leq \sum_{\tau=1}^{t} \left(\hat{F}_{r\tau} + \sum_{d \in \mathbb{D}} \Big[\sum_{\rho \in \mathbb{C}_r^d} \hat{x}^d_{\rho\tau} - \hat{x}^d_{r\tau} \Big] + \sum_{u \in \mathbb{U}} \bigg(\sum_{\substack{t^* \leq t \\ \tau \in \mathcal{M}^u_{t^*}}} \sum_{d \in \mathbb{D}} \Big[\sum_{\rho \in \mathbb{C}_r^d} K^d_{\rho t^* u\tau} - K^d_{rt^* u\tau} \Big] + I^S_{ur} F^V_{r\tau} \bigg) \delta_{u\tau} \right) \tag{20}$$

$$+ M^0_r \leq \overline{M}_{rt},$$

and similar transformations as described in Section 4.1 can be made.

4.3. The objective function

At last, the definitions in (13) are employed on the objective function (1). Also inserting (10) yields

$$\mathbb{E}_\delta \left[T^H \sum_{t \in \mathcal{T}} \beta_t \Big(\hat{\pi}_t + \sum_{\tau \in \mathcal{T}_t} \Pi_\tau \delta_{P\tau} \Big) \sum_{r \in \mathcal{R}} \sum_{d \in \mathbb{D}} E^d \Big(\hat{x}^d_{rt} + \sum_{u \in \mathbb{U}} \sum_{\tau \in \mathcal{M}^u_t} K^d_{rtu\tau} \delta_{u\tau} \Big) \right]. \tag{21}$$

Exploiting that (21) is a linear expression and expanding the multiplication, it is possible to take advantage of the facts that $\mathbb{E}[\delta_{ut}] = 0$ and thus $\mathbb{E}[\delta_{Pt}\delta_{u\tau}] = \mathrm{Cov}(\delta_{Pt}, \delta_{u\tau})$, to get

$$\sum_{t \in \mathcal{T}} \sum_{r \in \mathcal{R}} \sum_{d \in \mathbb{D}} T^H \beta_t E^d_{rt} \left(\hat{\pi}_t \hat{x}^d_{rt} + \sum_{u \in \mathbb{U}} \sum_{\tau \in \mathcal{M}^u_t} \Pi_t K^d_{rtu\tau} \mathrm{Cov}(\delta_{Pt}, \delta_{u\tau}) \right). \tag{22}$$

This objective function contains two terms, the first one equivalent to the deterministic objective in Section 2, while the second one accounts for the uncertainty. The price depends on $\Pi_t \delta_{Pt}$, and the decisions depend on $K^d_{rtu\tau}\delta_{u\tau}$, so the expected value of their product is needed. This is exactly what the covariance provides for zero-mean variables, so this term allows the objective function to take into account both cross- and autocorrelations between the stochastic parameters.

5. Illustrative example

A two-reservoir system in cascade is considered with a 52-week planning horizon starting in January. The larger reservoir upstream (M7) is connected to a pumped-storage facility, and has the highest degree of regulation (DoR). The downstream reservoir (M6) can only generate power. M7 and M6 belong to different inflow series, denoted in this example by their respective reservoir names. The capacities for storage and generation are listed in Table 1. As price and inflow forecasts, the mean of 3000 scenarios is used. They are generated by a time series model, tuned to match data from a sofisticated market model [20]. The time series formulation is similar to what is outlined in [21]. The weekly standard deviation in inflow ranges from 40 % to 140 % of the forecasted value, averaging at

[2] Here, binary parameters I^S_{ur} are 1 if u is the inflow series reservoir r belongs to, i.e. $u = V(r)$.

78 %. The uncertainty polytope is also defined as the convex hull of these scenarios, so the problem is feasible for each of them.

		M7	M6
Res.vol.	[Mm^3]	300	100
U^q_{rt}	[MW]	150	50
DoR		1.00	0.14

Table 1: Reservoir volume, maximum production capacity (U^q_{rt}) and degree of regulation (DoR). Pumping capacity at M7 is equal to the generation capacity

The LDR formulated in Section 2 have been implemented and solved in IBM ILOG CPLEX Optimization Studio 12.6.1.0 on a 64-bit Windows 7 Enterprise, Intel® Core™i7-3770 3.40 GHz, 16.0 GB RAM. The formulations yield an LP of 131 021 rows, 295 128 columns, and a run time of 249 seconds. The model is implemented with a full memory set, i.e. every previous time period is in the memory of every decision. Additionally, every reservoir has a target volume of 66.8 % in week 52, a normal reservoir level in this time period. The optimal policies are simulated on the 3000 scenarios. Percentiles of the resulting reservoir trajectories are plotted in Fig. 5a and Fig. 5b.

As can be seen, the LDR yield rather similar trajectories regardless of the inflow and price development. This is indicative of the model trying to counteract deviations. A natural conclusion is that it is easier to find robust continuations if the reservoir level is known with some degree of certainty. In M7 a gradual dispersal of scenarios can be seen, while in M6 the weeks 22-40 see the biggest spread of trajectories. In weeks of greater uncertainty, such as the snow melting season around weeks 17 to 33, more variation in trajectories can be observed. Since M6 has the lowest DoR, it is intuitively the reservoir most likely to spill. Considering this, it makes sense that M6 is drained first, resulting in a rather sharp volume decrease in the initial weeks. Then, as M6 moves close to its lower bound, M7 is gradually emptied.

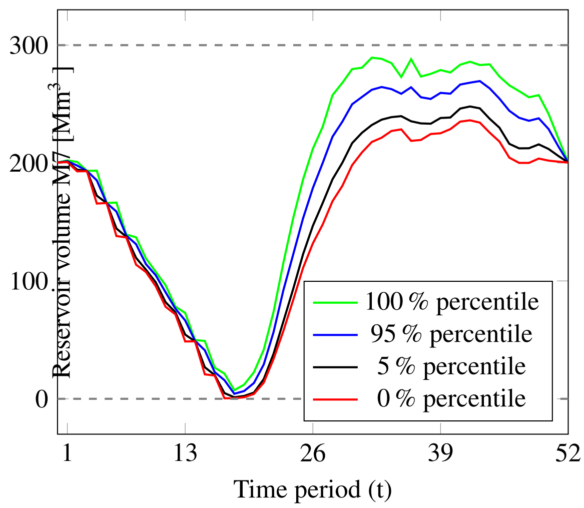

(a) Percentiles for reservoir volumes of M7

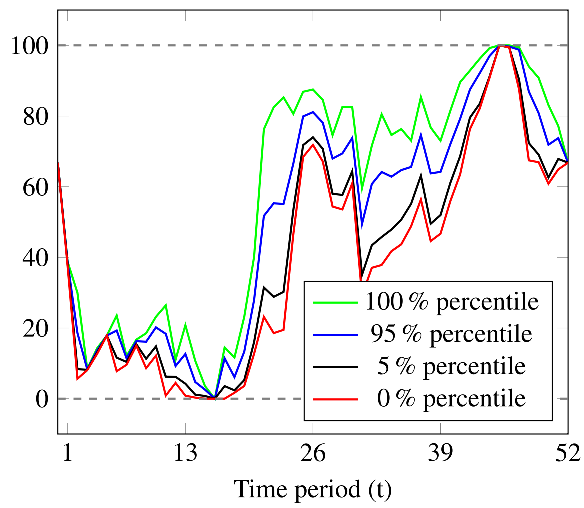

(b) Percentiles for reservoir volumes of M6

Fig. 5: Reservoir trajectories

Looking at the forecasted price in Fig. 6, it also seems logical that both reservoirs are emptied during the inital high price, and then start filling up at the very low prices during summer. In summary, the reservoir trajectories seem sensible.

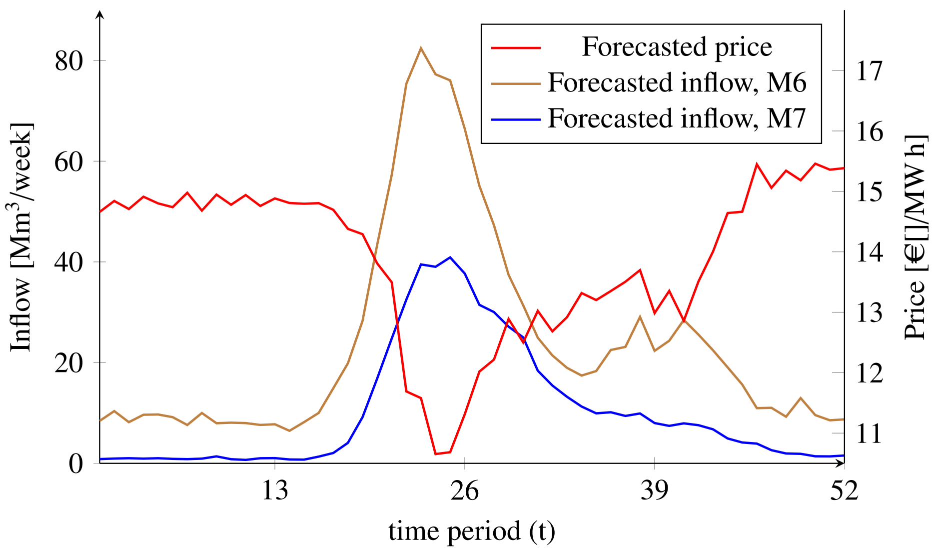

Fig. 6: Forecasted inflows and price

The next series of plots (Fig. 7-8) illustrate the slopes for discharge and pumping in every (t, τ)-combination for all uncertainty series. They indicate the amount of adjustment of a decision, given the realization of the price error, inflow error in M7 and inflow error in M6, respectively. The greater the magnitude of coefficients, the more the realization of $\delta_{u\tau}$ affects discharge and pumping in t. As one cannot react to the realization of future uncertainty, no elements appear above the diagonals.

In general, the discharge adjustment decisions are positively dependent on both the price and inflow deviation in the current time period. This means that if the realized price or inflow exceeds the forecast, the discharge is adjusted upwards, generating more power. Intuitively, the strongest dependencies are found in the reservoir's own inflow series. Generation decisions at M7 (Fig. 7) depend only on its own inflow series, while pumping decisions (Fig. 8) are entirely independent of both inflow disturbances, only reacting to the price in the previous period. Here one must note that since pumping is a negative variable, negative slopes imply an upward adjustment in pumping volume. Discharge decisions at M6 (Fig. 9), on the other hand, depend mainly on inflow at its own reservoir, but also on inflow to M7. These observations are reasonable, as M6 is the bottom reservoir, and its reservoir level depends on decisions made at M7. M7 does not need to take M6 into account, and pumping is used mainly as a financial opportunity, rather than as a tool to balance reservoir levels.

The possibility of reducing memory has been mentioned. Fig. 7-8 show that there are mainly two periods where long memory matters. First, there are dependencies between decisions around week 30 and realizations of inflow disturbance for up to the 10 previous weeks. The second period is the end of the scheduling horizon, where dependencies on many earlier periods are seen. Common for both periods is a need for a specific reservoir level. Around week 30, the reservoirs are approaching their maximum capacity, so accurate decisions keep the reservoirs from spilling. At the end, one aims for the target volume. Keeping more water than the minimum requirement does not give any benefits, but reduces the total generation, so here, higher accuracy pays off. It seems that the memory in the other time periods can safely be reduced, without affecting the optimal objective value much.

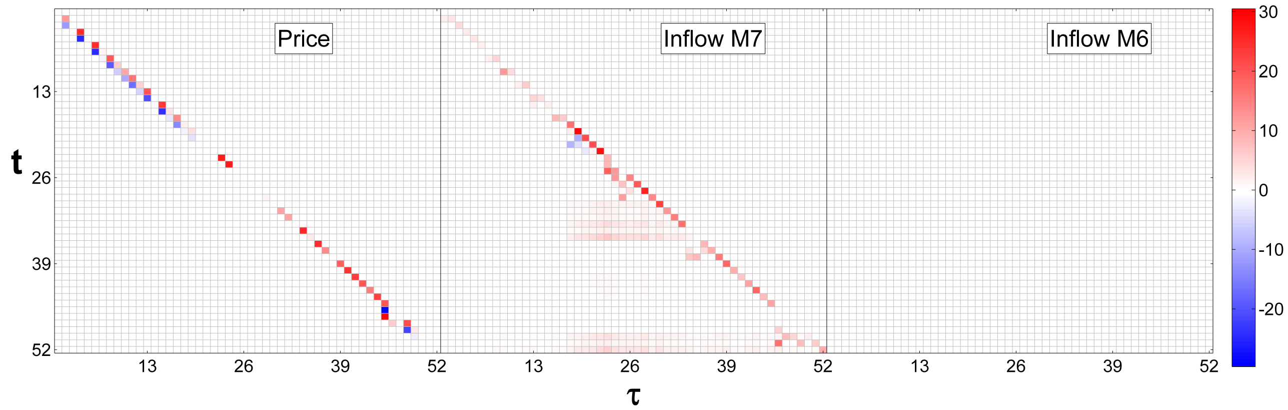

Fig. 7: Optimal slope matrix for discharge from the large upstream reservoir, $K^{q}_{M7tu\tau}$

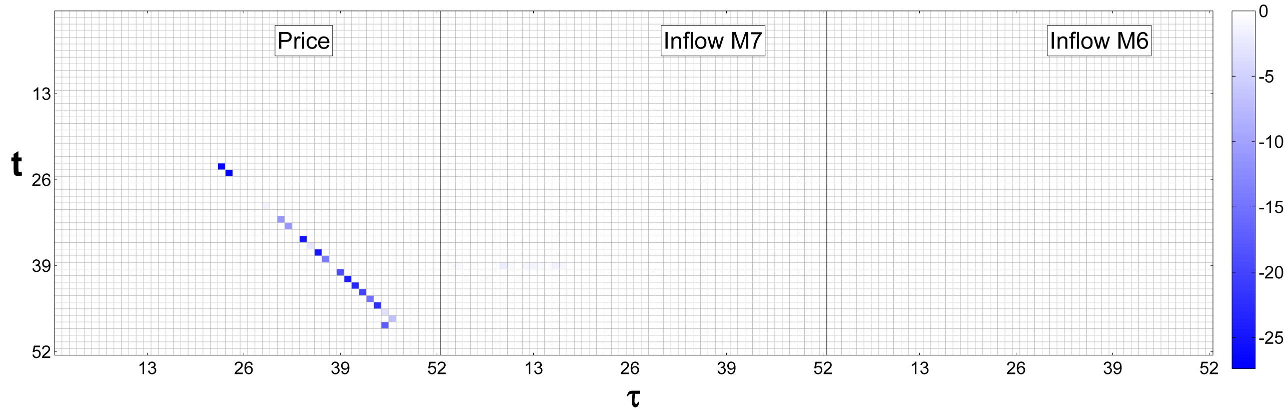

Fig. 8: Optimal slope matrix for pumping from the downstream to the upstream reservoir, $K^{p}_{M7tu\tau}$

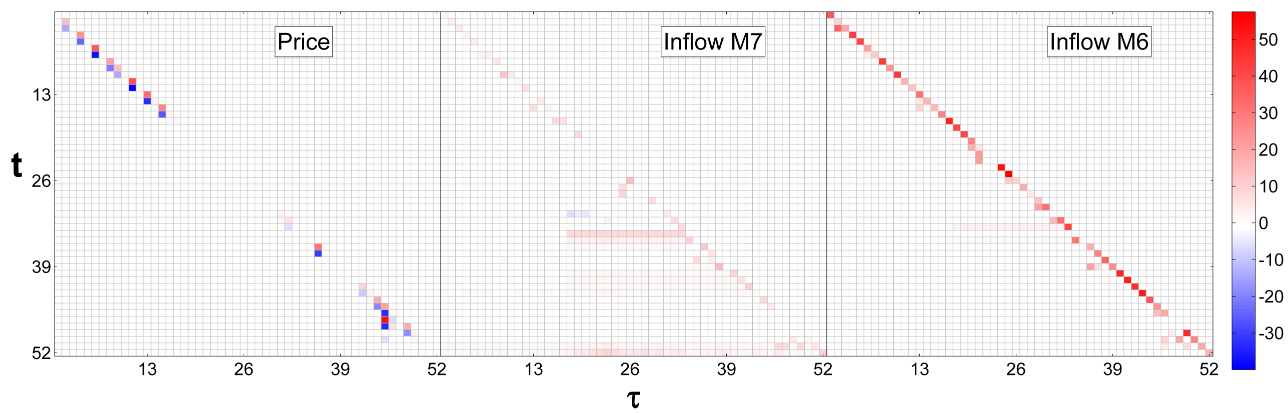

Fig. 9: Optimal slope matrix for discharge from the smaller downstream reservoir, $K^{q}_{M6tu\tau}$

6. Concluding remarks

In this work, the seasonal planning problem for a hydropower producer is considered. Decisions regarding the amount to discharge for generation, pumping and spillage are made in the face of uncertainty in inflow and price. A robust optimization model using LDR is proposed. Here, recourse decisions are decided by affine functions of the stochastic disturbances, rendering a tractable pure LP-formulation.

