

Optimizing day-ahead bid curves in hydropower production

Ellen Krohn Aasgård¹ · Christian Øyn Naversen² ·
Marte Fodstad² · Hans Ivar Skjelbred²

Received: 24 May 2016 / Accepted: 7 February 2017 / Published online: 15 February 2017
© Springer-Verlag Berlin Heidelberg 2017

Abstract In deregulated electricity markets, hydropower producers must bid their production into the day-ahead market. For price-taking producers, it is optimal to offer energy according to marginal costs, which for hydropower are determined by the opportunity cost of using water that could have been stored for future production. At the time of bidding, uncertainty of future prices and inflows may affect the opportunity costs and thus also the optimal bids. This paper presents a model for hydropower bidding where the bids are based on optimal production schedules from a stochastic model. We also present a heuristic algorithm for reducing the bid matrix into the size required by the market operator. Results for the optimized bids and the reduction algorithm are analyzed in a case study showing how uncertain inflows may affect the bids.

Keywords Unit-commitment · Stochastic programming

1 Introduction

In deregulated electricity markets, hydropower companies optimize their production independently and must determine the optimal supply curve in order to maximize revenue. This supply curve determines the bids to be sent to the day-ahead energy market. The main contribution of this paper is the demonstration of how short-term hydropower scheduling models can be extended to explicitly give optimized bid curves.

✉ Ellen Krohn Aasgård
Ellen.Aasgard@iot.ntnu.no

¹ Norwegian University of Science and Technology, Alfred Getz vei 3, 7491 Trondheim, Norway

² SINTEF Energy Research, Sem Sælands vei 11, 7034 Trondheim, Norway

With this method, we consider the role of inflow uncertainty when determining the optimal bids.

For price-taking producers, it is optimal to offer energy according to the marginal costs of production. However, finding the true marginal costs for hydropower is difficult due to the storability of the water resources. Precipitation is a “free” resource that only has a cost due to its limited availability. The marginal costs are therefore given by the opportunity cost of using water today that could have been stored for production in the future. The cost of using water therefore has to be based on the optimal production schedule over the time horizon given by the size of the reservoir. The long-term optimal release of water is typically calculated on a 3- to 5-year horizon by aggregated models. The results are used in medium- and short-term models where the physical system is more precisely described as the operating hour is approached [1,2]. The challenge for short-term scheduling is to balance the profit obtained from producing within the short-term horizon with the expected value of storing the water [3]. When bidding into the day-ahead market, this balance is struck by finding the optimal supply curves.

Prices and inflows are naturally uncertain, and long- and medium-term scheduling account for these uncertainties. In the short-term, however, prices and inflows have traditionally been taken as deterministic to allow for greater detail in the modelling of the physical system. Several efforts [3–8] have tried to incorporate the uncertain nature of prices and/or inflows also in the short-term. Price uncertainty is to some degree addressed by the bids themselves due to the price dependency in the bid curves. The supply curve described by the bids expresses the willingness to produce certain volumes at a certain price—more production is offered as the price increases. Even if the realized price is different from the forecasted value, bids according to the marginal costs should give satisfactory results. The bid curves are however not inherently inflow dependent, even though inflow may affect the *level* of the marginal costs. If a reservoir is small compared to the inflow or the discharge capacities of connected plants, so that the reservoir can be filled or depleted weekly or even daily, the availability of water and hence the opportunity costs may be affected by uncertain inflows in the short-term. If a large amount of water is available, either stored in the reservoir or as inflow, the marginal costs of using water for production are generally low. If less water is available, the marginal costs are higher as the resource is scarcer. Formulating a model for optimal bidding, we demonstrate that it might be beneficial to account for uncertainty in inflows as well as prices.

In the next section, we give some background on the hydropower bidding and scheduling problem. Section 3 outlines the modelling of bids. Sections 4 and 5 show results for a case study, where we first look at stylized examples of how inflow uncertainty may affect the bids and then show results obtained with realistic input for part of a Norwegian reservoir system. Section 6 gives some concluding remarks.

2 Short-term hydropower scheduling and bidding

Hydropower scheduling is a subset of the general optimal power scheduling and unit-commitment problem [9]. The bidding and scheduling problems are closely linked due to the fact that the cost of production is determined by the opportunity cost.

