Evaluation of Methods for Long-term Hydropower Planning in Electricity Market Models

***Elis Nycander*[1]*, *Erik Böhlmark*[1]**

[1]*Svenska kraftnät, Stockholm, Sweden*
**E-mail: elis.nycander@svk.se*

Keywords: ELECTRICITY MARKET MODELS, HYDROPOWER, WATER VALUES, BID3, EMPS

Abstract

This paper evaluates three different methods for long-term planning of hydro dominated power systems. Two established electricity market models using water values, EMPS and BID3, are compared, as well as a look-ahead dispatch which relies on the average optimal reservoir curves computed from a long-term steady-state optimization of the system. All three methods are evaluated using out-of-sample simulations using 100 sequential weather years, to compare the resulting dispatch of the system. It is found that the proposed look-ahead dispatch comes quite close to the established models in terms of average electricity costs. Of the evaluated models, BID3 gives the best results, both in terms of average electricity costs and considering that it does not require any load shedding in the evaluations, unlike EMPS and the look-ahead dispatch. However, the results from all three models can be significantly improved by calibration. While BID3 is a bit too risk averse using the initial settings, the other models have too little risk aversion, which results in significant load shedding.

1 Introduction

Operation planning of power systems with large amounts of hydro production is a dynamic stochastic optimization problem due to the uncertain nature of future inflow and the ability of reservoirs to store water for future time periods. Accordingly, many different methods have been developed to deal with hydropower uncertainty in electricity market models [1–3]. Most models use water values which are calculated from a long-term optimization of the given system, often with lower resolution than the detailed dispach simulation [1]. Essentially, the water value is a state dependent variable, which gives the expected marginal value of the last energy unit stored in the reservoirs for a given system state, in terms of the current reservoir levels and time of year. In the Nordic countries, the water value curves follow a distinct annual pattern. The reservoirs are emptied during the winter and reach their lowest level just before the spring flood. The reservoirs should be emptied enough to accommodate the upcoming inflow but not too much to risk a shortage of water, and this information is incorporated into the water values which are then used in the short-term optimization.

Established models that make use of water values for optimizing hydropower production include EMPS [4–6] and BID3 [7, 8], which are the two electricity market models most commonly used by transmission system operators (TSO) and large hydro producers in the Nordics. The methods used for the water value calculation differ between these models and is an important component of the models due to the importance of hydropower in the Nordic system. However, to our knowledge no comprehensive comparison the of the water values generated by different models exist. Thus, there is no information available for potential users of these models about how well they deal with the stochasticity of hydropower production.

In this paper we present a comparison of the water values generated by the EMPS and BID3 models, by testing how well the water values perform when used for planning the power system over an extended time period. We also include a comparison of the water-value approach of the two models with a look-ahead dispatch that uses information from the optimal long-term reservoir curve derived from a steady-state optimization of the system. All approaches are evaluated using out-of-sample evaluation for 100 sequential inflow years, meaning that the performance of the models is tested using a different set of scenarios than that used to compute the water values/optimal reservoir curves. The test system used is a fictitious system with properties similar to the Swedish test system, with inflows generated from a statistical model to allow out-of-sample evaluation.

The results show that BID3 gives the lowest costs of the models. Since it is the most risk averse it avoids load shedding in the evaluation, even before doing any calibration of the water value curves. On the other hand, both EMPS and the look-ahead dispatch give load shedding in some scenarios which increases the average electricity costs. By calibrating the water value curves for EMPS and the steady-state reservoir curve used in the look-ahead dispatch the load shedding can be decreased significantly and the average costs reduced. The costs from BID3 can also be reduced by decreasing the water values for low reservoir levels, thereby making the water values less risk averse. Hence, an important conclusion is that all three models need some calibration to give results closer to the optimal values.

An advantage of the look-ahead dispatch is that it does not require any specialized algorithms for stochastic dynamic programming (SDP) to compute the water value curves. Thus, this method may have potential for being incorporated into electricity market models that include significant amounts of short-term flexibility, such as batteries and hydrogen storages, which can be difficult to include in SDP-based algorithms. Combining stochastic hydropower modelling with such short-term flexiblity will become more important for electricity market models with increasing amounts of renewable energy and energy storages in power systems.