Results from an illustrative example provide credible decisions and reservoir trajectories. The obtained policies are robust with respect to deviations in inflow and price, and the level of discharge is positively dependent on the inflow and price disturbance.

Future research should include testing with larger water systems, and comparison of LDR with alternative solutions methods, both in terms of run time, objective value and schedule. It was seen that long memory was only significant in certain periods, suggesting that memory could be limited in others. Reducing memory is one way to improve the performance of LDR, so further research into this topic is recommended. Furthermore, quantifying the approximation error will be important to test LDR's applicability in seasonal hydropower scheduling.

7. Acknowledgments

The authors would like to thank Powel AS for providing the data, and employee Gro Klæboe for invaluable help. Additionally, our supervisor Stein-Erik has contributed significantly with useful ideas and suggestions. We also owe a thanks to Marco Zugno, who has co-written [18] which has been an inspiration, and kindly answered our questions.

References

[1] International Energy Agency, . IEA report sets a course for doubling hydroelectricity output by 2050; ???? URL: https://www.iea.org/topics/renewables/subtopics/hydropower/; available at: https://www.iea.org/topics/renewables/subtopics/hydropower/.

[2] Bagher, A.M., Vahid, M., Mohsen, M., Parvin, D.. Hydroelectric energy advantages and disadvantages. American Journal of Energy Science 2015;2(2):17–20.

[3] Fosso, O.B., Gjelsvik, A., Haugstad, A., Mo, B., Wangensteen, I.. Generation scheduling in a deregulated system. the Norwegian case. IEEE Transactions on Power Systems 1999;14(1):75–81.

[4] Wallace, S.W., Fleten, S.E.. Stochastic programming models in energy; vol. 10; chap. Ch: 10. Elsevier; 2003, p. 637–677.

[5] Fleten, S.E., Haugstvedt, D., Steinsbø, J.A., Belsnes, M., Fleischmann, F.. Bidding hydropower generation: Integrating short-and long-term scheduling. PCSS 2011;1:1–8.

[6] Fleten, S.E., Wallace, S.W.. Delta-hedging a hydropower plant using stochastic programming. Springer; 2009.

[7] Gjelsvik, A., Haugstad, A.. Considering head variations in a linear model for optimal hydro scheduling; 2005. Unpublished paper from SINTEF Energy Research, Trondheim, Norway.

[8] Stedinger, J.R., Sule, B.F., Loucks, D.P.. Stochastic dynamic programming models for reservoir operation optimization. Water Resources Research 1984;20(11):1499–1505.

[9] Rebennack, S., Iliadis, N.A., Pereira, M.V., Pardalos, P.M.. Electricity and CO_2 emissions system prices modeling. In: 2009 IEEE Bucharest PowerTech. 2009, p. 1–6.

[10] Pan, L., Housh, M., Liu, P., Cai, X., Chen, X.. Robust stochastic optimization for reservoir operation. Water Resources Research 2015;51:409–429.

[11] Rocha, P.. Medium-term planning in deregulated energy markets with decision rules. Ph.D. thesis; Imperial College London; 2013.

[12] Grønvik, I., Hadziomerovic, A., Ingvoldstad, N., Egging, R., Fleten, S.E.. Feasibility of linear decision rules for hydropower scheduling. In: 2014 International Conference on Probabilistic Methods Applied to Power Systems (PMAPS). 2014,.

[13] Ben-Tal, A., El Ghaoui, L., Nemirovski, A.. Robust Optimization. Princeton University Press; 2009.

[14] Kuhn, D., Wiesemann, W., Georghiou, A.. Primal and dual linear decision rules in stochastic and robust optimization. Mathematical Programming 2011;130(1):177–209.

[15] Bertsimas, D., Brown, D.B., Caramanis, C.. Theory and applications of robust optimization. SIAM review 2011;53(3):464–501.

[16] Georghiou, A., Wiesemann, W., Kuhn, D.. The decision rule approach to optimisation under uncertainty: methodology and applications in operations management. Mathematical Programming 2011;1:1–40.
[17] Warrington, J., Goulart, P., Mariethoz, S., Morari, M.. Policy-based-reserves for power systems. IEEE Transactions on Power Systems 2013;28(4):4427–4437.
[18] Zugno, M., Morales, J., Madsen, H.. Robust management of combined heat and power systems via linear decision rules. In: 2014 IEEE International Energy Conference (ENERGYCON). 2014, p. 479–486.
[19] Saltzman, M.J.. Linear programming duality; 1998. Available at: http://www.ces.clemson.edu/~mjs/courses/mthsc.440/duality.pdf.
[20] Wolfgang, O., Haugstad, A., Mo, B., Gjelsvik, A., Wangensteen, I., Doorman, G.. Hydro reservoir handling in Norway before and after deregulation. Energy 2009;34(10):1642–1651.
[21] 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. Springer; 2010, p. 33–55.

Appendix

In Section 4.2, the reservoir bounds seen in (4) are rewritten in terms of the linear decision rules. In the following section, a description of this procedure is given.

In (3), m_{rt} is defined in terms of other variables. In each week, the change in the reservoir levels from the last week is the sum of three parts. First, natural inflow increases the level. Second, decisions at all reservoirs directly upstream provide inflow. In the case of pumping, this is negative inflow, so it represents water leaving for an upstream reservoir. Third, decisions at the reservoir in question reduces the reservoir level (once again with the exception of pumping, which is negative reduction, adding water). This relation could be expressed

$$m_{rt} = T^S \sum_{d \in \mathbb{D}} \left(\sum_{\rho \in \mathbb{C}_r^d} x_{\rho t}^d - x_{rt}^d \right) + F_{rt} + m_{r(t\text{-}1)} \tag{23}$$

$$m_{rt} = T^S \sum_{d \in \mathbb{D}} \left(\sum_{\rho \in \mathbb{C}_r^d} \left[\hat{x}_{\rho t}^d + \sum_{u \in \mathbb{U}} \sum_{\tau \in \mathcal{M}_t^u} K_{\rho t u \tau}^d \delta_{u\tau} \right] - \left[\hat{x}_{rt}^d + \sum_{u \in \mathbb{U}} \sum_{\tau \in \mathcal{M}_t^u} K_{rtu\tau}^d \delta_{u\tau} \right] \right) + \left[\hat{F}_{rt} + F_{rt}^V \delta_{v(r)t} \right] + m_{r(t\text{-}1)}. \tag{24}$$

This formulation could be rolled out recursively, by expanding the definition of $m_{r(t\text{-}1)}$ on the second line. This definition, in turn, relies on an even earlier reservoir balance, and so on, until m_{r1} depends on the initial level $m_{r0} = M_r^0$. Regrouping into intercept and slope terms yields

$$\begin{aligned} m_{rt} = {} & \hat{F}_{rt} + \sum_{d \in \mathbb{D}} \left[\sum_{\rho \in \mathbb{C}_r^d} \hat{x}_{\rho t}^d - \hat{x}_{rt}^d \right] + \sum_{u \in \mathbb{U}} \sum_{\tau \in \mathcal{M}_t^u} \sum_{d \in \mathbb{D}} \left[\sum_{\rho \in \mathbb{C}_r^d} K_{\rho t u \tau}^d - K_{rtu\tau}^d \right] \delta_{u\tau} + F_{rt}^V \delta_{v(r)t} \\ & + \hat{F}_{r(t\text{-}1)} + \sum_{d \in \mathbb{D}} \left[\sum_{\rho \in \mathbb{C}_r^d} \hat{x}_{\rho(t\text{-}1)}^d - \hat{x}_{r(t\text{-}1)}^d \right] + \sum_{u \in \mathbb{U}} \sum_{\tau \in \mathcal{M}_{(t\text{-}1)}^u} \sum_{d \in \mathbb{D}} \left[\sum_{\rho \in \mathbb{C}_r^d} K_{\rho(t\text{-}1)u\tau}^d - K_{r(t\text{-}1)u\tau}^d \right] \delta_{u\tau} + F_{r(t\text{-}1)}^V \delta_{v(r)(t\text{-}1)} \\ & + \ldots \quad . \end{aligned} \tag{25}$$

Here, it is clear that the intercepts for each of the earlier time periods will appear once, and each $\delta_{u\tau}$ will appear in the expression for period t^* as long as $\tau \in \mathcal{M}_{t^*}^u$, the memory set of t^*. Let these periods t^* define the usage horizon $\mathcal{H}_\tau^u$ of $\delta_{u\tau}$. Additionally, let I_{ur}^S be 1 when $u = V(r)$ and 0 otherwise. This parameter is used to include natural inflow from the inflow series r belongs to, but exclude inflow from any others. Then, the long chain in (25) can be collapsed into

$$m_{rt} = \sum_{\tau=1}^{t}\left(\hat{F}_{r\tau} + \sum_{d\in\mathbb{D}}\left[\sum_{\rho\in\mathbb{C}_r^d}\hat{x}^d_{\rho\tau} - \hat{x}^d_{r\tau}\right] + \sum_{u\in\mathbb{U}}\left(\sum_{\substack{t^*\in\mathcal{H}^u_\tau\\ t^*\leq t}}\sum_{d\in\mathbb{D}}\left[\sum_{\rho\in\mathbb{C}_r^d}K^d_{\rho t^* u\tau} - K^d_{rt^* u\tau}\right] + I^S_{ur}F^V_{r\tau}\right)\delta_{u\tau}\right) + M^0_r. \tag{26}$$

Now, (4) can be expanded by inserting (26), and transformed in the same manner as presented for inequalties in Section 4.1.

5 Intermezzo

In Article 1, we presented the basics of the LDR transformation. To complement this, we will in Sections 5.1 and 5.2 give some explanatory notes. Then, in Section 5.3, a more detailed explanation of uncertainty sets is given. This topic was briefly utilized, but never elaborated upon in Article 1. Section 5.4 explains how our uncertainty set is generated. Last, theory is given on how upper bounds on the optimality gap can be found. Such upper bounds are important when evaluating the quality of LDR, and an introduction to the topic is needed before the implementation in Article 2.

5.1 Inequality constraints

The most interesting part of the LDR transformation is reformulating the inequality constraints. Equality constraints will be covered briefly in the next section, and the objective function is easily handled, so it will not be covered again here. Consider a set of inequality constraints

$$\sum_j A_j x_j(\delta) \leq b(\delta) \quad \forall \delta \in \mathcal{U}. \tag{1}$$

Rewriting in terms of affine functions, and separating intercepts and slopes yields

$$\sum_j \sum_i (A_j K_{ij}^x - M_i)\delta_i \leq \hat{b} - \sum_j A_j \hat{x}_j \quad \forall \delta \in \mathcal{U}, \tag{2}$$

which is shown in Article 1 to be equivalent to

$$\sum_k h_k \mu_k \quad \leq \hat{b} - \sum_j A_j \hat{x}_j \tag{3}$$

$$\sum_k H_{ik} \mu_k^{Uq} = \sum_j A_j K_{ij}^x - M_i, \quad \forall i \tag{4}$$

where k is the index of a constraint in the uncertainty polytope. Here, the importance of the polyhedral uncertainty set, fixed recourse, and affine right hand side is shown. These make LP duality theory applicable, to reformulate (2) into (3) and (4). If A depended on δ, or b was not affine, LP theory could not have been utilized. Where (2) has infinite cardinality and δ multiplied with K_{ij}^{x}, (3) and (4) are of finite cardinality and completely linear.

In the context of the problem at hand, the formulation in Article 1 has some notable properties. First, since most constraints in the original problem presented there, are of the simple form

$$\hat{x}_j + \sum_i K_{ij}^{x}\delta_i \leq \hat{b}.$$

This makes constraints of form (3) and (4) relatively simple, with only one variable, $A_1 = 1$ and $M_i = 0$. However, some subtleties arise from the introduction of *limited memory*. Specifically, each decision in time period t has a set $\mathcal{M}_t$ of time periods on which it depends. Every other time period is ignored. The reason for the expression 'limited memory' is that one natural way of implementing this, is by letting each time period *remember* only the last few periods. This arises as a natural way to reduce the size of the problem, and leads to some attractive simplifications. If x_j is independent of a certain δ_i, it means that the slope K_{ij}^{x} does not exist, and a column is cut from the problem. On the other hand, δ_i still needs to be taken into account when deciding the worst case value of x_j, so none of the constraints (4) are removed. Additionally, special care needs to be taken in constraints where variables with different memories appear together, like the reservoir balances that take decisions at other reservoirs into account. Limited memory could be implemented by explicitly requiring that the affected slopes $K_{ij}^{x} = 0$, but our approach utilizes human knowledge of the problem at hand, and reduces the need to trust a potentially unknown engine presolve procedure to remove these variables.

5.2 Equality constraints

As pointed out in [59], constraints on the form

$$\hat{x} + \sum_i K_i^x \delta_i = \hat{y} + \sum_i K_i^y \delta_i$$

can be split term by term, when uncertain parameters are linearly independent, and are therefore equivalent to

$$\begin{aligned} \hat{x} &= \hat{y} \\ K_i^x &= K_i^y, \quad \forall i. \end{aligned}$$

This is because the linear independence implies that a certain δ_i cannot be expressed by the other δ, so for the expressions to stay equal when δ_i varies, the slopes K_i on each side must be the same. These equations are generally not very complex, and simply act as definitions of placeholder variables. As we saw in Article 1, we opted for a slightly different approach when handling the equality constraints in the reservoir balance. In this case, since the reservoir balance simply defined the reservoir level, we substituted this expression into the bounds on the reservoir level.

5.3 Uncertainty sets

To fully utilize the benefits of LDR, a meaningful representation of the uncertainty set $\mathcal{U}$ is needed.

One often faces a trade-off between covering every imaginable realization of the uncertain parameters, and minimizing the complexity and size of the uncertainty set. The easiest to interpret, and the most robust choice, is a *box uncertainty set* that contains the full range of realizations for every single element δ of the uncertainty. However, it is highly unlikely that every uncertain parameter takes its worst-case value simultaneously. Thus, smaller uncertainty sets which are still nearly impossible to violate, can often be described. For instance, ellipsoidal uncertainty sets are used to solve an LDR hydropower planning formulation in [34]. These disallow the corners in box uncertainty sets, and are therefore tighter, without being much harder to describe. Nevertheless, as we saw, polyhedral sets are the most popular and natural choice. These can be described by a set of linear constraints, and are convex and simply connected. There are many types of polytopes, and here we will look at the linear equations that are used to describe ours.

We use a typical class of polyhedral sets, known as *budgeted uncertainty sets* or *Bertsimas and Sim* uncertainty sets [7], on the form:

$$\mathcal{U} = \{\delta : ||\delta||_1 \leq \Gamma \quad ||\delta||_\infty \leq 1\}, \tag{5}$$

where $\delta \in \mathbb{R}$. Here, the uncertain parameters δ are normalized deviations from the forecast. These are continous, and can take both positive and negative values, with a maximal magnitude of 1. The parameter Γ in budgeted uncertainty sets is often interpreted as the highest number of worst case periods, but with δ free, it takes a somewhat different meaning. Slightly simplified, it is here the highest allowable *difference* between the numbers of periods with 1 and -1 realizations. This is complicated by the fact that values between -1 and 1 are also allowable, but the basic principle is the same.

The basic framework described above is sufficient for understanding, and was how the uncertainty polytope was presented in Article 1. However, in the final version of our uncertainty polytope, we have refined the formulation somewhat. We have added the possibility of restricting with Γ-parameters for the whole planning horizon, or for other sets U_i of weeks. Here, the index i simply denotes such a set of weeks. For example, a lot of the uncertainty is related to snow melting in the spring, but most of this uncertainty is related to *when* the inflow will be coming, not its magnitude [52]. Therefore, it might be reasonable to restrict the total deviation in this period. Additionally, restrictions could be added to constrain the *actual* inflow deviations, instead of the normalized ones. Simply using the mapping values F^V_{rt} from Article 1, enables us to limit the number of Mm^3 of deviation to the range $(-\Lambda^L_i, \Lambda^U_i)$. Similarly, these restrictions can be defined for the whole or subsets of the planning horizon. Equivalent restrictions could be expressed for price, but would not have a similarly intuitive interpretation, and we do not have them in the definition of our polytope. Implementing these changes, the final description of the uncertainty polytope becomes

$$-1 \leq \delta_{ut} \leq 1 \qquad u \in \mathbb{U}, \quad t \in \mathcal{T} \tag{6}$$

$$-\Gamma_i^L \leq \sum_{t \in U_i} \delta_{ut} \leq \Gamma_i^U \qquad u \in \mathbb{U}, \; U_i \subseteq \mathcal{T} \tag{7}$$

$$-\Lambda_i^L \leq \sum_{t \in U_i} F_{vt}^V \delta_{vt} \leq \Lambda_i^U \qquad v \in \mathcal{V}, \; U_i \subseteq \mathcal{T}. \tag{8}$$

Note that the bound is not necessarily equal for downward and upward deviation, since the uncertainty distribution could be asymmetrical. Most parameters can only drop to zero, but theoretically increase infinitely. For simplicity of notation, (6) range from -1 to 1, but in reality the upper and lower bounds might be asymmetrical here as well. The slackest bound is set to have magnitude one, and the other is scaled proportionally.