We refer to [10] for a recent review of methods applied to both problems. In the following we concentrate on a few studies that are similar to ours. In [4], a model for determining optimal offer curves using stochastic dynamic programming is presented. Prices are stochastic, but inflows are taken to be deterministic as it is assumed that future inflows will come mainly from rain that has already fallen due to the short time horizon. Depending on local weather conditions and terrain, this may not be a valid assumption and it might be necessary to also consider stochastic inflow. Aimed at medium-term scheduling, [5] extends the dynamic programming framework to also include stochastic inflows, and notes that the optimal bids for one trading period depend on what happens in the rest of the time horizon of the reservoir. The authors of [3] formulate a model of optimal bidding that accounts for uncertain prices and inflows, but the effect of including inflow uncertainty is not investigated. The related work in [6] involves no bidding but shows examples of how uncertain inflow affects the spatial distribution of production between the plants in a river cascade. In [7], stochastic inflow is considered in an effort to balance the conflicting objectives of high pressure height and risk of spillage. Higher head means increased efficiency which gives greater returns on the water released. On the other hand, keeping the reservoir level high when inflows vary may lead to losses due to spillage. For small reservoirs, this dynamic may be important on a daily or even hourly basis and should be reflected in the bid curves. Larger hydropower systems may consist of several different sized reservoirs and plants along the same river system, and the total offer strategy may be constrained by the decisions for the smaller reservoirs. To get a correct assessment of the cost of production, uncertainty in inflow should be included for reservoirs where it may have an effect—but this is difficult to determine a priori. As adding a new stochastic parameter increases the complexity of a stochastic model, the sensitivity of the marginal costs to uncertainty in inflow should be investigated.

We base our modelling on the current Nordic market setting, where large producers solve the short-term scheduling problem by deterministic successive linear programming (SLP) [11]. The SLP method involves solving MIP and LP's in iterations to account for nonlinearities in the problem, and is available as a general software that is in operational use for a range of different topologies [12]. A stochastic version of successive linear programming (SSLP) has been implemented within the framework of the current software [8]. The method is based on a discrete representation of uncertainty in the form of scenario trees [13]. This current work extends the SSLP model to also include decision support for optimal bidding. The model is implemented in C and can be solved with CPLEX, Gurobi or Coin solvers. Alternative formulations for the short-term hydropower scheduling problem are used in other parts of the world, and the implementation of optimal bids based on the optimized production schedules may be a relevant extension also in other models.

3 Bid matrix generation and reduction

To explain how bids are modelled, a simplified presentation of the full mathematical problem [8] is given. Focus is on the basic aspects of the model, and difficult topologies, optimization of pressure height and other aspects typical and important for hydropower

scheduling are not presented. We define T to be the set of time steps and S to be the set of scenarios. We consider a river chain with R reservoirs, and assume for simplicity that all reservoirs have an associated generating unit. The objective is to maximize the revenues from selling power within the short-term model horizon and the value of water left in the reservoirs at the end of the horizon, less the costs related to start-ups. Revenues are calculated as the price of electricity (ρ_{ts}) times the volume produced (y_{tsr}). The value of storage $\pi_{s,r}$ describes the expected value of saving water for production in later periods. The value of storage is a function of the end-of-horizon water levels and should be understood as $\pi_{s,r}(s_{T,s,r})$. However, in the remainder of the text this functional form is dropped to ease notation. The value of storage is calculated from long- or medium-term models such as [14] and are taken as known input to the bidding model. Each time a generating unit is started, the variable v_{tsr} will take on a value of 1 and a cost of C_r^{Start} is incurred. The objective may thus be expressed as in Eq. (1).

$$\max \sum_{TSR} \rho_{ts} y_{tsr} + \sum_{SR} \pi_{sr} - \sum_{TSR} C_r^{Start} v_{tsr} \tag{1}$$

Storage in reservoirs must adhere to a mass balance equation between time steps, as expressed in Eq. (2).

$$s_{tsr} = s_{t-1sr} - q_{tsr}^P - q_{tsr}^B - q_{tsr}^S + \sum_{R^{up}} (q_{tsr}^P + q_{tsr}^B + q_{tsr}^S) + Q_{tsr} \tag{2}$$

The storage level at the end of this hour, s_{tsr} , must equal the storage level at the end of the previous hour less any discharge used for production (q_{tsr}^D), bypass (q_{tsr}^B), or spillage (q_{tsr}^S) within this hour. The system consist of a series of connected reservoirs, and any discharge from connected upstream reservoirs, R^{up} , must be added to the balance equation. For simplicity, Eq. (2) does not consider water flow travel time between reservoirs, but this feature is included in the full model. Inflow, Q_{tsr} , which may be stochastic, is also added. Reservoir levels also have to be within known maximum and minimum bounds.

The release of water is linked to power production by a piecewise linear concave production function for each generator. This production function has I line segments, each represented with a discharge volume q_{tsri} limited by an upper bound Q_{ri}^{max} and a power output rate E_{ri} . The sum of power produced from all segments must equal the power produced and requirements for minimum generation level, P_r^{min} , if the unit is on. This is expressed by Eq. (3).

$$\sum_I E_{ri} q_{tsri}^P = y_{tsr} + P_r^{min} u_{tsr} \tag{3}$$

Here, u_{tsr} is a binary variable that describes the on/off decision of each generating unit. This binary variable is also used in Eq. (4) to determine if a start-up has occurred between two consecutive time steps. If so, the auxiliary variable v_{tsr} takes on a value of 1 and a cost is incurred in the objective function.

$$v_{tSR} \geq u_{tSR} - u_{t-1SR} \quad (4)$$

Optimizing the objective in Eq. (1) subject to the restrictions in Eqs. (2)–(4) gives results in the form of optimal production volumes for every hour of every scenario.