2 Test System

The test system used has two price areas, North and South, with an annual demand as shown in Table 1 and Fig. 1. The generators in the system are shown in Table 2.

Table 1 Price areas in test system.

Area	Profile	Annual demand (TWh)	Average demand (MWh)	Peak demand (MWh)
North	SE2	0.8760	100	199
South	SE3	3.5040	400	683

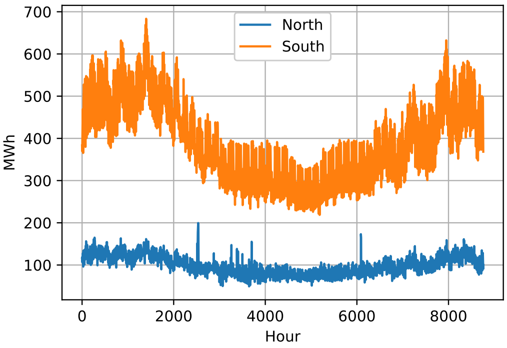

Fig. 1. Hourly demand for test system.

Table 2 Generators in test sytem.

Name	Type	Gen. capacity (MW)	Marg. cost (EUR /MWh)	Reservoir capacity (TWh)	Avg. inflow (TWh)
GT1_N	Thermal	50	50	-	-
Hydro_N	Hydro	300	0	0.5840	1.7520
Nuclear_S	Thermal	200	25	-	-
Hydro_S	Hydro	200	0	0.1752	0.8760
GT1_S	Thermal	200	100	-	-

The system has been designed to have similar characteristics as the Swedish power system, with low load and significant hydro resources in North. Thus the average demand in North is 100 MWh/h, but the average hydro inflow is 200 MWh/h, meaning that a bit more than 50% of the hydro generation must be exported to South through a 150 MW interconnector. South has an average load of 400 MWh/h, and an average hydro inflow of 100 MWh/h. Additionally, it has a 200 MW nuclear generator with a low cost of 25 EUR/MWh, and a 200 MW gas turbine with a high cost of 100 EUR/MWh. North has a 50 MW gas turbine with a cost of 50 EUR/MWh.

The total generation capacity in the system is 950 MW and the maximum hourly load is 844 MWh. South has a generation capacity of 600 MW combined with an import capacity of 150 MW and a maximum load of 683 MW.

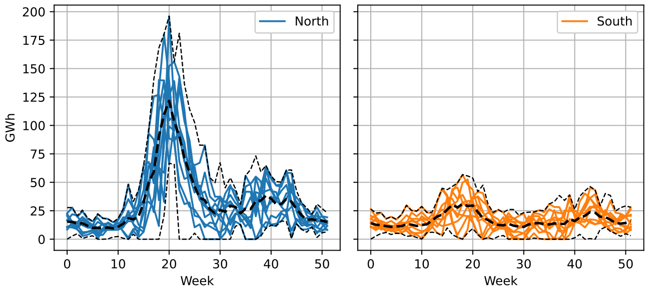

Fig. 2 Weekly inflows for test system. 10 out of the 35 scenarios used for the water value calculation are shown, together with the average of the 35 scenarios, as well as the weekly maximum and minimum values.

Fig. 2 shows the inflow scenarios used to compute the water values/optimal reservoir curves. These have been generated from a model that was obtained by fitting separate multivariate normal distributions to historical inflow data for the price areas SE2 and SE3. The model thus considers correlation between inflows for different weeks for both price areas. In order to do out-of-sample validation of the different methods it was necessary to use synthetic inflows. First 35 scenarios were used to compute the water values/optimal reservoir curves using the different methods, and then 100 different scenarios generated from the same statistical model were used for the evaluation. Note that the demand profiles are the same for every inflow scenario, i.e., no uncertainties apart from the inflows are modelled in the test system.

3 Water Value Calculation

Both BID3 and EMPS use implementations of stochastic dynamic programming (SDP) to compute the water values which are then used in the weekly dispatch simulations. EMPS uses backward induction to compute the value of stored water for a discrete set of reservoir levels in all time periods, starting with the last time period [6]. To reduce the complexity of the problem the water values are calculated independently for each area, even if the water values in one area are dependent on the reservoir levels in all other areas of the system. The exchange between different areas is then adjusted iteratively so that the import/export assumptions of different areas are consistent with each other.