5.4 The uncertainty at hand

In the uncertainty set we employ, all of the parameters described above, Γ, Λ, and F_{vt}^V, are estimated from a set of 3000 scenarios. As each of the parameters represents a measurable size, the bounds on them can be estimated from their realizations in the data.

Scenarios were generated by transforming price and inflow scenarios generated by EMPS into a time series model and simulating scenarios from this. The EMPS model is a fundamental model that describes hydrodominated systems like the Nordic countries well [61]. The price scenarios it generates retain a natural correlation between price and inflow.

Time series parameters for the inflow were estimated from historical data using a semiparametric AR1 model as described in [18]. Then a price model was estimated as a function of local inflow. Parameters were found using panel data methodology with an AR1 error process. With these models, a standardized random error was drawn and used as input to generate the inflow series. The resulting inflow scenarios were passed to the price model, and a corresponding price scenario was generated. The appropriate correlation between error terms was achieved by multiplying the standard random error with the Cholesky decomposed covariance matrix.

The simplest way of generating the polytope from these scenarios is to simply use the extreme values of all scenarios, as mentioned above. To trade robustness for less over-conservative solutions, percentiles of the scenarios can be used. This is illustrated in Figure 6 for inflow deviations of the scenarios. Here, it is seen that removing the 30 most extreme values in every time period, makes the uncertainty set substantially tighter. Still, only 1 % of the observations are removed, so the probability of scenarios outside of the polytope occurring is small.

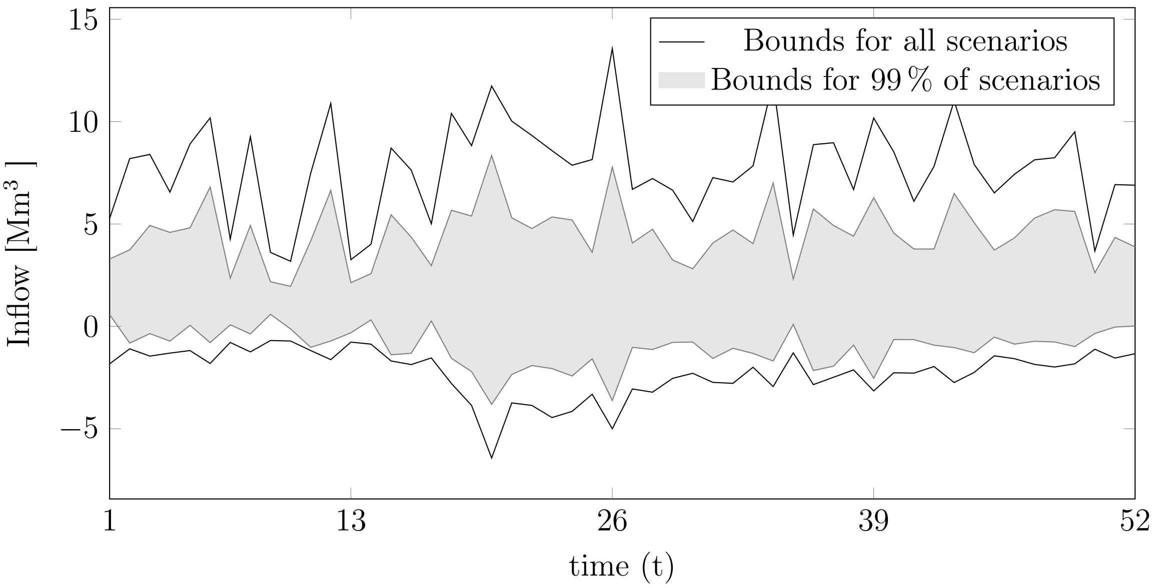

Figure 6: Deviation from inflow forecast in the 3000 scenarios.

This has been a short explanation of the uncertainty set employed in our thesis. The modelling and implementation details of generating the polytope receive closer attention in [40].

5.5 Upper bounds

While LDR have shown promising signs at reducing complexity, this approximation method might lead to a significant loss of optimality [30]. To evaluate the method, then, a measure of the optimality gap ϵ is necessary. Ideally, ϵ could be computed as the difference between the LDR solution's and the optimal solution's value. However, solving the original problem LDR was applied to exactly is potentially very hard [55]. Thus, upper bounds on ϵ are often of interest, to quantify the maximal potential suboptimality.

Such upper bounds could take the form of *a priori* bounds on the maximal value of ϵ for a whole group of instances, or *a posteriori* for a specific problem instance [16]. In the setting of LDR, a priori bounds have only been obtained for a handful of problems, limited to either a two-stage nature [21] or one-dimensional control [6]. A posteriori bounds are therefore more practically attainable for the seasonal planning problem. There are many ways of finding such upper bounds.

In the context of LDR, one possibility is the *dual LDR approximation* [30]. Here, an LDR transformation is applied to the dual version of the original LP. Strong duality ensures that the dual optimum has the same objective value as the primal one. Since the problem in question is a maximization problem, weak duality assures that any feasible approximation of the dual will be an upper bound on the optimal primal solution, and a new non-negative error, ϵ^D, emerges. The joint primal and dual approximation error ($\epsilon + \epsilon^D$) can be calculated as the difference between the primal and dual LDR solutions, and obviously constitutes an upper bound on ϵ. However, calculating this bound requires a whole new LDR approximation, and the quality of these bounds is highly sensitive to the distribution governing the uncertain parameters [23].

Alternatively, by sampling a sufficient number of scenarios, and solving each of them deterministically, an estimate of the expected value given perfect information (EV|PI) is obtained. Obviously, EV|PI is an upper bound on the expected attainable value without certain information, so the gap up to this value constitutes an upper bound on ϵ. Though these *perfect information upper bounds*

(PIUB) are simple to grasp and easy to compute, they often lack the necessary accuracy. Therefore, this thesis strengthens PIUB to the somewhat more promising *dual upper bound* (DUB) [43].

DUB is based on the theory developed in [11] and [24]. The basic principle is to give the solution algorithm access to future information, but punishing its use. The concept of *dual feasible* penalty terms is central, and a proper treatment of the topic is found in [11]. In a maximization problem, such penalty terms are subtracted from the objective function for any action, but should not penalize policies that do not use the extra information.

To briefly introduce the concept, consider a time period t. Let $r_t(\pi_i)$ be the reward or revenue generated by a feasible policy $\pi_i \in \Pi$, and $\mathcal{I}_t$ the information that is naturally known at t. That is, Π is the set of policies depending on $\mathcal{I} = \{\mathcal{I}_1, \mathcal{I}_2, ...\}$. Here, a policy π_i is a vector of actions $\{a_{i1}(\delta), a_{i2}(\delta), ...\}$, one for each time period t. It is said to depend on $\mathcal{I}$ if for any t, $a_{it}(\delta)$ only depends on information available in $\mathcal{I}_t$. Then, introduce an *information relaxation* $\mathcal{I}^R$, a set of more information than would naturally be known, so $\mathcal{I}_t \subset \mathcal{I}_t^R$. This additional information makes a new set Π^R of policies possible, that can potentially depend on future information, so $\Pi \subset \Pi^R$. Now, for any policy $\pi_i \in \Pi^R$ define a penalty $z_t(\pi_i)$, which, as mentioned, should not, in expectation, punish any non-anticipative policies. Thus, this punishment can, and should, be positive for anticipative policies in Π^R, but importantly, $\mathbb{E}[z_t(\pi_i)] \leq 0$ if $\pi_i \in \Pi$. With this simple requirement, we have that for any feasible, non-anticipative policy π_i,

$$\mathbb{E}\left[r_t(\pi_i) \;\middle|\; \mathcal{I}_t\right] \leq \mathbb{E}\left[r_t(\pi_i) - z_t(\pi_i) \;\middle|\; \mathcal{I}_t^R\right] \leq \max_{\pi_j \in \Pi^R} \mathbb{E}\left[r_t(\pi_j) - z_t(\pi_j) \;\middle|\; \mathcal{I}_t^R\right].$$

The first inequality holds because z_t is non-positive in expectation for $\pi_i \in \Pi$. The second holds because the best policy that is allowed to use $\mathcal{I}_t^R$ cannot do worse than a policy using less information. (Or even more fundamentally, π_i is a permittable choice for π_j in the maximization). This is important, because it is now possible to solve

$$\max_{\pi_i} \sum_t \mathbb{E}\left[r_t(\pi_i) - z_t(\pi_i) \;\middle|\; \mathcal{I}_t^R\right], \tag{DUB}$$

in order to get an upper bound for

$$\max_{\pi_i} \sum_t \mathbb{E}\left[r_t(\pi_i) \;\middle|\; \mathcal{I}_t\right].$$

That is, DUB provides an upper bound on the non-anticipative value. The strength of these bounds depend heavily on the penalty functions, but for real option valuation in gas markets, DUB has shown to give fairly tight bounds that are much tighter than PIUB [31]. In Article 2, a specific penalty scheme is implemented, that makes calculation of DUB computationally tractable.

Linear decision rules for seasonal hydropower planning: Implementation and computational performance

Simen V. Braaten, Ola Gjønnes, Knut Hjertvik

June 11, 2015

Abstract

Seasonal hydropower scheduling involves finding the optimal reservoir schedules, while taking uncertainties and non-linearities into account. This paper presents the implementation of a relatively new approach in this field, known as linear decision rules (LDR). In order to solve the problem in a satisfactory fashion, hydropower specific concepts like head effects and water values are taken into consideration through iterative methods. Furthermore, the performance of the LDR model is compared to a rolling intrinsic heuristic (RI), and a novel upper bound to quantify the loss of optimality. The test case used is an interconnected water system with 8 reservoirs, involving uncertainties in price as well as in 4 inflow series. Results show that the LDR model, especially when limited memory is incorporated, provide short computational times compared to RI. However, the gap between the LDR solutions and the upper bound is over 40 %, while RI comes within 4 % of this bound. Therefore it is not clear that LDR are a good fit for the seasonal hydropower planning problem, unless running times are of large importance.

1 Introduction

In this paper a linear decision rules (LDR) formulation for the seasonal hydropower planning problem is tested against both a standard solution algorithm, and theoretical upper bounds. The aim is to present how an LDR method can be used to handle the requirements of hydropower scheduling, and based on this, decide whether LDR are suited for use in seasonal hydropower scheduling. To this end, an alternative approach to solving the same problem is presented. This model is a Rolling Intrinsic (RI) formulation, a fairly standard tool for use in commmodity storage management [34]. The two solution methods are compared in a computational study.

The problem at hand is of interest to planners in the hydropower industry. There are hydropower plants in more than 150 countries around the world, and they account for 16 % of the total electricity generation [18]. Since 1990, global hydropower generation has increased by more than 50 % to 3427 TW h in 2010 [2]. Still, the technically exploitable potential is estimated at more than 16 400 TW h annually, and the steady growth of demand is expected to continue [19]. Thus, ensuring optimal operation is becoming increasingly important, and has major financial and environmental benefits.

Seasonal planning has a horizon of one to five years [37]. The problem is one of the early uses of operations research [35], and solution methods have often been based on stochastic dynamic programming (SDP) [20, 36]. These solutions are trusted by the industry and deliver good results, but suffer from the curse of dimensionality, related to sampling scenarios.

LDR do not need explicit scenario formulations, and the aim of the approach is to create simpler, smaller implementations, to reduce running time. This has been done successfully in other fields [6] such as supply chain management [5] and combined heat and power production [44].

In recent years, several other papers regarding LDR in hydropower scheduling have been published. An iterative procedure that can better handle the non-linear objective functions often arising when variable head of water is included, is presented in [23]. In [31,39], LDR approximations are applied on both the primal and dual of the seasonal hydropower problem, thereby also presenting values for the optimality gap. Compared to [23], the LDR formulation in this paper is applied to a more intricate water system. Additionally, in [23], electricity price is ignored, and output power is maximized, while in this paper, prices are considered stochastic and earnings are maximixed. This paper also builds on [39], but uses zero-mean uncertainty, and a new method for upper bounds. A similar LDR formulation is presented in [44], but this work focuses on application in hydropower scheduling, rather than combined heat and power systems.

The main contribution of this paper is testing the suitability of LDR in hydropower planning. Multiple adjustments are made, so that the LDR formulation handles the problem at hand satisfactorily. Computationally, it is shown that the loss of optimality is significant, but that running times can be much lower than those of competing methods. This means that LDR are not clearly suited for seasonal hydropower scheduling, but can have some value if running time is of the essence.

The rest of this paper is structured as follows. In Section 2 the important hydropower concepts of head and water value are introduced. A brief litterature review is also presented in this section. Then, the theories of LDR and upper bounds are given in Section 3 and 4 respectively. Details regarding the implementation of the head corrections, water values, limited memory, RI and upper bounds are found in Section 5. Lastly, the implementations of the LDR and RI formulations are tested on the 8-reservoir case provided in section 6, and results are presented and analyzed in section 7.

2 Hydropower scheduling

In this section an introduction is given to the problem at hand and its background.

2.1 Problem formulation

The challenge considered involves deciding a seasonal production plan. Given forecasted weekly electricity prices and inflows, profit should be maximized. For each reservoir r, discharge for generation, x^q_{rt}, pumping, x^p_{rt}, and spill, x^s_{rt} should be decided at all times t. For simplicity of notation, the three types of decisions constitute the *decision set* $\mathbb{D} = \{q, p, s\}$. The decisions that are made lead to reservoir levels m_{rt}, and every one of these variables need to be within certain bounds that may change from week to week, but are known on beforehand. Pumping, reservoir level and generation have upper and lower limits, while spill only needs to be non-negative. Last, each reservoir face limits on *total discharge*, i.e. water for generation together with water that bypass the connected power station. All reservoirs are considered to have a power station immediately downstream. Reservoirs that do not are modelled with dummy stations without generation capabilities. With this simplification, power stations can therefore be omitted from the model formulation.

$$\max \quad T^H \sum_{t \in \mathcal{T}} \beta_t \pi_t \sum_{r \in \mathcal{R}} \sum_{d \in \mathbb{D}} E^d_{rt} x^d_{rt} \tag{1}$$

$$\text{s.t.} \quad m_{r0} = M^0_r \qquad r \in \mathcal{R} \tag{2}$$

$$m_{rt} = m_{r(t\text{-}1)} + F_{rt} + T^S \sum_{d \in \mathbb{D}} \left(\sum_{\rho \in \mathbb{C}^d_r} x^d_{\rho t} - x^d_{rt} \right) \qquad r \in \mathcal{R}, t \in \mathcal{T} \tag{3}$$

$$\underline{M}_{rt} \leq m_{rt} \leq \overline{M}_{rt} \qquad r \in \mathcal{R}, t \in \mathcal{T} \tag{4}$$

$$L^q_{rt} \leq E^q_{rt} x^q_{rt} \leq U^q_{rt} \qquad r \in \mathcal{R}, t \in \mathcal{T} \tag{5}$$

$$L^p_{rt} \leq E^p_{rt} x^p_{rt} \leq U^p_{rt} \qquad r \in \mathcal{R}, t \in \mathcal{T} \tag{6}$$

$$\underline{D}_{rt} \leq x^q_{rt} + C^B_{rt} \leq \overline{D}_{rt} \qquad r \in \mathcal{R}, t \in \mathcal{T} \tag{7}$$

$$x^s_{rt} \geq 0 \qquad r \in \mathcal{R}, t \in \mathcal{T} \tag{8}$$

Total profit is maximized in the objective function (1), as the present value of the difference between revenues and costs for all decision types. Here, prices are given per MW h, while decisions are in m^3/s. Energy coefficients E^d are used to convert these volumetric flows to electric power. The energy coefficients are defined to ensure that generation produces electric power, while pumping costs power, and spill does neither. The number of hours in a time period, T^H, is used to get the total energy level in MW h.