3.1 Bidding in the Nordic day-ahead market

For participants in the Nordic day-ahead market organized by Nord Pool, the bids are sent to the market operator in the form of price-volume pairs, summarized in a bid matrix, see [3] for details. This matrix describes the supply curve of the producer for each hour the coming day. The idea behind our modelling of bids is previously described in [15, 16], and is quite straightforward: Production is optimized in every hour of every scenario, and it is from these production schedules that the bid matrix is created. In this matrix, every price found in the scenarios gives a column and every hour gives a row, while the optimized net power production for that hour and price is the bid. According to the market rules of NordPool, there are two criteria that must be fulfilled for all submitted bids:

1. All bids within an hour must be non-decreasing for an increasing price
2. The number of columns in the bid matrix cannot exceed 64.

The first requirement is handled by adding constraints for the production variables, while the second is addressed by post-calculations. This is further described in the following two subsections. In the last subsection, we describe the options and requirements for the shape of the scenario tree. Our modelling of the shape of the bid curve (e.g. piecewise linear or piecewise constant) is flexible and may be adapted to criteria for other market systems.

Our formulation of bids differ from the method used in [3], which also describes hydropower bidding for the Nordic area. The core of the problem is that simultaneous optimization of both the bid prices and bid volume is nonlinear. In [3], this is solved by selecting a set of fixed prices for which to bid. These price points make up the columns of the bid matrix and may be selected based on the distribution of (price) scenarios or other heuristics. In our formulation, we give all the scenario prices as potential bid points to the optimization model, and select the most important points *after* the optimization has found the production for every price found in the scenario tree. This gives a large bid matrix, and a heuristic algorithm is used to reduce the number of columns in the matrix. The new formulation reduces the number of variables since no explicit bid variables are needed. Further, post-selecting the bid prices in the reduction algorithm rather than selecting prices a priori may give stronger guidance to the selection due to the knowledge of the optimized bid curve. A comparison of our proposed method to the method in [3] is given in Sect. 5.2.

3.2 Optimization constraints

Consider the three scenarios in Fig. 1. The model will optimize the production in all 12 points, corresponding to the 3 scenarios covering 4 h. Looking at each hour separately,




Fig. 1 Illustration of prices in three different scenarios. Note that the scenarios may cross each other

we can ensure that the bid is always non-decreasing for increasing prices. In the first hour, the production in scenario s_1 cannot exceed that of s_2 , as the price in s_1 is lower than in s_2 . Similarly, the production in s_2 cannot exceed that of s_3 . In the second hour, however, the order has changed since the price in s_1 now is greater than in s_2 . This leads to restrictions of the type expressed in Eq. (5).

$$y_{ts} \leq y_{ts^+}, \quad \text{if } \rho_{ts} < \rho_{ts^+}. \tag{5}$$

Here, ρ_{ts} is the price in time step t and scenario s , and y_{ts} is the corresponding net production. The index s is part of the set $S \setminus s^{max}$, where S is the set of all scenarios and s^{max} is the scenario with the highest price within the hour. The scenario s^+ is defined as the scenario with the lowest price that is still greater or equal to ρ_{ts} in that hour. In the case where $\rho_{ts} = \rho_{ts^+}$, the production is required to be equal in the two scenarios, as expressed in Eq. (6).

$$y_{ts} = y_{ts^+}, \quad \text{if } \rho_{ts} = \rho_{ts^+}. \tag{6}$$

In fact, any two scenarios that share the same price should have equal volumes in that hour, as only one volume can be bid for each price. In Fig. 1, s_1 and s_3 have the same price in the fourth hour and should thus offer the same production, regardless of the volumes offered in the other hours. With additional stochastic parameters, such as inflows or prices in other markets, there might be a situation where two scenarios have the same price, but differ in the values for other parameters. To get a consistent bid curve, we require that all (multivariate) scenarios that share a common price must bid the same volume for that price.

Equations (5) and (6) give $(N_s - 1)T$ new constraints, where N_s is the number of scenarios and T is the number of hours with bidding. To add these constraints to the optimization model, we first define

$$y_{ts} = P_{ts} - C_{ts} \tag{7}$$

where P_{TS} is generation and C_{TS} is load for the producer. The variables P_{TS} and C_{TS} in Eq. (7) can be used on any level of aggregation: plant, group of plants or total system. In general, the plants of the system are divided into N_g groups so that

$$P_{TS} = \sum_{g \in G_n} P_{TSi} \tag{8}$$

$$C_{TS} = \sum_{g \in G_n} C_{TSi} \tag{9}$$

where G_n is the set of plants in group $n \in N_g$.

Letting all plants be part of the same group enforces the net production of the entire system to obey the requirement of non-decreasing bids for increasing price, and a bid matrix for the total system is the output. Creating a group for each plant will result in bids on plant level, which is useful for validation of the results as operating points for each plant can be recognized from the supply curves.