BID3 uses a version of SDP known as constructive dual dynamic programming (CDDP) [2]. This first involves computing the amount of water that would be used/demanded in each time period as a function of the water value, referred to as a ”demand curve for release”. These demand curves are then used to compute the water values using backward induction starting from the last time period, in a somewhat similar manner as in EMPS. Unlike EMPS, BID3 computes water value curves which are dependent on the reservoir content of the other reservoirs in the system [2]. This is done using an approach that divides the system into two reservoir regions, where one is the region for which water values are being computed and the other is an aggregation of other reservoirs that will affect the water values. Thus, for a two area system like the one used in this paper the water value algorithm used by BID3 can give optimal results, while for a system consisting of more than two reservoir regions some suboptimality may arise from the grouping of several price areas into one reservoir region.

Thus, both EMPS and, in particular, BID3, compute the water values for areas that represent a relatively large geographical region, with all hydro plants in that area aggregated into a single hydropower equivalent. Computing water values for individual hydropower plants is possible, but requires other techniques such as stochastic dual dynamic programming (SDDP), that can avoid discretization of the state space and the very large size of the resulting optimization problems [3, 9].

Fig. 3 and Fig. 4 show the water value curves for EMPS and BID3, for North, and South, respectively. Note that for BID3, the water values are also dependent on the reservoir level in the other reservoir. The water value curves shown in the figures correspond to the curves when the level of the other reservoir follows the average reservoir curves obtained from the steady-state optimization, shown in Fig. 6.

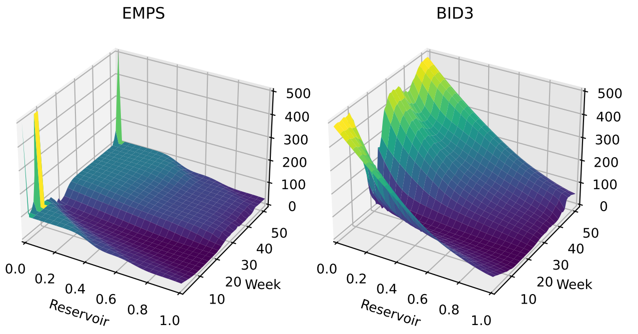

Fig. 3 Water values for North. The values from EMPS have been capped at 500 EUR/MWh to get consistent coloring.

The water values shown in Fig. 3 and Fig. 4 were obtained without doing any calibration to improve the results in the evaluation. However, as described in Section 5, it turned out that some calibration could improve the results from both models. For BID3, the water values were too risk-averse, leading to higher dispatch costs for the thermal generators in the system, while in EMPS the water values were too low, leading to some load shedding in the evaluation.

Table 3 shows the parameters that were changed when calibrating the water values. For BID3 an important setting influencing the water values are the user-defined water value bidding steps, which define the price steps for which the water value curves are computed. The highest bidding step puts a cap on the obtained water values in BID3, and thus reducing the highest step from 500 EUR/MWh to 150 EUR/MWh reduced the water values, thus making the reservoir curves less conservative and lowering operation costs. The spacing of the steps was set to 10 EUR/MWh, to get a total of 16 steps.

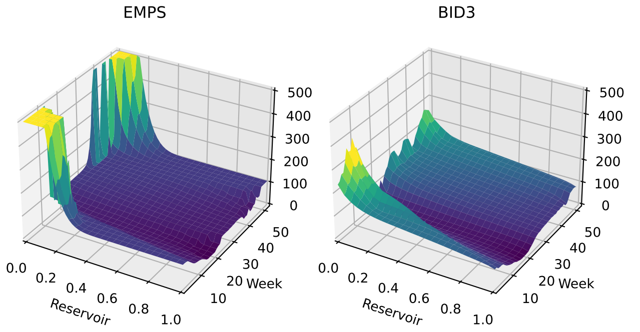

Fig. 4 Water values for South. The values from EMPS have been capped at 500 EUR/MWh to get consistent coloring.