Equalities (2) establish initial reservoir levels, while (3) update them between time periods. Decisions made, both in connected upstream reservoirs, $\mathbb{C}^d_r$, and in the current reservoir r, together with natural inflow F_{rt} and the previous level decide the current reservoir level. Constraints (4)-(8) impose upper and lower bounds on reservoir level, generation, pumping, discharge, and spill respectively. Energy coefficients are needed in the generation constraints (5) and pumping constraints (6), as pumping and generation limits are given in MW. It is worth noting that pumping is considered to be negative generation, decided at the upstream reservoir. Thus, inequalities (6) often have an upper bound of zero and a negative lower bound. Since reservoir levels and natural inflows are given in Mm^3, the decisions that are given in m^3/s need to be multiplied by the length of a time period, T^S. Constraints (7) are the total discharge limits, so here bypass (C^B_{rt}) is considered in addition to the water for generation.

Here, all constraints are given as deterministic, but obviously, future prices and inflows are not known a year in advance. The problem should therefore be considered to have uncertain price and inflow. How this is done is explained closer in Section 3.

2.2 Head effects

Maintaining optimal head is one of many considerations when operating a hydropower system [26]. A higher reservoir level gives a higher head and energy coefficient, which means that more energy is generated per discharged unit of water. One of the reasons standard linear programming has seen somewhat limited use in the field is the inherent non-linearity this introduces.

One way of including variable head is by correction coefficients, in the form of head sensitivities in the objective, as in [16, 17]. Another approach is solving the model for specific heads of water, then updating these based on the given reservoir trajectories from the optimization [10, 12]. The first method uses a potentially inaccurate point estimate of the value of head, and the constraints are not affected by the change of generation in the objective. The drawback of the second approach is that updating the energy coefficient parameter *before* reoptimizing, gives the model an incentive to switch between high and low reservoir volumes for every iteration, ultimately leading to convergence issues [16].

In this paper, the two methods are combined, so that the deficiencies of both models are mitigated. A similar approach for including head variations can be seen in [23].

2.3 Water value

One of the primary applications of a seasonal planning model for a power producer, is to obtain water values to use in short-term operations. Finding water values can also be useful within seasonal planning itself, as they can help valuate the water in the end period, thus avoiding unwanted effects such as reservoirs emptying in the final stage. Various heuristics can be used to estimate these values, for example heuristics based on the information in long-term forward contracts [21]. Alternatively, one can use derivatives of the value function in SDP with respect to reservoir level, as in [13]. In LP solutions, the shadow prices of the reservoir balance is a natural estimate of the marginal benefit of additional water [11].

In this paper, water values are found by fixating different storage levels, before optimizing with and without a marginal unit of water, and calculating the difference in objective value. This is because, even though purely linear, the standard reservoir balance equation is not present in the LDR-formulation. As seen in [44], equality constraints are either decomposed or embedded in inequality constraints. However, the completely linear structure makes it possible to store the optimal basis, so reoptimizing after minor changes is very quick.

2.4 Standard approaches

The earliest approaches for solving the seasonal scheduling problem with uncertainty rely on SDP [36, 42]. The approach consists of solving a sequence of subproblems, each linked to a specific time and state, and combining them to full solutions. It does not require linear or even convex problems, but is most applicable for small systems, due to the curse of dimensionality [4]. Many techniques exist to reduce the complexity [3], thus SDP is still an active field of research [27], and much used in reservoir management [38].

One popular way of tackling the curse of dimensionality is stochastic dual dynamic programming (SDDP) [24, 25]. Based on nested Benders' decompostion, SDDP collects optimal dual variables from later stages, and use them to estimate the value in states other than those that are solved. Although probably the most successful method available, SDDP is not easily able to handle jointly stochastic inflows and prices [30], as this can destroy the convexity of the objective function. To overcome this, independence between price and inflow can be assumed [17], but a layer of complexity is still added.

As another approach to reduce complexity, *approximate dynamic programming* [33] relies on simplified value functions or Monte Carlo-simulation [34] to estimate future value. However, since these approaches require sampling scenarios, an accurate distribution of the uncertain parameters might be necessary. Often, historical scenarios will suffice, so this problem is reduced, but real distrubutions are rarely available [23].

In this work, computational tests are run against RI. RI is an iterative approach, based on Monte Carlo simulation and receding horizons. Also known as *model predictive control* or *receding horizon deterministic optimization* [29], it has seen use within hydropower scheduling [43], as well as other power and energy systems facing uncertain forecasts [1]. In the approach of this paper, inspired by [34] in the context of commodity storage, scenarios for the uncertain parameters are sampled. Then, a forecast is generated for the planning horizon, and for each scenario, a deterministic model is solved, using this forecast as certain data. From the solution, the first stage decisions are locked, and the forecasts for the remaining stages are updated. Using these new forecasts, the deterministic model is reoptimized for the rest of the horizon. Repeating this as the remaining horizon decreases toward zero, yields a final solution consisting of the locked solution from each stage. This constitutes one simulation run. To get a good approximation on the optimal value, a large number of simulations are run, and the objective values averaged.

3 Linear decision rules

LDR approximations have emerged as an alternative solution method, that does not rely on discretization of the uncertainty. Rather than solving a problem given certain states or scenarios, recourse decisions are modelled as functions of the uncertain parameters. This way, it is possible to find decisions for any realization of the uncertaininty. If all types of functions are to be considered, finding the optimal recourse function is very difficult [15]. However, only looking at functions that are affine in the parameters leaves the search computationally tractable [7, 22].

In practice, denoting uncertain parameters by δ, this means that a recourse variable x_j can be defined as

$$x_j = \hat{x_j} + \sum_i K_{ij}\delta_i.$$

Here, $\hat{x}$ denotes the intercept, and K_{ij} the slopes of an affine function. As described in [44], substituting these rules into a restriction on the form

$$\sum_j A_j x_j \leq b$$

which should hold for all possible realizations of δ, given by

$$\sum_i H_{ik}\delta_i \leq h_k, \quad \forall k \tag{$\mathcal{U}$}$$

leads to the equivalent formulation, μ_k being the dual variables of the uncertainty set.

$$\sum_k h_k \mu_k \leq b - \sum_j A_j \hat{x}_j \tag{9}$$

$$\sum_k H_{ik}\mu_k = \sum_j A_j K_{ij}, \quad \forall i \tag{10}$$

$$\mu_k \geq 0, \quad \forall k. \tag{11}$$

In this paper, the parameters δ are considered to be normalized deviations from forecasted levels, for both price and inflow. The interested reader can refer to [8] or [44] for a closer introduction to this transformation.

4 Upper bounds

While LDR have shown very promising signs at reducing computational complexity, a considerable loss of optimality is possible [22]. This means that to evaluate the method's usefulness, a measure of the optimality gap ϵ is needed. Clearly, it is not trivial to calculate ϵ directly, as it necessitates the exact solution of the original problem LDR was applied to, which is potentially very hard [40]. Thus, upper bounds on ϵ are often of interest, to quantify the maximal potential suboptimality.

To do this, several methods are applicable. In the context of LDR, an especially interesting choice is the *dual LDR* bound [32]. Here, LDR are applied to the dual of the original problem, and the difference between the solutions is an upper bound on both optimality gaps. Another, simple method is that of perfect information upper bounds (PIUB). A set of scenarios are deterministically solved with perfect information, and the expected value is an upper bound for the optimal objective value with uncertainty. In this article, the PIUB is strengthened by introducing *dual feasible* penalty terms [9] in the objective function. The solution value with these penalties constitutes a potentially tighter upper bound [34], and in Section 5, a specific penalty scheme is introduced that makes the calculation of such a *dual upper bound* (DUB) computationally tractable.

5 Implementation

In this section follows a discussion of the implementation details. Both the model and the testing framework are described, with a focus on handling the challenges mentioned above.

Applying the transformation of (9)-(11) to (2)-(8) and employing expectancy in (1) yields the LDR model used in this paper, with weekly resolution and a 52-week scheduling horizon. It is implemented in IBM ILOG CPLEX Optimization Studio 12.6.1 using IBM OPL on a computer running 64-bit CentOS 6.6, with two AMD Opteron 2431 2.4 GHz processors and 24 GB RAM.

5.1 Variable memory sets

The policies generated by the LDR model are specified by a memory length, i.e. how many periods back in time the decisions can react on. Zero memory implies that the policy only takes the present into consideration, while a memory length of 52, considered as full memory, means that the decisions can react on every previous uncertainty parameter. The memory length can vary across reservoirs, uncertainty series and time periods, e.g. based on topology, known periods of high uncertainty, or reservoir size. Several memory sets have been tested, including simple sets where every reservoir has identical memory length across every uncertainty series and time period, and more complex sets where water course specific information is utilized. Two such sets are the *perfect memory* set, and the *flood memory* set.

The perfect memory set is not a feasible implementation strategy, but rather a benchmark test. Here, analysis is done on the results of running with full memory. The perfect memory set consists of exactly those slopes that take on a value in this run. That is, the perfect memory set contains those dependencies that the model chooses, given the possibility to use any slopes. Of course, the information about which slopes to choose would not be available to the planner ahead of time. Thus, this is not a fair set to compare against, but it reveals an interesting property of LDR. By keeping only about 5% of the available slopes, the same objective value can theoretically be achieved.

The flood memory set is based on the idea that some periods are more important than others. Specifically, much of the uncertainty lays in the spring flood, which is often based around week 20 to 33 in the test case presented in 6. Thus, the flood memory is a modified version of memory one. In addition to having memory one, every period in the spring flood is dependent on every earlier period in the spring flood. On top of this, the last three periods are also dependent on the spring flood. This way, the most critical periods get flexibility from knowing the outcome of the most uncertain periods.

5.2 Head effects

When solving the LDR-model, the energy coefficients have to be considered constant to keep the formulation linear. Thus, iterative reoptimizations and updates of these coefficients are made, to include the dynamics of variable head. The model is initially solved with energy coefficients corresponding to reference head levels. After optimization, updated head levels are calculated from the expected reservoir levels, and the coefficients are updated accordingly, as is done in [12]. Pseudocode for the implementation of this procedure is described in Algorithm 1. Without any other changes the solutions will diverge, as a greater coefficient motivates higher generation, which in turn leads to a lower reservoir level the next iteration. It is thus important to give the model an incentive to try to obtain a beneficial head level. In this implementation, the incentive is given by the term presented in [16, 17] in the objective. The term is a linearization of the marginal increase in the objective given a unit of extra head. Thus, it rewards high storage levels for reservoirs with generation immediately downstream, and penalizes high storage levels for reservoirs with generation immediately upstream. For pumping, this is reversed.

Algorithm 1 Update energy coefficient

input: $\mathcal{T}$ - set of time periods
$\mathcal{R}$ - set of reservoirs
m_{rt} - expected volume level in reservoir r at end of time period t
s_r - reservoir immediately downstream r
w_r - tailwater level of reservoir r
$\overline{h_r}$ - reference head of reservoir r

procedure UPDATE ENERGY COEFFICIENTS
  for all $t \in \mathcal{T}$ **do**
    for all $r \in \mathcal{R}$ **do**
      $z_{rt} \leftarrow \mathtt{GetSurfaceHeight}(m_{rt})$ # Get meters above sea level from reservoir volume
    for all $r \in \mathcal{R}$ **do** # Update the head level
      $h_{rt} \leftarrow z_{rt} - \mathtt{max}(z_{s_r t}, w_r)$
      $h'_{rt} \leftarrow \mathtt{mean}(h_{r(t-1)}, h_{rt})$ # The average head of time period t
      $E_{rt} \leftarrow E_{r0} \cdot (h'_{rt}/\overline{h_r})$
  return

5.3 Water values

Water values as functions of the reservoir level are important outputs of a seasonal hydro power scheduling model. Here, these functions are estimated based on multiple sampled points consisting of water values and reservoir levels for a given time period, i.e states, as presented in Algorthim 2. To retrieve the values for different states, the regulated water levels are constrained. Depending on the regulated levels, the uncertainty in inflow might now lead to an infeasible problem. This issue is avoided by changing the uncertainty set so that the change in water in the reservoir is accounted for by the uncertainty. Additionally, the slopes reacting on the now constrained δ have to be forced to zero. If not, the problem is unbounded, as these slopes now have the possibility to counterbalance an infinitely large production intercept. The model is then optimized with and without a marginal unit of water in the reservoirs, by changing the upper bounds of the transformed version of (4). Finally, the water values at the given scenario are obtained by taking the difference between the objective values before and after the marginal change. Every reoptimization of the LP is done very quickly, as CPLEX warm starts with the previous optimal basis.

Algorithm 2 Calculate water values

input: `model` - the LDR model
$\mathcal{S}$ - set of states where water values are wanted
$\mathcal{R}$ - set of reservoirs
m_{rs} - desired level for reservoir r in state s
ϵ - marginal unit of water

procedure CALCULATE WATER VALUES
  for all $s \in \mathcal{S}$ **do**
    Restrict δ such that it is in compliance with water levels in s
    for all $r \in \mathcal{R}$ **do**
      $(\overline{M}_{r1}, \underline{M}_{r1}) \leftarrow (m_{rs}, m_{rs})$ # Restrict reservoir levels
      `model.Slope`(r, δ)`.`$Limits \leftarrow 0$ # Prevent relevant slopes from taking value
    $z_1 \leftarrow$ `model.`$solve()$
    for all $r \in \mathcal{R}$ **do**
      `model.ReservoirBalance`(r)`.Bounds.`$Add(\epsilon)$ # Add marginal unit
      $z_2 \leftarrow$ `model.`$solve()$
      `model.ReservoirBalance`(r)`.Bounds.`$Remove(\epsilon)$ # Reset bounds
      $w_{rs} \leftarrow (z_2 - z_1)/\epsilon$ # Water value reservoir r in state s
    `model.`$Reset()$ # Reset changes done to model
  return w

5.4 Rolling intrinsic heuristic

The RI heuristic is implemented in Matlab r2014a. A deterministic model is built, taking the number of periods remaining, future inflows, and prices as input. Initial forecasts are generated as the mean of 3000 scenarios based on EMPS [41], and multiple update methods are available. Deviations from the forecast could be used as indicators of trend, such that stretches of low inflow leads to pessimistic forecasts. Symmetrically, deviations could be compensated for, so low inflow leads to higher forecasts going forward [13]. The simplest option is staying with the mean throughout the horizon, only updating with correct information for the current week. This forecasting method is chosen in the tests since it is widely used by practitioners. The specifics of the procedure are outlined in Algorithm 3.

A known weakness of RI is that the locked solutions may lead to infeasibilities at later stages, when realized inflow differs from expectations [28]. To counteract this, the RI problem is given buffered bounds to avoid going too close to the actual bounds. When the buffered problem turns out to be infeasible, the buffers are removed, and the model is solved with its actual bounds. Additionally, to give the algorithm a conservative view of the future, the forecast is adjusted to its worst case for the last 5 weeks in the horizon. This way, RI proceeds with more caution, and rarely makes choices leaving future periods infeasible (one or more periods are infeasible in only ~3‰ of all scenarios).

Algorithm 3 Rolling Intrinsic

input: `DetProb`(T', P, I, b) - the deterministic problem, given remaining time periods, price, inflow, and constraint buffers.
`Forecast`$(t_i : t_j)$ - a way to forecast inflow and price for periods i to j
T - number of time periods,
N - number of Monte Carlo simulations,
b - buffer values for restrictions.

procedure RI
  for $i = 1 : N$ **do**
    $P, I \leftarrow$ generated price and inflow scenario
    for $t = 1 : T$ **do**
      $\hat{P}_t, \hat{I}_t \leftarrow P_t, I_t$ # Information of this period is known
      $\hat{P}_{t+1:T}, \hat{I}_{t+1:T} \leftarrow$ `Forecast`$(t+1 : T)$ # Information for rest of horizon is forecasted
      $x \leftarrow$Solve `DetProb`$(T+1-t, \hat{P}, \hat{I}, b)$ # Solve with buffers
      if x is infeasible **then**
        $x \leftarrow$Solve `DetProb`$(T+1-t, \hat{P}, \hat{I}, 0)$ # Solve without buffers
      $x_t^{lock} \leftarrow x_t$ # Lock solution of current period
    $z \leftarrow$`Simulate`(x^{lock}, P, I) # Evaluate locked solution under realized price and inflow
    `MeanResult` $\leftarrow$ `UpdateMean`(z) # The average objective value
  return `MeanResult`

5.5 Upper bounds

Both the PIUB and DUB approaches benefit from the same deterministic model that is used in RI. When calculating PIUB, each scenario is deterministically solved with perfect information for the full planning horizon. The mean objective value constitutes the bound.