Our modeling assumes that all production is sold in the day-ahead market so that net production corresponds to sale. The method can easily be adapted to handle fixed delivery or purchase obligations by subtracting the net committed load from the right-hand-side of Eq. (7). To allow for optimization of trade in multiple markets with delivery in the same hour, new variables would have to be added to allocate the production to the different markets.

3.3 Bid matrix reduction

The restriction on the total number of columns, or prices, in the final bid matrix is not enforced by the optimization. The bid matrix directly from optimization will have one column for each price found in the scenario tree. This means that a case with 25 scenarios over 24 h may have up to 600 columns, which is far more than the allowed number of 64. A method for removing the least significant price columns, based on a greedy algorithm, is therefore used to reduce the optimal bid matrix.

Before reducing the number of columns, the bid matrix needs to be completed. In the given example, we have 600 columns for any given hour, but only 25 optimized bids corresponding to the prices in any hour. The 575 remaining columns are chosen to be linear interpolations of their neighbouring optimal bids, as interpolation is what is done in the market clearing if the realized price falls between two bids. This means that the 575 extra columns do not change the geometric shape of the bid matrix. The lowest and highest optimized bids within each hour are set as the endpoints.

From the completed matrix with the optimized and linearized bids, we start the reduction algorithm. The consequence of removing a column is illustrated in Fig. 2. Removing any of the optimized columns in this hour will, in general, change the shape of the bid curve. Removing one of the linearized bid columns will not change the shape of the bid curve in this hour, but it will affect the curve in other hours as each column is significant in at least one hour.

The goal of the reduction algorithm is to make the minimum geometric change to the total bid matrix. The immediate cost of removing a price column in a given hour

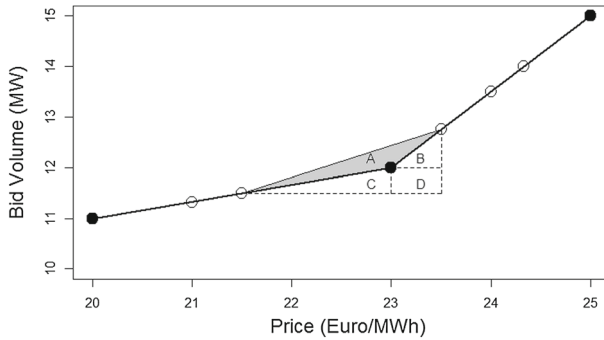


Fig. 2 Part of a bid curve for a single hour with three optimized bid points (*black dots*) and a number of linearized bids (*white dots*). Removing the middle optimized bid from the matrix changes the geometric shape of the curve by the area A. The cost can be calculated from areas B, C and D

is defined as the area between the original curve and the reduced curve, as shown by area A in Fig. 2. Bid number i has the coordinates (ρ_i, y_i) , where ρ_i is the price of column i and y_i is the bid volume in column i . To calculate the cost of removing the middle point in Fig. 2, area A must be calculated. The area of the total triangle is

$$A + B + C + D = \frac{1}{2}(y_3 - y_1)(\rho_3 - \rho_1) \tag{10}$$

where

$$B = \frac{1}{2}(y_3 - y_2)(\rho_3 - \rho_2), \tag{11}$$

$$C = \frac{1}{2}(y_2 - y_1)(\rho_2 - \rho_1), \tag{12}$$

$$D = (y_2 - y_1)(\rho_3 - \rho_2). \tag{13}$$

Solving for area A results in

$$A = \frac{1}{2}y_3(\rho_2 - \rho_1) - \frac{1}{2}y_2(\rho_3 - \rho_1) + \frac{1}{2}y_1(\rho_3 - \rho_2). \tag{14}$$

The expression of the cost Γ_{ti} of removing column i in time step t is thus

$$\Gamma_{ti} = \left| \frac{1}{2}y_{ti}(\rho_i - \rho_{i-1}) - \frac{1}{2}y_{ti}(\rho_{i+1} - \rho_{i-1}) + \frac{1}{2}y_{ti-1}(\rho_{i+1} - \rho_i) \right| \tag{15}$$

The absolute value in Eq. (15) is necessary to account for both concave and convex areas of the bid curve. The total cost of removing column i from the bid matrix will be the sum over all hours as in Eq. (16).

$$\Gamma_i = \sum_{t \in T} \Gamma_{ti} \tag{16}$$

Equation (15) is not valid for the first and last columns, and the costs of removing these are set to a large value as they are often decisive to the structure of the bid matrix and should not be removed. Using the sum in Eq. (16), the total immediate cost of removing a column is found. The final reduced bid matrix is found by removing the cheapest column one by one until the bid matrix has sufficiently few columns left. Always removing the cheapest column is a greedy heuristic that may not give the overall optimal strategy. It should be noted that the term cost or immediate cost as used above is simply a measure of the area between the full and reduced curves, not the actual cost of being dispatched according to alternative curves. A change in any offer curve will have real cost consequences not only for the period the offer curve is put into, but also for subsequent periods due to time-dependencies of discharges along the river. The immediate cost of removing any column of the bid matrix is thus internal to the optimization model, while the actual cost of dispatches resulting from alternative bid curves are very real to the producer. The greedy algorithm we propose is only concerned with removing the cheapest columns one by one, and does not consider any other consequences of removing the column. The performance of using the reduced versus the full curve in actual operations is a topic for further research, but we do show some results for the reduction algorithm in Sect. 5.1.