For EMPS, using the default water values resulted in significant load shedding in the evaluation, since the reservoirs were run too low. An important parameter affecting the level of the reservoir curves in EMPS is the so called feedback factor. For each area, the feedback factor controls how large the demand used in the water value calculation is in relation to the hydro production. The default value for North was $0.50387 \approx 0.5$, since the annual inflow is twice the size of the annual demand. Increasing the feedback factor in North will increase the demand seen in the water value calculation, and hence increase the computed water values. The value of 0.7 that was used was the value that gave the lowest total costs in the evaluation, when testing values in the range [0.55, 0.8], with a step size of 0.5.

Note that there can still be further room for improving the evaluation results by using different values for the parameters in Table 3, or by changing other model settings. For example, in EMPS it would also be possible to adjust the feedback factor in South, and several other calibration parameters exist.

Table 3 Water value calibration parameters.

Model	Parameter	Initial value	Calibrated value
BID3	Water value bidding steps	[0, 20, 30, 35, 30, 45, 50, 55, 60, 65, 70, 90, 140, 300, 500]	[0, 10, 20, 30 40, 50, 60, 70, 80, 90, 100, 110, 120, 130, 140, 150]
EMPS	Feedback factor (North)	0.50387	0.7

Fig. 5 shows the calibrated water value curves for North. Compared to the initial results, the water value curves of from the two models are now much more similar. It should be noted that using these curves improved the results for both models, compared to the initial settings. There are still notable differences between the curves from EMPS and BID3. For EMPS the water value curve is relatively flat at a level of 100 EUR/MWh for reservoir levels between 20-55%, while for BID3 there is a continuous decline over the whole reservoir range. For very low reservoir levels the water values from EMPS are much higher, up to 2400 EUR/MWh, while the water values in BID3 are effectively capped at the maximum water value bidding step of 150 EUR/MWh. Comparing the calibrated curves in Fig. 5 with the original curves in Fig. 3 we observe that for EMPS the effect of the calibration has been to extend the flat part of the curve at 100 EUR/MWh to higher reservoir levels, while for BID3 the whole curve has been pushed down from around 500 EUR/MWh to around 150 EUR/MWh for lower reservoir levels.

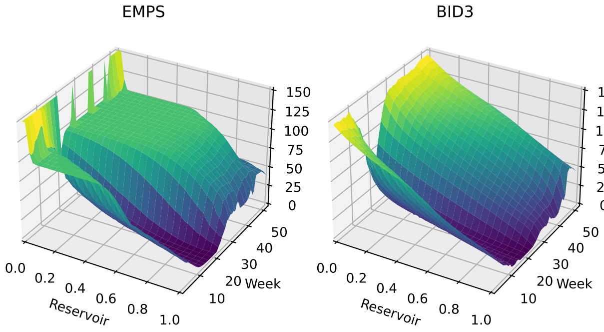

Fig. 5 Water values for North after using the calibrated settings in Table 3. The values from EMPS have been capped at 150 EUR/MWh to get consistent coloring.

4 Look-ahead Dispatch

As an alternative to using water values for scheduling the hydropower, we suggest a method that is based upon finding the optimal average yearly reservoir curve for all reservoirs. The optimal reservoir curves are then used to set the final reservoir levels in the weekly dispatch problems, which are solved for a time period stretching some weeks beyond the week under consideration, which we call a look-ahead dispatch.

To find the optimal reservoir curves, the optimal dispatch problem is solved for all 35 inflow years in series. This means that the final reservoir level in the first year is equal to the initial reservoir level in the second year, and so on. In the same manner, the final reservoir level in the last year is equal to the initial reservoir level in the first year, so that we get a cyclic sequence of whether years, which we call the ”steady-state” problem. This means that the model is completely free to find the optimal allocation of water throughout the modelled years, without having to impose any exogenous reservoir level requirements. We choose to solve the steady-state problem using weekly resolution, but it would also be possible to use a higher resolution, e.g., to better model short-term variability in the system. Fig. 6 shows the optimal reservoir curves obtained from the steady-state problem, as well as the average reservoir curves, which are used as input for the look-ahead dispatch. Notice that the reservoir curve in the first year is a continuation of the reservoir curve for the last year.