For DUB, optimal policies are calculated by using a rolling horizon, similar to RI. However, DUB should also allow anticipative policies, although with a penalty. The idea behind the particular scheme utilized here is to use the perfect information relaxation. That is, the problem is solved with all information available, but the algorithm is penalized for using future price information. As can be seen in Algorithm 4, the method bears close resemblance to RI, but operates with perfect inflow information. Intuitively, this is an upper bound, since the planner knows the future inflow. Technically, it implements a scheme to penalize marginal net production due to better price information. The objective function can be expressed as $P \cdot (q - p)$, where P is price, q is generated energy, and p is energy consumed for pumping. Now, define the penalty term

$$(P - \hat{P}) \cdot ((q - \hat{q}) - (p - \hat{p})). \tag{12}$$

Here, $\hat{P}$ is the forecasted price, while $\hat{q}$ and $\hat{p}$ are the best solutions not using future price information. Such a scheme will only penalize the marginal net production $(q - \hat{q})$ that is scheduled based on the knowledge that realized price exceeds expectation (and marginal net reduction or pumping $(p - \hat{p})$ when realized price is low). Thus, in expectation, it should not punish any policies that do not use future price information. Since perfect inflow information is used all the way from the beginning, the problem will never become infeasible. This means that the penalties are dual feasible, and the solutions are primal feasible, so DUB can be found using the following maximization:

$$
\begin{aligned}
& \max_{q,p} P \cdot (q-p) - (P-\hat{P}) \cdot ((q-\hat{q}) - (p-\hat{p})) && (13)\\
= \; & \max_{q,p} P \cdot (q-p) - P \cdot (q-p) + \hat{P} \cdot (q-p) + P \cdot (\hat{q}-\hat{p}) - \hat{P} \cdot (\hat{q}-\hat{p}) && (14)\\
= \; & \max_{q,p} \hat{P} \cdot (q-p) + (P-\hat{P}) \cdot (\hat{q}-\hat{p}). && (15)
\end{aligned}
$$

Since $\hat{q}$ and $\hat{p}$ are independent of this problem, the second term of (15) is constant with respect to the decision variables q and p. By this, the solution (q, p) to (13) should coinicide with

$$\operatorname*{argmax}_{q,p} \hat{P} \cdot (q-p), \tag{16}$$

which is exactly what $\hat{q}$ and $\hat{p}$ are defined to be. Thus, the optimal value of (15) will be

$$
\begin{aligned}
& \hat{P} \cdot (\hat{q}-\hat{p}) + (P-\hat{P}) \cdot (\hat{q}-\hat{p}) && (17)\\
= \; & P \cdot (\hat{q}-\hat{p}), && (18)
\end{aligned}
$$

which simply amounts to solving the problem with forecasted prices, and then evaluating the solution using realized prices. This is what Algorithm 4 does, and what constitutes the dual upper bound on the optimal solution value.

Algorithm 4 Dual Upper Bound

input: $\texttt{DetProb}(T', P, I)$ - the deterministic problem, given remaining time periods, price, and inflow
$\texttt{Forecast}(t_i : t_j)$ - a way to forecast price for periods i to j
T - number of time periods,
N - number of Monte Carlo simulations.

procedure DUB
  for $i = 1 : N$ **do**
    $P, I \leftarrow$ generated price and inflow scenario
    for $t = 1 : T$ **do**
      $\hat{P}_t \leftarrow P_t$ # Price information of this period is known
      $\hat{P}_{t+1:T} \leftarrow \texttt{Forecast}(t+1 : T)$ # Price information for rest of horizon is forecasted
      $x \leftarrow$Solve $\texttt{DetProb}(T+1-t, \hat{P}, I)$ # Solve with perfect inflow information
      $x_t^{lock} \leftarrow x_t$ # Lock solution of current period
    $z \leftarrow \texttt{Simulate}(x^{lock}, P, I)$ # Evaluate locked solution under realized price
    $\texttt{MeanResult} \leftarrow \texttt{UpdateMean}(z)$ # The average objective value
  return $\texttt{MeanResult}$

6 Example case

Fantasi is a hydropower plant developed by Powel AS as a testing and benchmarking case. Its topology is illustrated in Figure 1, and it is specifically designed to contain many of the most

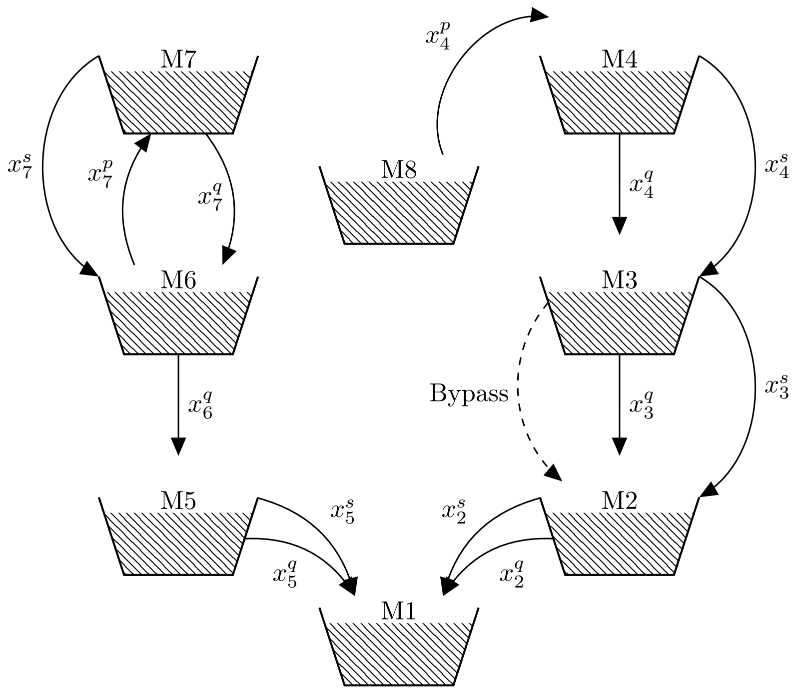

Figure 1: An illustration of the water course "Fantasi".

common components of a water course. Examples include a large spread in reservoir sizes, multiple inflow series, a pumped storage facility, and a pure pumping station. Additionally, M3 has a fixed minimal bypass, and some reservoirs spill to other reservoirs, while some spill to the ocean.

Key data points on the example case are given in Table 1, and both the expected price and standard deviation in inflow is presented in Figure 2. The latter is obtained by taking the average of the standard deviations of the different inflow series, and is included to depict the development in the uncertainty over time. In addition, all reservoirs are required to start and end at 66.8 % of their capacity, an average degree of filling in January.

7 Computational results and conclusion

In this section, the main results and findings from running LDR on the example case are presented. Finally, the performance of LDR is evaluated in comparison to RI and DUB.

Table 1: Key characteristics of the water course Fantasi

Reservoir	Inflow series	Capactiy [Mm^3]	Average storeable inflow [Mm^3/year]
M1	4	50	150
M2	3	50	200
M3	1	500	1000
M4	1	400	500
M5	2	150	300
M6	2	100	700
M7	1	300	300
M8	1	10	260

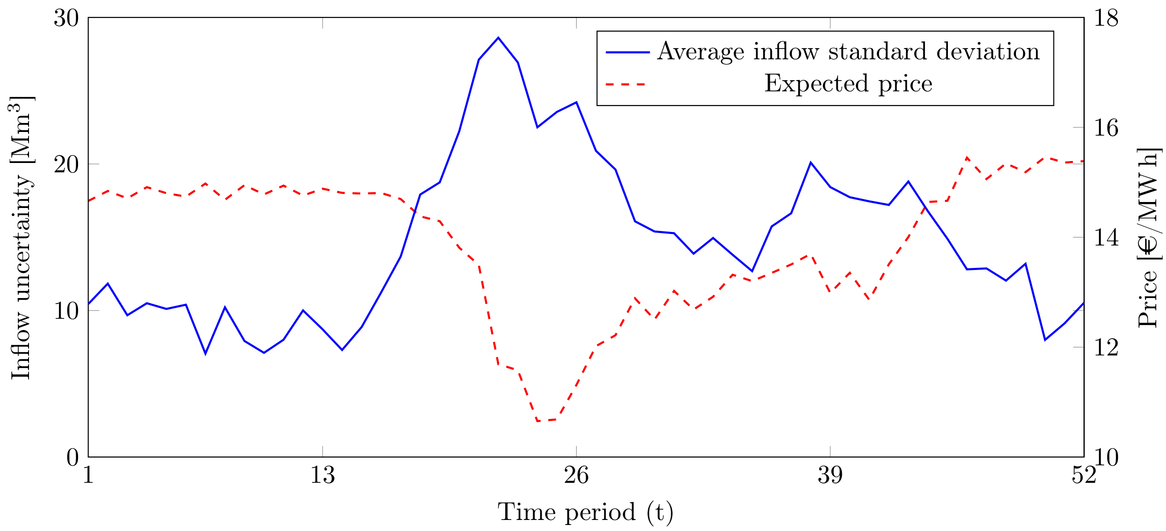

Figure 2: The expected price plotted with the standard deviation of inflow

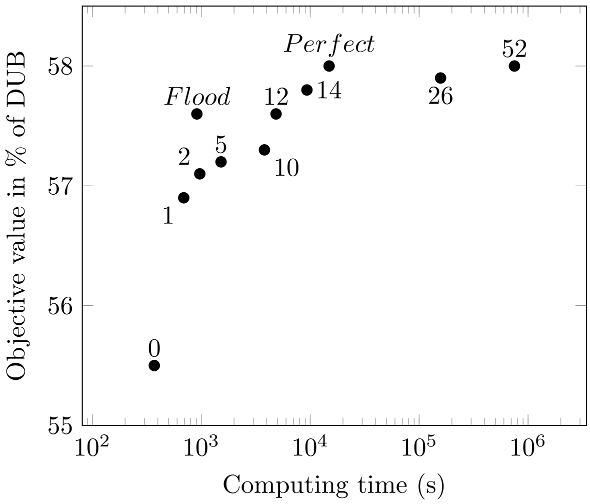

Figure 3: The computational times of different memory sets plotted on a logarithmic scale against the objective value in % of DUB

7.1 Memory sets

The LDR model is solved with a number of different memory sets, among them the flood memory and perfect memory set. The results from these test are presented in Figure 3, where computational times are plotted on a logarithmic scale against the objective value. It is evident that larger memory sets lead to significantly higher complexity in the problem. This is not surprising, as each slope is a variable, and the number of variables drives solution time in Simplex.

When it comes to the objective value, the biggest gap is between the 0-memory and the 1-memory solution. This is sensible, as it should be valuable to be able to take the past into account at all, relatively to only reacting to the present. The relative differences between the memory lengths decrease slightly towards the 10-memory solution. Somewhat surprisingly, an increase in the relative differences between the memory lengths can be observed from the 10-memory to the 14-memory implementation. This could be because the snow melting period, in which the uncertainty is highest, is approximately 14 weeks long, around weeks 16-30 in the spring and summer (see Figure 2). Taking this whole period into account might add significant flexibility to the solution, yielding less conservative plans. From memory length 14 up to full memory at 52 weeks, the relative improvements are once again minor, suggesting that beyond this, there is no large benefit in very long memory lengths.

The optimal slope values from the 52-memory solution, illustrated in Figures 4 and 5, support these arguments. Here, it is clear that the slopes closest to the main diagonal are the most important ones. Most periods are only dependent on their own uncertainty, and at most the last few weeks. The notable exceptions are periods close to the end of the spring flood, and at the end of the horizon. In these periods, many previous weeks are taken into consideration, to accurately react to total inflow. At the end of the spring flood, the reservoirs are usually close to full, making it important to avoid

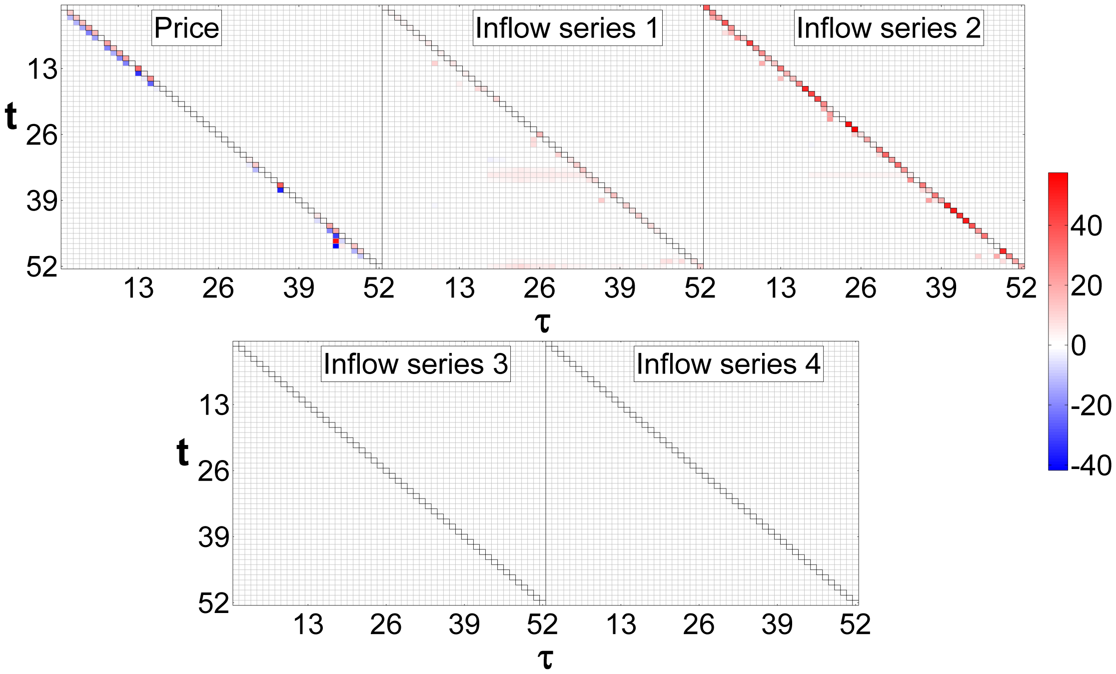

Figure 4: Optimal slope matrix for discharge from reservoir M6

spilling. When coming close to the end of the planning horizon, the reservoirs levels should be as close as possible to final lower bounds, to optimally utilize the water available.

Another interesting and intuitive observation, is that the decisions at a reservoir r only depend on the uncertainty in the reservoir's own inflow series, and the series at the reservoirs upstream to r, in addition to the price. Consider the slopes of reservoir M6 presented in Figure 4. These slopes depend on the realization of the price, the first, and the second inflow series. Reservoir M7, the only reservoir upstream to M6, belong to the first inflow series, while the second inflow series is M6's own series. In contrast, study the plot for the bottom reservoir M1, found in Figure 5. The decisions of this reservoir are clearly dependent upon every uncertainty series in the model. This is a natural result, as the total inflow to a reservoir is dependent upon the preceeding reservoirs, which naturally react to their own inflow series. Thus, as long as the decisions can be made based on what has happened upstream, robustness can be achieved. This could be the basis of further memory reductions, removing dependencies to redundant inflow series.