3.4 Scenario tree structure

Input for the stochastic parameters is described in the form of multivariate scenario trees. In this work the day-ahead market price and inflows to the reservoirs are considered to be stochastic. Price uncertainty is needed for our formulation of the bid decisions, but inflow can be chosen to be stochastic or deterministic depending on how important it is for the specific topology. Other stochastic parameters may be included in the model in the future, such as loads, maintenance and prices in other markets.

Regarding the scenario tree structure, the model handles both two-stage (fan) and multi-stage trees. We recommend having a high branching factor prior to the bidding period, since the prices should span out a bid curve that covers the operating area of the plants. With fewer scenarios, there is less information to build a good bid curve from. Figure 3 shows different tree structures that are allowed in our formulation. In the case study in Sects. 4 and 5, we show results for fan trees of the type to the left in Fig. 3. We also recommend adding (extreme) minimum and maximum scenarios to make sure that the bid curves are meaningful even for prices outside the forecasted values, similarly to what is done in [17].

4 Case study: examples with exaggerated inflow uncertainty

This section presents results for a case study from part of a Norwegian watercourse. A sketch of the topology is given in Fig. 4. For all cases, a 1-week horizon is used and we bid for the second day, i.e., hours 25–48. The marginal value of water is given as a constant value for each reservoir. We first analyze how inflow uncertainty may affect the bid curves by looking at special situations for the downstream reservoir. By looking




Fig. 3 Examples of scenario trees. Bidding occurs for the hours enclosed in the *grey box*. To the *left*, a two-stage (fan) tree with six scenarios is shown. The *middle* tree also has six scenarios in total, but only three individual scenarios during the bidding period. This tree may therefore give a less precise bid curve, as there is less information to construct the curve from. The tree to the right shows that branching also may occur within the bid period




Fig. 4 The topology of the system under study. The upstream reservoir represents the aggregated capacity of several upstream reservoirs and is very large. Inflow can be stored here between seasons. Water is drawn from the upstream reservoir to two plants in series with a small reservoir in between, which can be depleted and refilled almost daily. The small reservoir is a bottleneck in the system. The difficulty when determining the production schedule for this system is to find the correct balance between production from the upper plant, high pressure height and risk of spillage from the downstream reservoir

at problem instances with stylized inflow uncertainty, we are able to isolate and explain the effects of considering uncertain inflows. We also investigate the sensitivity of the bid curves to uncertain inflows depending on the size of the downstream reservoir.

4.1 Special situations with high and low inflow

To look at the effect of inflow uncertainty, we look at two special cases: (1) possible high inflows and (2) possible low inflows. In both cases, the scenario tree consists of five scenarios for price combined with two scenarios for inflow. In the first instance, the highest inflow scenario provides inflow of the same size as the discharge capacity of the plant below. In the deterministic case, mean inflow is used which is less than the capacity of the plant. The bids are optimized with and without considering inflow uncertainty and the results can be seen in Fig. 5.




Fig. 5 The bid curves with (black) and without (grey) inflow uncertainty for the problem instance with possible high inflows. The results show that the bid curve based on deterministic inflows is below the bids with stochastic inflows, meaning that a somewhat higher marginal cost is calculated in the deterministic case




Fig. 6 The bids obtained with and without inflow uncertainty for the same situation as in Fig. 5, but for a single hour (Hour 7). The results for the other hours are comparable

To explain why the bid curves differ, we look at how the offered energy is divided between the plants. This is accomplished by running the same problem instance with the restriction in Eqs. (5) and (6) on plant level instead of for the total system. In the face of possible spill in the downstream reservoir, it is desirable to release less water from the upstream plant while releasing more water from the downstream plant. In fact, neither the case with deterministic nor stochastic inflow release any water from the upstream plant in this case, so the volumes offered in Fig. 5 are all supplied from the downstream plant.

Looking at Fig. 6 which shows the bids for a selected hour, it can be seen that for low prices, the bid curve based on deterministic inflow reflect that some water should be saved, while the stochastic inflow case produce close to maximum production of 150 MW almost independent of the price. The instance with deterministic inflow does not see the full risk of spill, and will not produce at lower efficiency above best-point for the lower prices. The stochastic inflow model, however, recognizes that in the face of potential spill, it is better to produce at lower efficiencies to avoid spill.