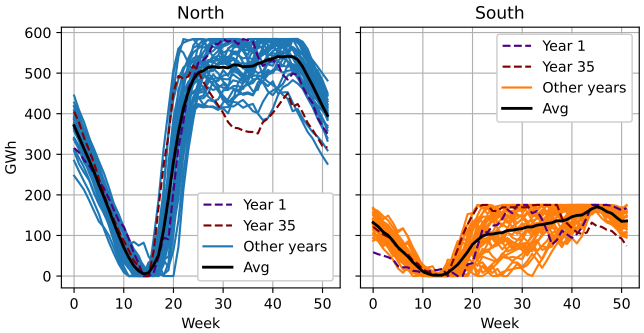

Fig. 6 Optimal reservoir curves obtained from the steady-state problem.

In the look-ahead dispatch we solve the dispatch problem for the week under consideration as well as 4 weeks further into the future. Since only the inflow during the current week is assumed to be known, the inflow during the 4 additional weeks is assumed to be the average inflow during the 35 weather years for which data is available. The final reservoir level is set by assuming that we move 50% closer to the average steady-state reservoir curve during the period under consideration. If, e.g., the current reservoir level is 100 GWh below the steady-state curve, then the final reservoir level will be 50 GWh below the steady-state curve. In this way, when the reservoir values are low less water will be used than normal, and the opposite if reservoir levels are high, giving a similar behaviour as when using water values to plan the hydro production.

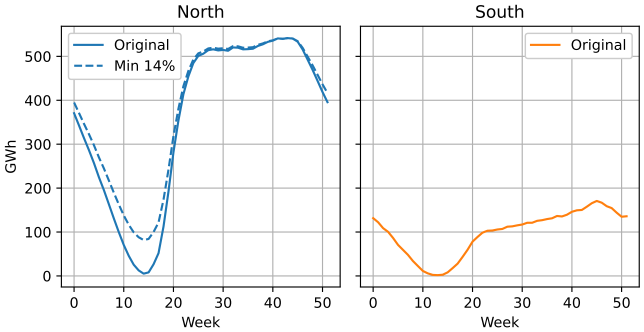

Fig. 7 Steady state reservoir curves rescaled for North to a minimum of 14% of the reservoir size.

The look-ahead dispatch can also be calibrated to give improved results. Most importantly, the reservoir curves obtained from the steady-state optimization assume perfect foresight over the whole optimization period, which allows the reservoirs to be emptied more than when planning under uncertainty. Hence, it will be necessary to shift up the steady-state curves used in the look-ahead dispatch to consider the effect of this uncertainty. A simple way is to rescale the whole curve while keeping the maximum value fixed, and increasing the minimum point of the curve to a given level, as shown in Fig. 7. Here, the calibration was only applied to North, while testing minimum values for the steady-state curve of 2%-26% with a step size of 4%. It was found that using 14% as the minimum reservoir level for North gave the lowest costs in the evaluation, and hence this value was used.

5 Evaluation of Models

To evaluate the different methods for planning the system we use out-of-sample simulations with 100 inflow years generated by the same model as in the water value calculation. For EMPS and BID3, the water values obtained were used to decide the value of the stored water at the end of each week, while for the look-ahead dispatch (LOA), the final reservoir level at the end of the dispatch was set considering the steady-state reservoir curve as described in Section 4. However, note that all of the models have the same information available when solving a particular week, namely the inflow for that week, as well as the inflow of the 35 weather years used for the water values.

Fig. 8 shows the average electricity cost over the 100 simulated years using the default model settings. For each model two different values are shown: the electricity cost from the model itself, as well as the electricity cost obtained when simulating each year using perfect foresight (PF). When simulating with perfect foresight, the initial and final reservoir levels for each year are taken from the results when running the specific model, but the whole year is optimized with perfect foresight, as opposed to simulating one week at a time. The total cost shown is divided into the generation costs, i.e., the fuel costs of the thermal plants, as well as the loadshed costs, which arise for those hours and scenarios when load shedding is required.

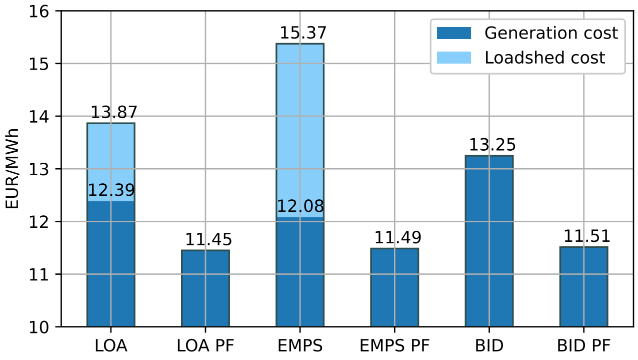

Fig. 8 Average electricity costs for the different models with default settings. The cost shown is the average electricity production cost over the 100 inflow years used for out-of-sample evaluation.