7.2 Head effects

The iterative modifications presented have an effect on the reservoir trajectories and objective values. In Figures 6 and 7, trajectories of two reservoirs are shown. M6 is one of the smaller reservoirs in the water course, while M3 is the largest. Both of the trajectories seem suboptimal before any corrections are made. Notably, M6 discharges a large amount of water very early, leaving it with

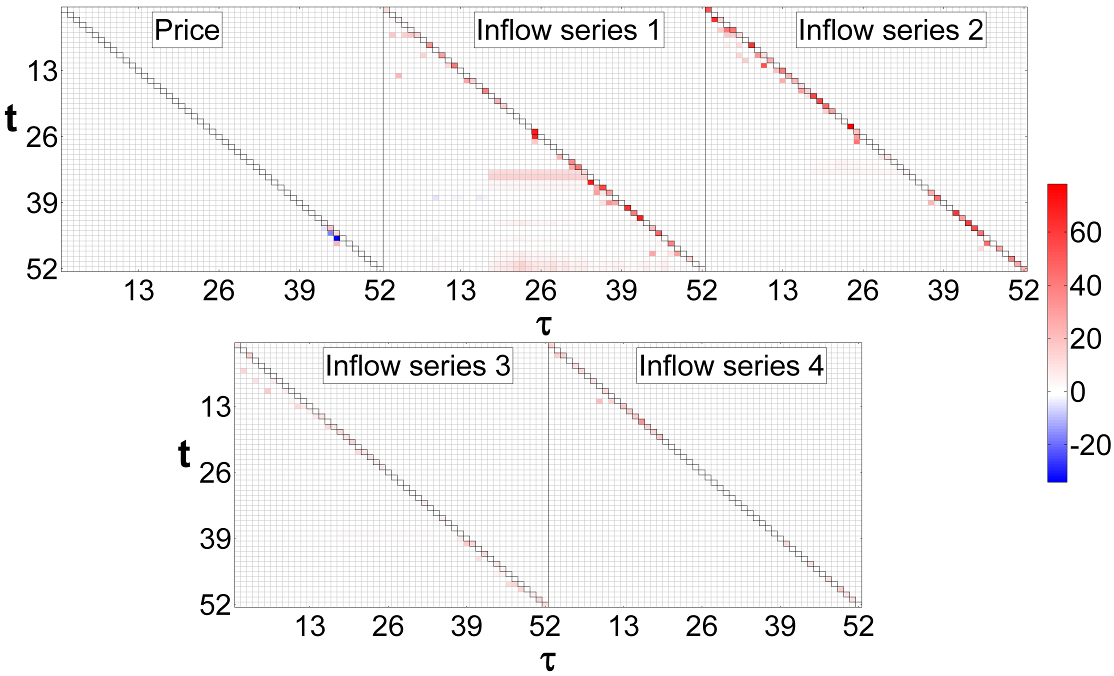

Figure 5: Optimal slope matrix for discharge from the bottom reservoir M1

a low head during the profitable winter months, while M3 stops its filling period with a negative spike around period 25.

After adding the head value term in the objective function, and iterating until convergence, a different pattern emerges. Now M6 maintains an almost maximal head through the three first months, while M3 completes filling up without interruption. This is evidence that the LDR model is able to take non-linearities into account if handled correctly, and makes it more usable. A planner would be able to trust more in a policy that takes head variations into account, modelling the reality more closely.

7.3 Water values

Another iterative procedure, namely the water value calculation, also seems to provide sensible results. This is important if considering to use output from the model in shorter term scheduling models. As seen in Figure 8, the water value curves look reasonable. In periods when inflow is high, an extra unit of water is worth less than in periods where inflow is low. Similarly, when the price is high in the winter, and the reservoirs are slowly emptying, additional water is valuable. Here, the point value of the water at the mean reservoir levels is shown, but the water value at each reservoir is actually a multidimensional function of reservoir levels at all reservoirs in the whole system.

7.4 The problem with linearity

One of the areas where the LDR model seemingly produces irrational results is in the relationship between power generation and pumping. In several time periods the model suggests both pumping and generating power at the same reservoir. Given a constant price throughout the period, and that pumping consumes more power than discharging generates, this is not something a real production planner would do.

As seen in Figure 9, the result is natural when considering how LDR are defined. The uncertain parameters can go below their expected values, in which case a positive slope will lead to decisions that are below their nominal levels. For this to still be feasible, the decisions need to have intercepts above their lower bounds. The consequence of this, is that if there is a time period where both pumping and power generation have slopes, then both needs intercepts as well. Often it is a good idea to have non-zero slopes for both these decisions, in order to fully exploit scenarios with extreme prices or inflows. In this case, the symmetry of reactions gives a situation where almost any realization of the uncertainty will give rise to both pumping and power generation in the same time period. If the LDR model is used as decision support, the planner should use the net decision in these cases.

To mitigate this problem, alternatives like piecewise linear [14] or quadratic [32] decision rules can be investigated. Both of these alternatives allow more asymmetric reactions to positive and negative realizations.

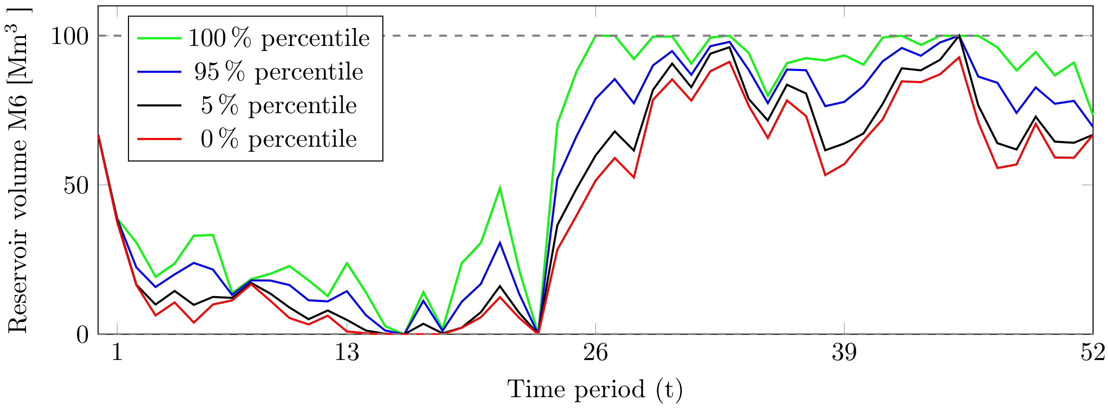

(a) Reservoir trajectories M6 prior to head corrections

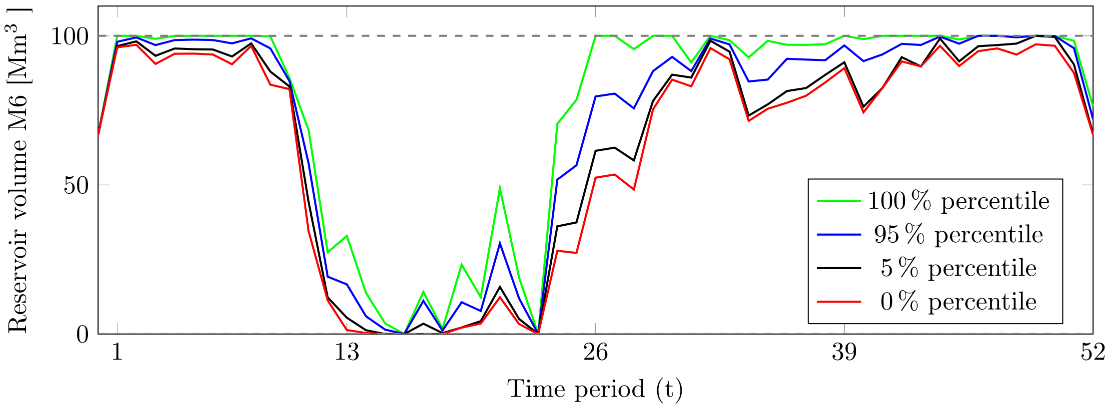

(b) Reservoir trajectories M6 after head corrections

Figure 6: Percentile trajectories for M6 obtained from simulating policies from LDR with five period memory

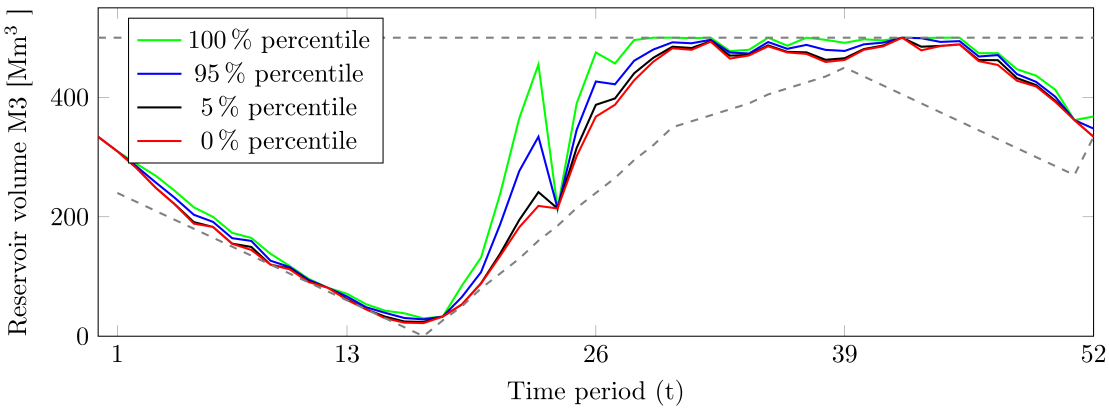

(a) Reservoir trajectories M3 prior to head corrections

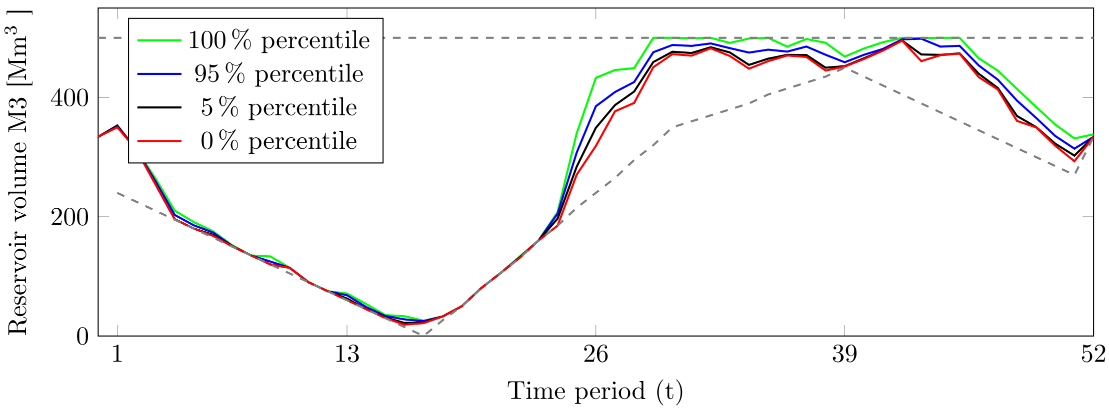

(b) Reservoir trajectories M3 after head corrections

Figure 7: Percentile trajectories for M3 obtained from simulating policies from LDR with five period memory

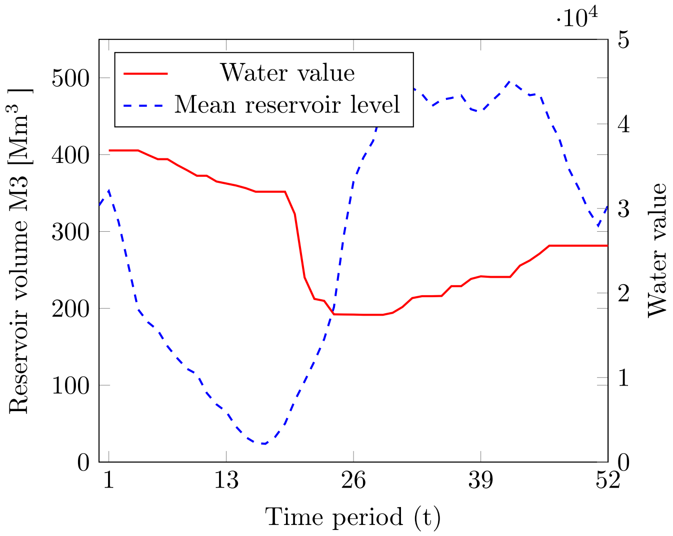

(a) Water value vs mean reservoir level M3

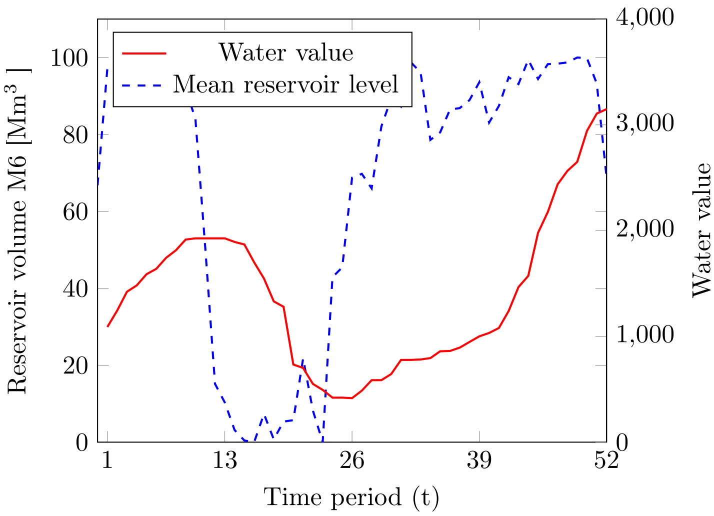

(b) Water value vs mean reservoir level M6

Figure 8: Comparison of water value and mean reservoir level for head corrected runs

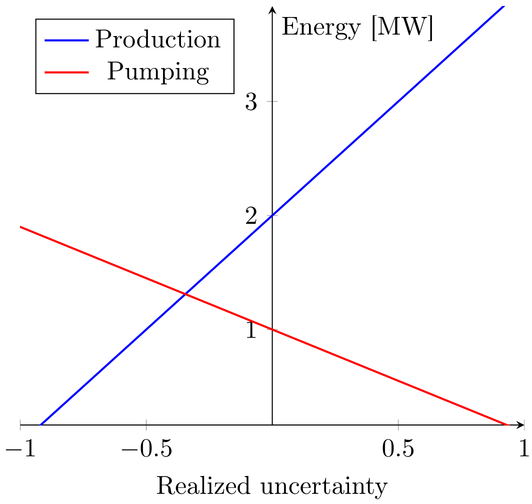

Figure 9: Generation and pumping plotted against the uncertainty parameter. Pumping is plotted as positive for the sake of understanding.

7.5 Comparison

The additions that are made to the LDR implementation have shown to have their intended effects. Now, to evaluate the usefulness of LDR in the context of seasonal hydropower planning, it is compared to RI. Both solution methods are benchmarked against the DUB described in Section 4, and the results are presented in Table 2 and Figure 10.

Here, it is clear that RI outperforms LDR. RI comes within 4 % of a theoretical upper bound, while the gap to LDR is more than 40 %. At full memory, the LDR model also takes much longer to compute, leaving RI the dominant solution. The running times of RI are comparable to what is reported in [17], so it can be assumed to be on par with the standard tools in use.

The time savings of LDR are evident when appropriate memory restrictions are applied. Relative to RI, the difference between full memory and no memory in LDR is small, at just a few percentage points. This means that LDR can be run with very restrictive memory with no significant loss of optimality. The flood memory suggested in Section 5.1 performs comparably with full memory at under 0.1 % of the running time. This supports the idea that if the spring flood is taken into account, periods of lower uncertainty do not need to be considered as closely, and indicates that variable memory lengths may be designed, that perform very well with shorter running times. The flood memory also clearly beats the running time of RI, although it is not evident that the reduced running time justifies the large optimality gap.

7.6 Conclusion

From the computational results, it is not clear that LDR is a good fit for the seasonal hydropower planning problem, even though the method potentially delivers some benefits. Examples of this

Table 2: Comparison of objective values and computational times from the different models and memory sets

Model	Simulated objective value (% of DUB)	Computational time (s)
PIUB	136 981 (102)	1 406
DUB	134 240 (100)	32 044
RI	128 830 (96,0)	32 310
LDR perfect memory	77 867 (58,0)	14 928
LDR flood memory	77 203 (57,5)	913
LDR 52 memory	77 867 (58,0)	750 154
LDR 26 memory	77 777 (57,9)	157 484
LDR 14 memory	77 542 (57,8)	9 337
LDR 10 memory	76 915 (57,3)	3 804
LDR 5 memory	76 838 (57,2)	1 520
LDR 2 memory	76 646 (57,1)	970
LDR 1 memory	76 375 (56,9)	691
LDR 0 memory	74 523 (55,5)	371

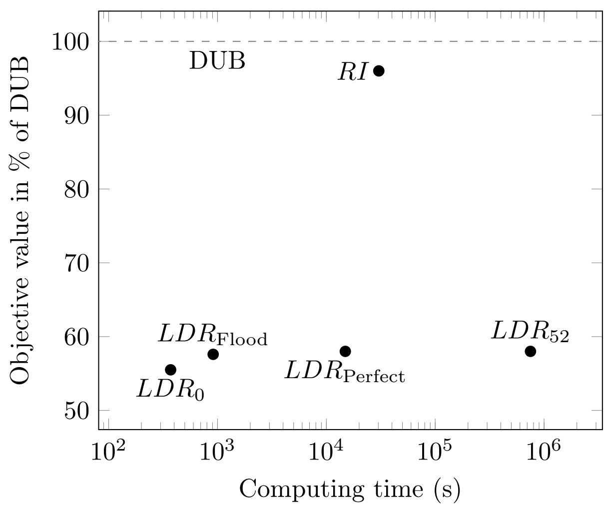

Figure 10: The computational times plotted on a logarithmic scale against the objective value in % of DUB

include coherent policies for any scenario, but also intuitively reasonable solutions, and potentially short running times compared to RI. Especially if efficient reductions of the memory can be developed, LDR outruns RI by orders of magnitude. However, this comes at a significant cost of optimality. Even with full memory, the LDR are not able to close the gap up to RI, leaving the difference at more than one third of the optimal value. This means that LDR cannot be considered a good replacement, or a good fit for seasonal planning in hydropower scheduling. Rather, LDR could be used as a quick diagnostic tool, or when a planner wishes to iteratively develop and test scenarios, or do analysis on a large number of potential scenarios. Here, exact solutions need not be of high importance, while quickly assessing results of changes might be interesting.