Table 1 The time to fill or deplete the downstream reservoir with maximum discharge from the connected upstream and downstream plants

Downstream reservoir size	Time to fill (h)	Time to deplete (h)
Normal	27.8	13.9
Large (double)	55.5	27.8
Small (half)	13.9	6.9

In the second situation, one of the inflow scenarios gives zero inflow and thus a potential for low pressure height in the downstream reservoir. This is unwanted as a higher head gives greater returns on the water released. In this case, it is desirable to release more water from the upstream plant in order to raise the pressure height in the small reservoir. By releasing water from the upstream plant, it is possible to fill the downstream reservoir in about 28 h, after which the situation is again similar to the high inflow case as even small amounts of inflow gives a minor potential for spill in the downstream reservoir. This strategy is reflected in the bid curves with both stochastic and deterministic inflow, which both suggest that water should be released at maximum production for the upstream plant independently of the price. For the downstream plant, the results are comparable to what is shown for the high inflow case in Figs. 5 and 6.

From these two special cases of possible high and low inflow, it is evident that uncertainty in inflow affects the production strategy. Including stochastic inflow in the optimization of bids gives offer curves that better reflect the true opportunity costs.

4.2 Sensitivity of bid curves to river chain flexibility

To investigate the effect of considering inflow uncertainty depending on the limited flexibility of the downstream reservoir, the size of the reservoir is changed. Particularly, we look at situations where the downstream reservoir is doubled or halved in size. Keeping the capacities of the upstream and downstream plants the same, the time to fill/deplete the reservoir for the instances tested is given in Table 1. The results for all test instances are summarized in Table 2.

We first show results for the situation with possible high inflow. When the downstream reservoir is larger, the increased storage capacity allows for some production from the upstream plant, which was not the case in the face of high inflows in the previous section. The bid curve for the upstream plant is shown in Fig. 7. For this situation, it is possible to release some water from the upstream plant without spillage downstream, and this is recognized by both bid strategies. For lower prices, more production is offered when inflow is deterministic because the model does not see that releasing water now might lead to production at lower prices or efficiencies later in the week to avoid spillage. This risk is recognized with stochastic inflow, and hence more modest volumes are offered. For the downstream plant, the bids for the larger reservoir are comparable to the results shown in Figs. 5 and 6 for the normal reservoir size.

Table 2 Summary of test results with exaggerated inflow uncertainty

Reservoir size	High inflow		Low inflow	
	Upstream plant	Downstream plant	Upstream plant	Downstream plant
Normal	No production	Similar to Fig. 6	Max production	Similar to Fig. 6
Large	Similar to Fig. 7	Similar to Fig. 6	Max production	Similar to Fig. 6
Small	No production	Similar to Fig. 6	Max production	Similar to Fig. 6

For all instances, the strategy is to produce as much as possible from the upstream reservoir to have a high pressure height in the downstream reservoir and then bid similarly as in Fig. 6 for the downstream plant

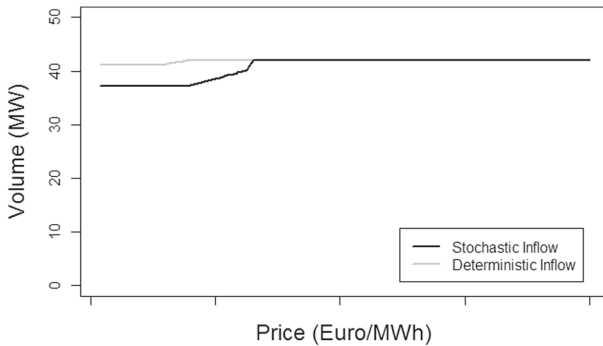


Fig. 7 The bid curve for the upper plant when the downstream reservoir is doubled in size and high inflows are expected

When the reservoir size is halved, there is again no production from the upstream plant when high inflow is expected. The bids for the downstream plant are again similar to Figs. 5 and 6.

If we instead look at the situation with possible low inflow, the strategy for all reservoir sizes is to release at maximum production from the upstream plant to refill the lower reservoir. After the lower reservoir is replenished, the bids are again comparable to Figs. 5 and 6.

5 Case study: realistic input

We now move away from the exaggerated examples of the previous section and investigate how the bid optimization model behaves for more realistic input for the topology presented in Fig. 4. To generate a scenario tree, 60 scenarios for price are obtained from a market analysis company [18] and 50 scenarios for inflow are taken as the ensemble forecast from the HBV model [19] for the same day. The input corresponds to 10 individual days in August 2012. We assume that prices and inflows are independent and combine them into a scenario fan with 3000 scenarios from which scenario trees are generated using the scenario tree reduction algorithm described in [20]. Prices and inflows may indeed not be independent, and the yearly or seasonally accumulated inflow within a price area is in fact a strong driver of the price in hydro-dominated




Fig. 8 Example of inflow scenarios to the downstream reservoir




Fig. 9 Example of scenarios for prices

systems. The correlation between local inflow and the price is much weaker, and may also vary from watercourse to watercourse. Independence is assumed in [3, 6] where it is found that prices and inflows are only contemporaneously correlated. Any correlations between inflows and prices may be considered by properly constructing the tree and combining scenarios. Examples of scenario values for inflow to the downstream reservoir and price can be seen in Figs. 8 and 9, respectively.