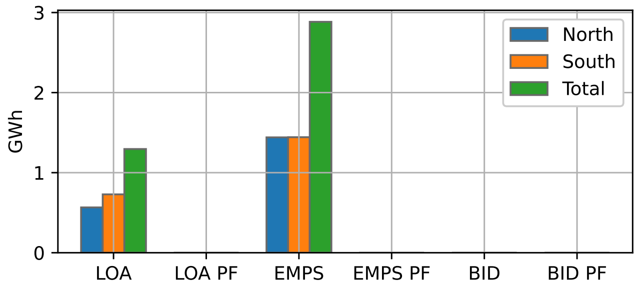

Fig. 9 Average yearly load shedding for models with default settings.

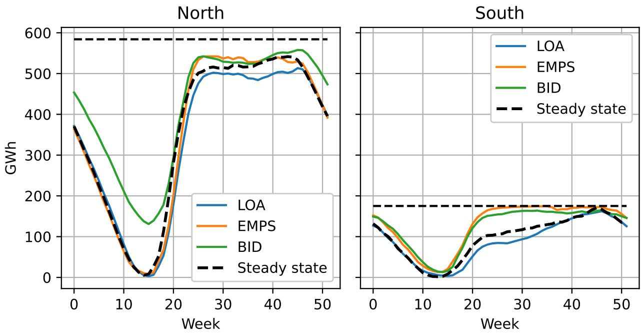

Fig. 10 Average reservoir levels for models with default settings.

Without calibrating the models, BID3 has the lowest costs followed by LOA and EMPS. For both LOA and EMPS, significant costs arise due to load shedding. Notice that load shedding never occurs in the perfect foresight evaluations, since the use of the reservoirs can be optimized over the year to avoid any load shedding. Fig. 9 shows the average yearly load shedding for the models. The total load shedding is less than 3 GWh or 0.07% of the total annual demand for both EMPS and LOA, but still has a large impact on the average costs due to the high value of lost load, which is 5000 EUR/MWh. Fig. 10 shows the average reservoir curves for the different models. It can be seen that LOA and EMPS run the reservoir in North on significantly lower levels during the winter, which results in the load shedding seen in the evaluation.

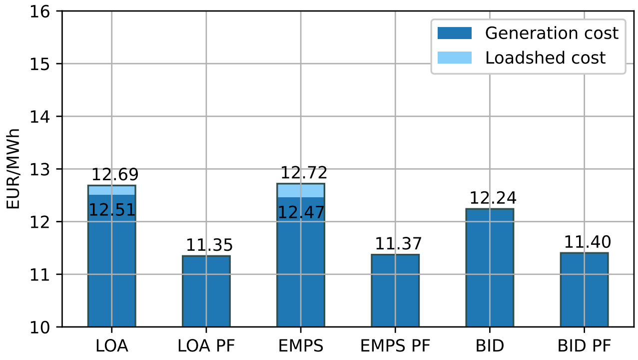

Fig. 11 Average electricity costs for the different models with calibrated settings.

Fig. 11 shows the average electricity costs after calibrating the models. The difference between the models is now much smaller, though the order of performance is the same, with BID3 giving the lowest total costs. Calibrating the models also reduces the cost with perfect foresight for all models, which means that the allocation of water between years is also improved. However, the main improvement is in the cost for the weekly evaluation, meaning that most of the improvement results from better allocation of water within a year. For both LOA and EMPS, there is a large decrease in the costs from load

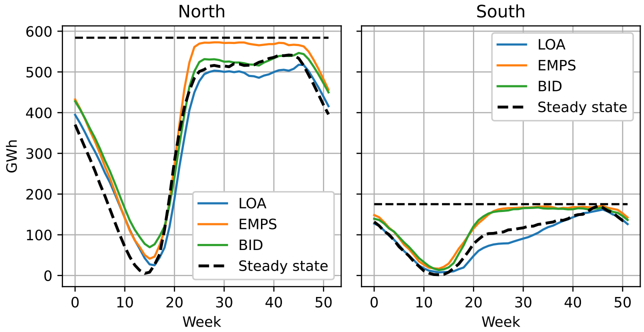

Fig. 12 Average reservoir levels for the different models after calibration.