The LDR formulation was shown to handle head effects well using an iterative approach, and deliver reasonable water values, which means it *can* be used to quickly produce output for other models to build upon.

Based on the results presented in this paper, further research is recommended into the reduction of memory sets in LDR. Since LDR do not seem fit to compete as a primary seasonal planning tool, attempting to get the quickest solution times possible might add value to the approach. Additionally, the water values output from the LDR model seemed reasonable. Further comparison of these against those returned by other methods is potentially interesting. If the marginal values of water are shown to be correct, they can be used as a quick way to input water values to a shorter term model. Last, the symmetry of purely linear decision rules seemed to give irrational results. More research on piecewise linear or quadratic decision rules could yield more fruitful policies.

8 Acknowledgements

The authors would like to thank supervisors Gro Klæboe and Stein-Erik Fleten for invaluable help and input during this work.

References

[1] M. Arnold and G. Andersson. Model predictive control of energy storage including uncertain forecasts. In *Power Systems Computation Conference (PSCC), Stockholm, Sweden*, 2011.

[2] A. M. Bagher, M. Vahid, M. Mohsen, and D. Parvin. Hydroelectric energy advantages and disadvantages. *American Journal of Energy Science*, 2(2):17–20, 2015.

[3] J. C. Bean, J. R. Birge, and R. L. Smith. Aggregation in dynamic programming. *Operations Research*, 35(2):215–220, 1987.

[4] R. Bellman. *Dynamic Programming*, volume 1. Princeton University Press, 1957.

[5] A. Ben-Tal, B. Golany, A. Nemirovski, and J.-P. Vial. Retailer-supplier flexible commitments contracts: A robust optimization approach. *Manufacturing & Service Operations Management*, 7(3):248–271, 2005.

[6] A. Ben-Tal, A. Goryashko, E. Guslitzer, and A. Nemirovski. Adjustable robust solutions of uncertain linear programs. *Mathematical Programming*, 99(2):351–376, 2004.

[7] A. Ben-Tal and A. Nemirovski. Robust convex optimization. *Mathematics of Operations Research*, 23(4):769–805, 1998.

[8] S. V. Braaten, O. Gjønnes, and K. Hjertvik. Linear decision rules for seasonal hydropower planning: Modelling considerations. *Working paper, Norwegian University of Science and Technology, Trondheim*, jun 2015.

[9] D. B. Brown, J. E. Smith, and P. Sun. Information relaxations and duality in stochastic dynamic programs. *Operations Research*, 58(4):785–801, 2010.

[10] D. De Ladurantaye, M. Gendreau, and J.-Y. Potvin. Strategic bidding for price-taker hydroelectricity producers. *IEEE Transactions on Power Systems*, 22(4):2187–2203, 2007.

[11] F. R. Førsund. *Hydropower Economics*, chapter 3. Springer Science & Business Media, 2 edition, 2015.

[12] O. B. Fosso and M. M. Belsnes. Short-term hydro scheduling in a liberalized power system. In *Proceedings of the International Conference on Power System Technology , 21-24 November*. IEEE, 2004.

[13] O. B. Fosso, A. Gjelsvik, A. Haugstad, B. Mo, and I. Wangensteen. Generation scheduling in a deregulated system. The Norwegian case. *IEEE Transactions on Power Systems*, 14(1):75–81, 1999.

[14] A. Georghiou, W. Wiesemann, and D. Kuhn. The decision rule approach to optimisation under uncertainty: Methodology and applications in operations management. working paper, 2011, 2011.

[15] A. Georghiou, W. Wiesemann, and D. Kuhn. Generalized decision rule approximations for stochastic programming via liftings. *Mathematical Programming*, early access:1–38, 2014.

[16] A. Gjelsvik and A. Haugstad. Considering head variations in a linear model for optimal hydro scheduling. Unpublished paper from SINTEF Energy Research, Trondheim, Norway, 2005.

[17] A. Gjelsvik, B. Mo, and A. Haugstad. 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*, pages 33–55. Springer, 2010.

[18] International Energy Agency. IEA report sets a course for doubling hydroelectricity output by 2050. Available at: https://www.iea.org/topics/renewables/subtopics/hydropower/, april 2015.

[19] International Energy Agency. Renewable energy essentials: Hydropower. Available at: http://www.iea.org/publications/freepublications/publication/hydropower_essentials, mar 2015.

[20] J. Kelman, J. R. Stedinger, L. A. Cooper, E. Hsu, and S.-Q. Yuan. Sampling stochastic dynamic programming applied to reservoir operation. *Water Resources Research*, 26(3):447–454, 1990.

[21] T. K. Kristoffersen and S.-E. Fleten. *Energy, Natural Resources and Environmental Economics*, chapter 4 in part II, pages 187–201. Springer Science & Business Media, 2010.

[22] D. Kuhn, W. Wiesemann, and A. Georghiou. Primal and dual linear decision rules in stochastic and robust optimization. *Mathematical Programming*, 130(1):177–209, 2011.

[23] L. Pan, M. Housh, P. Liu, X. Cai, and X. Chen. Robust stochastic optimization for reservoir operation. *Water Resources Research*, 51:409–429, 2015.

[24] M. Pereira. Optimal stochastic operations scheduling of large hydroelectric systems. *International Journal of Electrical Power & Energy Systems*, 11(3):161–169, 1989.

[25] M. Pereira and L. M. Pinto. Multi-stage stochastic optimization applied to energy planning. *Mathematical Programming*, 52(1-3):359–375, 1991.

[26] R. Ponrajah, J. Witherspoon, and F. Galiana. Systems to optimise conversion efficiencies at Ontario Hydro's hydroelectric plants. *IEEE Transactions on Power Systems*, 13(3):1044–1050, 1998.

[27] M. L. Puterman. *Markov decision processes: discrete stochastic dynamic programming*. John Wiley & Sons, 2014.

[28] S. J. Qin and T. A. Badgwell. A survey of industrial model predictive control technology. *Control engineering practice*, 11(7):733–764, 2003.

[29] J. B. Rawlings and K. R. Muske. The stability of constrained receding horizon control. *IEEE Transactions on Automatic Control*, 38(10):1512–1516, 1993.

[30] S. Rebennack, N. A. Iliadis, M. V. Pereira, and P. M. Pardalos. Electricity and CO_2 emissions system prices modeling. In *IEEE PowerTech, Bucharest, Romania*, pages 1–6, 2009.

[31] P. Rocha. *Medium-term planning in deregulated energy markets with decision rules*. PhD thesis, Imperial College London, 2013.

[32] P. Rocha and D. Kuhn. A polynomial-time solution scheme for quadratic stochastic programs. *Journal of Optimization Theory and Applications*, 158(2):576–589, 2013.

[33] D. F. Salas and W. B. Powell. Benchmarking a scalable approximate dynamic programming algorithm for stochastic control of multidimensional energy storage Problems. Technical report, Department of Operations Research and Financial Engineering, Princeton, NJ, 2013.

[34] N. Secomandi. Analysis and enhancement of practice-based policies for the real option management of commodity storage assets. Working paper, Carnegie Mellon University, Pittsburgh, 2014.

[35] S. Stage and Y. Larsson. Incremental cost of water power. *Power Apparatus and Systems, part iii. Transactions of the American Institute of Electrical Engineers*, 80(3):361–364, 1961.

[36] J. R. Stedinger, B. F. Sule, and D. P. Loucks. Stochastic dynamic programming models for reservoir operation optimization. *Water Resources Research*, 20(11):1499–1505, 1984.

[37] G. Steeger, L. A. Barroso, and S. Rebennack. Optimal bidding strategies for hydro-electric producers: a literature survey. *IEEE Transactions on Power Systems*, 29(4):1758–1766, 2014.

[38] J. A. Tejada-Guibert, S. A. Johnson, and J. R. Stedinger. The value of hydrologic information in stochastic dynamic programming models of a multireservoir system. *Water Resources Research*, 31(10):2571–2579, 1995.

[39] I. A. Utheim Grønvik, A. Hadziomerovic, N. Ingvoldstad, R. Egging, and S.-E. Fleten. Feasibility of linear decision rules for hydropower scheduling. *PMAPS*, 2014.

[40] L. G. Valiant. Completeness classes in algebra. In *Proceedings of the eleventh annual ACM symposium on Theory of computing*, pages 249–261. ACM, 1979.

[41] O. Wolfgang, A. Haugstad, B. Mo, A. Gjelsvik, I. Wangensteen, and G. Doorman. Hydro reservoir handling in Norway before and after deregulation. *Energy*, 34(10):1642–1651, 2009.

[42] S. Yakowitz. Dynamic programming applications in water resources. *Water Resources Research*, 18(4):673–696, 1982.

[43] T. Zhao, D. Yang, X. Cai, J. Zhao, and H. Wang. Identifying effective forecast horizon for real-time reservoir operation under a limited inflow forecast. *Water Resources Research*, 48(1), 2012. W01540.

[44] M. Zugno, J. Morales, and H. Madsen. Robust management of combined heat and power systems via linear decision rules. In *2014 IEEE International Energy Conference (ENERGYCON)*, pages 479–486, 2014.

6 Conclusion

We started this thesis by stating the following research question: How suitable are linear decision rules for seasonal hydro power scheduling?

This lead to three sub-questions:

1. Can hydropower specific implementation details, such as head effects and water values, be included in the solution process?
2. Will LDR be able to give quicker results or higher objective values than RI?
3. What bounds on the suboptimality imposed by the LDR approximation can be found?

Now we are able to draw conclusions, based on the computational tests we have conducted.

Our results indicate that head effects and water values can indeed be handled by the use of iterative procedures, so LDR models have sufficient expressiveness to take the details of hydropower scheduling into account. The computational times are reduced significantly in comparison with RI when limited memory is applied. It also seems relatively fast against other solution methods currently in use, even though this needs to be tested before conclusions can be made. However, this comes at a considerable cost. The LDR formulation is not able to give objective values that are competetive with RI's. The solutions values are well below both theoretical upper bounds, and feasible solutions from RI. Thus, the suboptimality of the LDR approximation is quantified and large.

Altogether, it is therefore not clear that LDR is a good fit for the seasonal hydropower planning problem. The topic should be further researched, since there are several interesting questions that arise from our tests.

In Article 2 [9], we saw that piecewise or quadratic decision rules should be investigated, since they potentially can reduce the symmetry, and thus the suboptimality, observed. Additionally, LDR could be fit for other purposes than returning policies to be implemented. If the water values are benchmarked against results from other methods, they could turn out to be usable as input to short-term models. If good limitations on the memory set are found, LDR could be used as a very fast diagnostic tool when quick results are more important than optimal results. Last, we saw in both articles that the bounds on final reservoir levels lead to many dependencies at the end of the horizon. Further research could reveal if LDR models will perform better with an end-of-horizon water valuation instead of a strict bound.

References

[1] J. C. Bean, J. R. Birge, and R. L. Smith. Aggregation in dynamic programming. *Operations Research*, 35(2):215–220, 1987.

[2] R. Bellman. *Dynamic Programming*, volume 1. Princeton University Press, 1957.

[3] A. Ben-Tal, L. El Ghaoui, and A. Nemirovski. *Robust Optimization*. Princeton University Press, 2009.

[4] A. Ben-Tal, A. Goryashko, E. Guslitzer, and A. Nemirovski. Adjustable robust solutions of uncertain linear programs. *Mathematical Programming*, 99(2):351–376, 2004.

[5] A. Ben-Tal and A. Nemirovski. Robust convex optimization. *Mathematics of Operations Research*, 23(4):769–805, 1998.

[6] D. Bertsimas, D. A. Iancu, and P. A. Parrilo. Optimality of affine policies in multi-stage robust optimization. *Mathematics of Operations Research*, 35(2):363–394, 2010.

[7] D. Bertsimas and M. Sim. The price of robustness. *Operations Research*, 52(1):35–53, 2004.

[8] H. Bing, Z. Lizi, and S. Jun. A survey of optimal scheduling of cascaded hydropower stations. *Modern Electric Power*, 1:1–17, 2007.

[9] S. V. Braaten, O. Gjønnes, and K. Hjertvik. Linear decision rules for seasonal hydropower planning: Implementation and computational performance. *Working paper, Norwegian University of Science and Technology, Trondheim*, jun 2015.

[10] S. V. Braaten, O. Gjønnes, and K. Hjertvik. Linear decision rules for seasonal hydropower planning: Modelling considerations. *Working paper, Norwegian University of Science and Technology, Trondheim*, jun 2015.

[11] D. B. Brown, J. E. Smith, and P. Sun. Information relaxations and duality in stochastic dynamic programs. *Operations Research*, 58(4):785–801, 2010.

[12] X. Chen, M. Sim, P. Sun, and J. Zhang. A linear decision-based approximation approach to stochastic programming. *Operations Research*, 56(2):344–357, 2008.

[13] M. Crona. Evaluation of flexibility in hydropower stations. Master's thesis, Uppsala Universitet, 2012.

[14] G. Doorman. *Hydro Power Scheduling*. NTNU, 2009.

[15] J. Downing, Y. Prairie, J. Cole, C. Duarte, L. Tranvik, R. Striegl, W. McDowell, P. Kortelainen, N. Caraco, J. Melack, et al. The global abundance and size distribution of lakes, ponds, and impoundments. *Limnology and Oceanography*, 51(5):2388–2397, 2006.

[16] A. Georghiou, W. Wiesemann, and D. Kuhn. The decision rule approach to optimisation under uncertainty: Methodology and applications in operations management. working paper, 2011, 2011.

[17] A. Gjelsvik and A. Haugstad. Considering head variations in a linear model for optimal hydro scheduling. Unpublished paper from SINTEF Energy Research, Trondheim, Norway, 2005.

[18] A. Gjelsvik, B. Mo, and A. Haugstad. 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*, pages 33–55. Springer, 2010.

[19] Global Energy Observatory. Current list of hydro powerplants. Available at: http://globalenergyobservatory.org/list.php?db=PowerPlants&type=Hydro.

[20] B. L. Gorissen, İ. Yanıkoğlu, and D. den Hertog. A practical guide to robust optimization. *Omega*, 53:124–137, 2015.

[21] C. E. Gounaris, W. Wiesemann, and C. A. Floudas. The robust capacitated

vehicle routing problem under demand uncertainty. *Operations Research*, 61(3):677–693, 2013.

[22] I. Grønvik, A. Hadziomerovic, N. Ingvoldstad, R. Egging, and S.-E. Fleten. Feasibility of linear decision rules for hydropower scheduling. In *2014 International Conference on Probabilistic Methods Applied to Power Systems (PMAPS)*, 2014.

[23] M. J. Hadjiyiannis, P. J. Goulart, and D. Kuhn. A scenario approach for estimating the suboptimality of linear decision rules in two-stage robust optimization. In *50th IEEE Conference on Decision and Control and European Control Conference (CDC-ECC)*, pages 7386–7391, 2011.

[24] M. B. Haugh and L. Kogan. Pricing American options: A duality approach. *Operations Research*, 52(2):258–270, 2004.

[25] J. L. Higle. Stochastic programming: Optimization when uncertainty matters. *INFORMS*, 1(1):30–53, 2005.

[26] F. S. Hillier and G. J. Lieberman. *Introduction to operations research*, chapter 11, pages 533–575. Tata McGraw-Hill Education, 2001.

[27] International Energy Agency. IEA report sets a course for doubling hydroelectricity output by 2050. Available at: https://www.iea.org/topics/renewables/subtopics/hydropower, april 2015.

[28] International Energy Agency. Renewable energy essentials: Hydropower. Available at: http://www.iea.org/publications/freepublications/publication/hydropower_essentials, mar 2015.

[29] J. Jacobs, G. Freeman, J. Grygier, D. Morton, G. Schultz, K. Staschus, and J. Stedinger. SOCRATES: A system for scheduling hydroelectric generation under uncertainty. *Annals of Operations Research*, 59(1):99–133, 1995.

[30] D. Kuhn, W. Wiesemann, and A. Georghiou. Primal and dual linear decision rules in stochastic and robust optimization. *Mathematical Programming*, 130(1):177–209, 2011.