For the input values given in Figs. 8 and 9, we obtain the bids shown in Fig. 10 for all hours and in Fig. 11 for one selected hour. The results show that slightly larger volumes are offered if inflow is considered stochastic rather than deterministic. The difference in the bid curves stem from the same effects described earlier for the example problems, and strikes a balance between production upstream, high head and risk of spillage for the downstream reservoir. As risk of spillage is best represented by stochastic modelling of inflows, bid curves obtained when considering uncertain inflow will better represent the true opportunity costs.




Fig. 10 Bid curves with (*black*) and without (*grey*) inflow uncertainty for the realistic case study. For all hours of the day, the same or a higher bid curve (lower marginal cost) is found for the stochastic inflow case




Fig. 11 Bid curves with and without inflow uncertainty for the realistic case study for a selected hour (Hour 1). When inflow is considered to be stochastic, slightly more production is offered for the same price than if inflow is considered deterministic. This means that the calculated marginal cost is lower for the stochastic inflow case

5.1 Bid matrix reduction

For the realistic case study, we show the effect of reducing the bid matrix from its initial size down to the 64 points currently allowed by Nord Pool. For each individual hour, the bid curve will be a piecewise linear function between the price-volume pairs given in the bid matrix. This can be seen from Fig. 12 which show the points from the complete and reduced bid matrix for a selected hour. For all hours and for all instances tested, we do not find any significant differences between the full and reduced matrix. The costs of reducing the matrix are very small, totaling less than a few euros even for the largest instances.

The costs of reducing bid matrices of various sizes down to the required size are plotted in Fig. 13. The size of the initial matrix is determined by the number of unique price scenarios in the tree. The cost is really a measure of the change in the shape of the bid curve due to reduction, but has units of Euro. The change in shape is very small and so are the costs.




Fig. 12 The bid curve for an individual hour with the points from the full and reduced matrix plotted




Fig. 13 The cost of reducing the full optimal bid matrix down to the required size of 64 columns

5.2 Comparison to fixing bid points prior to optimization

As discussed in Sect. 3.1, our model differs from the model presented in [3] by the fact that the columns in the final bid matrix are chosen after the optimization. In contrast, [3] choose bid points prior to the optimization based on the characteristics of the prices in the scenario tree. Our reasoning for selecting bid prices after the optimization is that knowledge of the optimal quantity for each price would give a better guidance for selecting the bid points that best represents the true curve given that only 64 points may be used. A comparison between our proposed method and the method in [3] is shown in Fig. 14. The same holds for other hours and days tested, which indicate that the two methods give comparable results and may be used interchangeably by producers. Indeed, the method in [3] is implemented as an option in our model, whereby the user gives a self-selected set of bid prices as input to the model. Even so, an advantage of our method is that it is less reliant on user input in the form of a good selection of bid points.




Fig. 14 The bid curves obtained from our proposed model compared to the method presented in [3], where bid prices are selected prior to the optimization. For the alternative method, 64 bid prices were selected based on the distribution of prices in the scenario tree, so that a finer resolution of points is used in the most likely range of prices. The bid curves obtained from the two methods are very similar. Tests using equidistant points in the alternative method also shows the same results

5.3 Problem size and calculation time

This section gives details of the computational performance of the model. The numbers given in this section has been obtained by using CPLEX 12.6.3 on a 3 GHz Intel PC with 8.0 GB RAM. A set of scenario trees of different sizes are generated from an initial fan tree of 3000 scenarios. For the instance with 25 scenarios, the problem has 287,151 variables (of which 3600 are binaries) and 111,195 constraints. The problem size grows linearly with the number of scenarios. Calculation times are shown in Fig. 15. The results show that the calculation time is higher when finding the optimal bids and not just the production schedule. The constraints expressed in Eqs. (5) and (6) that are added in order to optimize the bids couple decisions in individual scenarios that are otherwise independent. These constraints, although in another form, are also included in the method in [3], and for our instances we cannot assert which of the methods is fastest in general. The scenario-coupling constraints makes the problem harder to solve, but gives the bid matrix as a direct result from the optimization. This means that less manual adaptation of the results from either a stochastic model or multi-scenario runs of the deterministic model is needed. The time taken to reduce the matrix is less than 2 s for all instances and is thus negligible compared to solving the optimization problem for all but the smallest cases.

6 Conclusion

A method for determining optimal bids to the day-ahead market for a price-taking hydropower producer has been presented. The method builds on a model that finds the optimal short-term production schedule for hydropower. The new feature is that optimal bid curves are available as a direct result from the optimization.




Fig. 15 The calculation time for our method (*dots*) and the method in [3] (*triangles*) for bid optimization, and production scheduling only (*cross*). For more than 40 scenarios, a time limit of 2 h per iteration of the successive linear programming algorithm (see Sect. 2) is reached. The numbers shown are averages over the 10 days tested

The bids are an expression of marginal costs, which for hydropower are given by the opportunity cost of using water that could have been stored for future production. The cost of production therefore has a temporal element and may be affected by uncertain inflows within the inherent time horizon of the reservoir. For reservoirs that are small compared to the inflow and/or the capacity of connected plants, inflow uncertainty may be important when determining the supply curves. We have shown by examples that this is in fact the case. If the model is to be used for bidding in the industry, the sensitivity of the marginal costs to uncertainty in inflow should be assessed and included where it is important.

To get a bid matrix that complies with the current market rules set by Nord Pool, a reduction method based on a greedy algorithm is used. This method is not guaranteed to preserve the optimal bid curves, but performs well in our tests with only marginal deviations between the optimized and reduced bid matrixes.

In the future, a quantification of the gain from improved decision support for bidding should be done through a comparison with the methods currently used in the hydropower industry. Lastly, the computational burden of the bidding model in its current form is too excessive for operational use and could be improved by algorithmic refinements.

References

1. Fosso, O.B., Gjelsvik, A., Haugstad, A., Mo, B., Wangenstein, I.: Generation scheduling in a deregulated system. The Norwegian case. *IEEE Trans. Power Syst.* **14**(1), 75–81 (1999)
2. Pereira, M.V.F., Pinto, L.M.V.G.: Stochastic optimization of a multireservoir hydroelectric system: a decomposition approach. *Water Resour. Res.* **21**(6), 779–892 (1985)
3. Fleten, S.-E., Kristoffersen, T.K.: Stochastic programming for optimizing bidding strategies of a Nordic hydropower producer. *Eur. J. Oper. Res.* **181**(2), 916–928 (2007)
4. Pritchard, G., Zakeri, G.: Market offering strategies for hydro-electric generators. *Oper. Res.* **51**, 602612 (2003)

5. Pritchard, G., Philpott, A.B., Neame, P.J.: Hydroelectric reservoir optimization in a pool market. *Math. Program.* **103**(3), 445–461 (2005)
6. Fleten, S.-E., Kristoffersen, T.K.: Short-term hydropower production planning by stochastic programming. *Comput. Oper. Res.* **35**, 2656–2671 (2008)
7. Séguin, S., Fleten, S.-E., Côté, P., Pichler, A., Audet, C.: Stochastic short-term hydropower planning with inflow scenario trees. Technical Report G-2015-97. GERAD, Montreal, Canada (2015)
8. Belsnes, M.M., Wolfgang, O., Follestad, T., Aasgård, E.K.: Applying successive linear programming for stochastic short-term hydropower optimization. *Electr. Power Syst. Res.* **130**, 167–180 (2016)
9. Padhy, N.P.: Unit commitment—a bibliographical survey. *IEEE Trans. Power Syst.* **19**(2), 1196–1205 (2004)
10. Steeger, G., Barroso, L.A., Rebennack, S.: Optimal bidding strategies for hydro-electric producers: a literature survey. *IEEE Trans. Power Syst.* **29**(4), 1758–1766 (2014)
11. Fosso, O.B., Belsnes, M.M.: Short-term hydro scheduling in a liberalized power system. *Int. Conf. Power Syst. Technol.* **2**, 1321–1326 (2004)
12. SHOP info site, [online]. https://www.sintef.no/en/sintef-energy/softwareenglish/allsoftware_english/shop/. Accessed 21 Feb 2016
13. Birge, J.R., Louveaux, F.: *Introduction to Stochastic Programming*. Springer, New York (1997)
14. Gjelsvik, A., Belsnes, M.M., Haugstad, A.: An algorithm for stochastic medium-term hydrothermal scheduling under spot price uncertainty. In: *Proceedings 13th Power System Computation Conference*, Trondheim, 28 June–2 July 1999
15. Fodstad, M., Henden, A.L., Helseth, A.: Hydropower scheduling in day-ahead and balancing markets. In: *12th International Conference on the European Energy Market* (2015)
16. Plazas, M.A., Conejo, A.J., Prieto, F.J.: Multimarket optimal bidding for a power producer. *IEEE Trans. Power Syst.* **20**(4), 2041–2049 (2005)
17. Ottesen, S.Ø., Tomasgard, A., Fleten, S.-E.: Prosumer bidding and scheduling in electricity markets. *Energy* **94**, 828–843 (2016)
18. SKM Market Predictor AS P.O. Box 2637, 7414 Trondheim, Norway. Tel: (+47) 73 80 58 00, Fax: (+47) 73 80 58 01. Contact: info@skmenergy.com
19. Bergström, S.: *The HBV model—its structure and applications*. SMHI Reports RH, No. 4, Norrköping (1992)
20. Heitsch, H., Römisch, W.: Scenario reduction algorithms in stochastic programming. *Comput. Optim. Appl.* **24**(2–3), 187–206 (2003)