[31] G. Lai, F. Margot, and N. Secomandi. An approximate dynamic programming approach to benchmark practice-based heuristics for natural gas storage valuation. *Operations Research*, 58(3):564–582, 2010.

[32] B. Lehner, C. R. Liermann, C. Revenga, C. Vörösmarty, B. Fekete, P. Crouzet, P. Döll, M. Endejan, K. Frenken, J. Magome, et al. High-resolution mapping of the world's reservoirs and dams for sustainable river-flow management. *Frontiers in Ecology and the Environment*, 9(9):494–502, 2011.

[33] Y. Li, G. Huang, and X. Chen. Multistage scenario-based interval-stochastic programming for planning water resources allocation. *Stochastic Environmental Research and Risk Assessment*, 23(6):781–792, 2009.

[34] D. Mejia-Giraldo, J. Villarreal-Marimon, and H. Lopez-Mejia. Adjustable robust long-term energy planning under hydrological uncertainty. In *PES General Meeting*, pages 1–5. IEEE, 2014.

[35] D. P. Morton. An enhanced decomposition algorithm for multistage stochastic hydroelectric scheduling. *Annals of Operations Research*, 64(1):211–235, 1996.

[36] M. P. Nowak and W. Römisch. Stochastic lagrangian relaxation applied to power scheduling in a hydro-thermal system under uncertainty. *Annals of Operations Research*, 100(1-4):251–272, 2000.

[37] L. Pan, M. Housh, P. Liu, X. Cai, and X. Chen. Robust stochastic optimization for reservoir operation. *Water Resources Research*, 51:409–429, 2015.

[38] M. Pereira. Optimal stochastic operations scheduling of large hydroelectric systems. *International Journal of Electrical Power & Energy Systems*, 11(3):161–169, 1989.

[39] A. B. Philpott and Z. Guan. On the convergence of stochastic dual dynamic programming and related methods. *Operations Research Letters*, 36(4):450–455, 2008.

[40] D. Risberg. Robust optimization applied to seasonal planning of hydropower. Master's thesis, Kungliga Tekniska högskolan, Stockholm, 2015.

[41] P. Rocha. *Medium-term planning in deregulated energy markets with decision rules*. PhD thesis, Imperial College London, 2013.

[42] W. Rondeel. Hydro electricity and storage capabilities in norway - can they be useful for europe? Available at: http://www.uni-saarland.de/fak7/fze/AKE_Archiv/DPG2012-AKE_Berlin/Buch/DPG2012_AKE2.1Rondeel_HydroStorage-inNorway_p49.pdf, March 2012.

[43] N. Secomandi. Analysis and enhancement of practice-based policies for the real option management of commodity storage assets. Working paper, Carnegie Mellon University, Pittsburgh, 2014.

[44] S. Seguin, P. Cote, and C. Audet. Self-scheduling short-term unit commitment and loading problem. *IEEE Transactions on Power Systems*, Early Access(99):1–10, 2014.

[45] A. Shapiro, D. Dentcheva, and A. Ruszczyński. *Lectures on stochastic programming: Modeling and theory*, volume 16. SIAM, 2014.

[46] A. Shapiro and A. Philpott. A tutorial on stochastic programming. Manuscript: Available at http://www2.isye.gatech.edu/people/faculty/Alex_Shapiro/TutorialSP.pdf, 2007.

[47] W. S. Sifuentes and A. Vargas. Hydrothermal scheduling using benders decomposition: accelerating techniques. *IEEE Transactions on Power Systems*, 22(3):1351–1359, 2007.

[48] S. P. Simonovic. Reservoir systems analysis: closing gap between theory and practice. *Journal of Water Resources Planning and Management*, 118(3):262–280, 1992.

[49] S. Stage and Y. Larsson. Incremental cost of water power. *Power Apparatus and Systems, part iii. Transactions of the American Institute of Electrical Engineers*, 80(3):361–364, 1961.

[50] Statistics Norway. Electricity, annual figures, 2013. Available at: http://www.ssb.no/en/energi-og-industri/statistikker/elektrisitetaar/aar/

2015-03-25?fane=tabell\&sort=nummer\&tabell=222049, March 2015.

[51] G. Steeger, L. A. Barroso, and S. Rebennack. Optimal bidding strategies for hydro-electric producers: a literature survey. *IEEE Transactions on Power Systems*, 29(4):1758–1766, 2014.

[52] J. A. Steinsbø. Constructing bidding curves for a hydropower producer at Nord Pool. Master's thesis, NTNU, 2008.

[53] J. A. Tejada-Guibert, S. A. Johnson, and J. R. Stedinger. The value of hydrologic information in stochastic dynamic programming models of a multireservoir system. *Water Resources Research*, 31(10):2571–2579, 1995.

[54] U.S. Department of the Interior. Hydroelectric power. Available at:http://www.usbr.gov/power/edu/pamphlet.pdf, July 2005.

[55] L. G. Valiant. Completeness classes in algebra. In *Proceedings of the eleventh annual ACM symposium on Theory of computing*, pages 249–261. ACM, 1979.

[56] P. Vegard Hansen. Inflow uncertainty in hydropower markets. *The Scandinavian Journal of Economics*, 111(1):189–207, 2009.

[57] S. W. Wallace and S.-E. Fleten. *Stochastic programming models in energy*, volume 10, chapter 10, pages 637–677. Elsevier, 2003.

[58] I. Wangensteen. *Power system economics: the Nordic electricity market*. Tapir Academic Press, 2012.

[59] J. Warrington, P. Goulart, S. Mariethoz, and M. Morari. Policy-based-reserves for power systems. *IEEE Transactions on Power Systems*, 28(4):4427–4437, 2013.

[60] D. W. Watkins, D. C. McKinney, L. S. Lasdon, S. S. Nielsen, and Q. W. Martin. A scenario-based stochastic programming model for water supplies from the highland lakes. *International Transactions in Operational Research*, 7(3):211–230, 2000.

[61] O. Wolfgang, A. Haugstad, B. Mo, A. Gjelsvik, I. Wangensteen, and G. Doorman. Hydro reservoir handling in Norway before and after deregulation. *Energy*, 34(10):1642–1651, 2009.

[62] S. Yakowitz. Dynamic programming applications in water resources. *Water Resources Research*, 18(4):673–696, 1982.

[63] M. Zugno, J. Morales, and H. Madsen. Robust management of combined heat and power systems via linear decision rules. In *2014 IEEE International Energy Conference (ENERGYCON)*, pages 479–486, 2014.

Appendix A Assumptions

A number of assumptions and simplifications have been made when developing the mathematical formulation of the hydropower problem. These are detailed below.

All technical power station data, such as minimum and maximum generation capacities and generator efficiencies, are assumed fully known. Limitations on minimum and maximum reservoir levels are also classified as deterministic data, along with information describing the topology of the water system. The reason behind these assumptions is the fact that uncertainties in station and reservoir data are negligible compared to the uncertainty in the inflow and price parameters that we have chosen to model as stochastic.

Line congestion or line outage is hard to predict and is not considered. Instead the producer is expected to always get the weekly electricity price (on average) for every MWh generated.

The discount rate has been considered deterministic and constant, since the focus of this paper has been on hydropower modeling, not representating the interest rate.

Typical short term operational issues, such as ramping constraints, curtailment costs, and costs related to start-up and shutdown, are neglected due to the length of each time period. In the case of omitting start-up and shutdown costs, there is also a second motive, as including these costs would have turned the model into an NP-hard mixed integer programming problem, rather than a P-hard linear program.

A power plant's effienciency (η) is assumed constant, thus independent of the level of generation (q). Establishing a schedule where generating units operate at their best point is a task for a short-term model, the seasonal model only allocates generation to power stations.

Appendix B Full LDR formulation

By applying the LDR reformulation steps described in Article 1 to the mathematical model presented in both articles, we arrive at the following model.

$$\max \quad \sum_{t\in\mathcal{T}}\sum_{r\in\mathcal{R}}\sum_{d\in\mathbb{D}} T^H \beta_t E^d_{rt}\left(\hat{\pi}_t\hat{x}^d_{rt} + \sum_{u\in\mathbb{U}}\sum_{\tau\in\mathcal{M}^u_t} \Pi_t K^d_{rtu\tau}\mathrm{Cov}(\delta_{Pt},\delta_{u\tau})\right) \tag{B.1}$$

$$\text{s.t.} \qquad \sum_j h_j\mu^{UR}_{rtj} \le \overline{M}_{rt} - \sum_{\tau=1}^{t}\left(\hat{F}_{r\tau} + \sum_{d\in\mathbb{D}}\Big[\sum_{\rho\in\mathbb{C}^d_r}\hat{x}^d_{\rho\tau} - \hat{x}^d_{r\tau}\Big]\right) + M^0_r, \qquad r\in\mathcal{R}, t\in\mathcal{T} \tag{B.2a}$$

$$\sum_j h_j\mu^{LR}_{rtj} \le \sum_{\tau=1}^{t}\left(\hat{F}_{r\tau} + \sum_{d\in\mathbb{D}}\Big[\sum_{\rho\in\mathbb{C}^d_r}\hat{x}^d_{\rho\tau} - \hat{x}^d_{r\tau}\Big]\right) - \underline{M}_{rt} - M^0_r, \qquad r\in\mathcal{R}, t\in\mathcal{T} \tag{B.2b}$$

$$\sum_j H_{ju\tau}\mu^{bR}_{rtj} = \sum_{\substack{t^*\le t:\\ \tau\in\mathcal{M}^u_{t^*}}}\sum_{d\in\mathbb{D}}\Big[\sum_{\rho\in\mathbb{C}^d_r} K^d_{\rho t^* u\tau} - K^d_{rt^*u\tau}\Big] + I^S_{ur}F^V_{r\tau}, \qquad \begin{array}{l} r\in\mathcal{R}, t\in\mathcal{T},\\ u\in\mathbb{U}, \tau\in\mathcal{T}\\ b\in\{U,L\}\end{array} \tag{B.3}$$

$$\sum_j h_j\mu^{Uq}_{rtj} \le \frac{U^q_{rt}}{E^q_{rt}} - \hat{x}^q_{rt}, \quad r\in\mathcal{R}, t\in\mathcal{T} \tag{B.4a}$$

$$\sum_j h_j\mu^{Lq}_{rtj} \le \hat{x}^q_{rt} - \frac{L^q_{rt}}{E^q_{rt}}, \quad r\in\mathcal{R}, t\in\mathcal{T} \tag{B.4b}$$

$$\sum_j H_{ju\tau}\mu^{bq}_{rtj} = K^q_{rt\tau u}, \qquad r\in\mathcal{R}, t\in\mathcal{T}, u\in\mathbb{U}, \tau\in\mathcal{M}^u_t, b\in\{U,L\} \tag{B.5a}$$

$$\sum_j H_{ju\tau}\mu^{bq}_{rtj} = 0, \qquad r\in\mathcal{R}, t\in\mathcal{T}, u\in\mathbb{U}, \tau\notin\mathcal{M}^u_t, b\in\{U,L\} \tag{B.5b}$$

$$\sum_j h_j \mu_{rtj}^{UD} \leq \overline{D}_{rt} - C_{rt}^B - \hat{x}_{rt}^q, \quad r \in \mathcal{R}, t \in \mathcal{T} \tag{B.6a}$$

$$\sum_j h_j \mu_{rtj}^{LD} \leq \hat{x}_{rt}^q + C_{rt}^B - \underline{D}_{rt}, \quad r \in \mathcal{R}, t \in \mathcal{T} \tag{B.6b}$$

$$\sum_j H_{ju\tau} \mu_{rtj}^{bD} = K_{rt\tau u}^q, \quad r \in \mathcal{R}, t \in \mathcal{T}, u \in \mathbb{U}, \tau \in \mathcal{M}_t^u, b \in \{U, L\} \tag{B.7a}$$

$$\sum_j H_{ju\tau} \mu_{rtj}^{bD} = 0, \quad r \in \mathcal{R}, t \in \mathcal{T}, u \in \mathbb{U}, \tau \notin \mathcal{M}_t^u, b \in \{U, L\} \tag{B.7b}$$

$$\sum_j h_j \mu_{rtj}^{Up} \leq \frac{U_{rt}^p}{E_{rt}^p} - \hat{x}_{rt}^p, \quad r \in \mathcal{R}, t \in \mathcal{T} \tag{B.8a}$$

$$\sum_j h_j \mu_{rtj}^{Lp} \leq \hat{x}_{rt}^p - \frac{L_{rt}^p}{E_{rt}^p}, \quad r \in \mathcal{R}, t \in \mathcal{T} \tag{B.8b}$$

$$\sum_j H_{ju\tau} \mu_{rtj}^{bp} = K_{rt\tau u}^p, \quad r \in \mathcal{R}, t \in \mathcal{T}, u \in \mathbb{U}, \tau \in \mathcal{M}_t^u, b \in \{U, L\} \tag{B.9a}$$

$$\sum_j H_{ju\tau} \mu_{rtj}^{bp} = 0, \quad r \in \mathcal{R}, t \in \mathcal{T}, u \in \mathbb{U}, \tau \notin \mathcal{M}_t^u, b \in \{U, L\} \tag{B.9b}$$

$$\sum_j h_j \mu_{rtj}^{Ls} \leq \hat{x}_{rt}^s, \quad r \in \mathcal{R}, t \in \mathcal{T} \tag{B.10}$$

$$\sum_j H_{ju\tau} \mu_{rtj}^{Ls} = K_{rt\tau u}^q, \quad r \in \mathcal{R}, t \in \mathcal{T}, u \in \mathbb{U}, \tau \in \mathcal{M}_t^u \tag{B.11a}$$

$$\sum_j H_{ju\tau} \mu_{rtj}^{Ls} = 0, \quad r \in \mathcal{R}, t \in \mathcal{T}, u \in \mathbb{U}, \tau \notin \mathcal{M}_t^u \tag{B.11b}$$

$$\mu_{rtj}^{UR}, \mu_{rtj}^{Uq}, \mu_{rtj}^{UD}, \mu_{rtj}^{Up} \geq 0, \quad r \in \mathcal{R}, t \in \mathcal{T}, j \tag{B.12}$$

$$\mu_{rtj}^{LR}, \mu_{rtj}^{Lq}, \mu_{rtj}^{LD}, \mu_{rtj}^{Lp}, \mu_{rtj}^{Ls} \geq 0, \quad r \in \mathcal{R}, t \in \mathcal{T}, j \tag{B.13}$$

$$\hat{x}_{rt}^q, \hat{x}_{rt}^s \geq 0, \quad r \in \mathcal{R}, t \in \mathcal{T} \tag{B.14}$$

$$\hat{x}_{rt}^p \leq 0, \quad r \in \mathcal{R}, t \in \mathcal{T} \tag{B.15}$$

$$K_{rt\tau u}^q, K_{rt\tau u}^p, K_{rt\tau u}^s \quad \text{free} \quad r \in \mathcal{R}, t \in \mathcal{T}, u \in \mathbb{U}, \tau \in \mathcal{M}_t^u \tag{B.16}$$

Here, (B.1) is the objective function, as described in Article 1. The constraints are hard to understand or describe intuitively, but are grouped according to their origin. Generally, μ denotes a dual variable from the transformation, $\hat{x}$ an intercept, and K a slope. A dual variable $\mu^{U\cdot}$ corresponds to an upper bound, while $\mu^{L\cdot}$ corresponds to a lower bound. Where equivalent constraints differ only in the dual variables, the shorthand b is used, to stand for either U or L.

Constraints (B.2) and (B.3) are the transformed reservoir limits, with expressions for the reservoir level substituted in. We call the constraints on the form (B.2) *dual objective* constraints because they correspond to the objective value on the left hand side in the transformations. There is one dual objective constraint for each constraint in the original problem. (B.2a) correspond to upper bounds, while (B.2b) correspond to lower bounds. Equations like (B.3) are called *dual feasibility* constraints, because they define the feasible region of the left hand side in the transformations.

Constraints (B.4) and (B.5) are transformed from the bounds on the power generation. The two cases (B.5a) and (B.5b) separate between the uncertainty parameters $\delta_{u\tau}$ that the generation in period t depends on, and those it does not, respectively. As mentioned in Article 1, uncertainty parameters outside the memory set $\mathcal{M}_t^u$ lead to zero-valued right hand sides.

Similarly, (B.6) and (B.7) correspond to the discharge limits, while (B.8) and (B.9) originate from the pumping limits. (B.10) and (B.11) are the transformed limits on spill, and notably only one dual objective constraint is needed here, since spill is unconstrained from above. As a last note, (B.15) reflects the fact that pumping is defined as a non-positive variable, while all other intercepts and dual variables are non-negative, and slopes are free.