shedding, along with an increase in the generation costs. However, the decrease in the cost for load shedding is large enough to offset the increase in generation costs, giving a more optimal dispatch of the system. Note that for all calibration parameters tested there was always some load shedding from LOA and EMPS. One reason for this could be that it would be required to calibrate the parameters for South in order to completely remove the load shedding. For BID3, there is a decrease in the electricity cost of around 1 EUR/MWh after calibration, due to decreased generation costs. Notably, BID3 gives lower generation costs than LOA and EMPS, even if the cost for load shedding is neglected. It should be noted that, as expalined in Section 3, the two-area system allows the algorithm in BID3 to obtain optimal water values, which helps explain the higher performance of this method.

Fig. 12 shows the average reservoir curves after calibration. Now the reservoir curves for North align more closely during the winter, and all lie above the steady-state curve. This is to be expected, because the steady-state curve is obtained by optimizing all years using perfect foresight, and can thus afford to empty the reservoirs further compared to when doing the weekly optimization with uncertainty.

6 Conclusion and Discussion

We have compared three different methods for scheduling a hydro-based power system under uncertain reservoir inflows. The two commercial electricity market models, BID3 and EMPS, rely on water values computed using SDP to determine the value of water stored in reservoirs for each weekly scheduling problem. On the other hand, the proposed look-ahead dispatch uses the average steady-state reservoir curves combined with heuristics to set the final reservoir levels.

Of the three models, BID3 gives the lowest total system operating costs, while the costs from EMPS and the look-ahead dispatch are increased due to load shedding required in the evaluation. By calibrating the models, the load shedding from EMPS and the look-ahead dispatch can be reduced significantly, giving lower average electricity costs. The costs from BID3 can also be reduced, since the initial settings give too high water values which makes the model too risk averse. After calibration, the costs and reservoir curves from the different models align more closely, although BID3 still gives the lowest generation costs, also when neglecting the load shedding from the other models. It should be noted that the results obtained for the test system used in this paper do not necessarily extend to larger more realistic test systems. Further studies would be needed to determine the benefits of the different methods for systems with more price areas, and to further fine tune the calibration of the methods. One challenge of doing such studies is that the different data structures and interfaces of EMPS and BID3 makes the implementation of identical test systems in both models relatively time consuming.

Compared to the methods using water value calculations based on SDP, the look-ahead dispatch may have some advantages for modelling future electricity markets where other types of flexibility such as energy storage and demand side response become increasingly important. Since the steady-state reservoir curve used in look-ahead dispatch is computed using a deterministic scheduling problem, this problem can easily be made to incorporate other sources of flexibility as well. Also, the problem may be solved using a higher time resolution compared to the weekly resolution used here, to better model short-term flexibility. This will increase the computational burden of solving the problem, but avoid the need for any methodological changes in the method.

There is also potential for improving the scaling of the steady-state curve and the heuristics used in the look-ahead dispatch. For example, the rate at which the dispatch will try to move towards the steady-state reservoir curve can be varied depending on annual inflow patterns, and also the rate could be varied for positive and negative deviations from the steady-state curve, reflecting, e.g., that emptying reservoirs more than the steady-state curve during the winter can be costly as it increases the chance of load shedding.

The look-ahead dispatch could also be used for systems with more detailed hydropower modelling, providing optimal reservoir curves for individual hydropower plants. For these reasons, the look-ahead dispatch may provide a practical way of incorporating hydropower planning in electricity market models. Further research will be needed to test the approach on real systems, including other sources of flexibility and more detailed hydropower modelling, and for further refining the heuristics of the method.

7 Bibliography

[1] Charlotta Ahlfors, Mikael Amelin, *Weekly planning of hydropower in systems with large volumes and varying power generation: A literature review*, in 2021 IEEE Madrid PowerTech, 2021, pp. 1-7.

[2] E. Grant Read, Magnus Hindsberger, *Constructive Dual DP for Reservoir Optimization*, in *Handbook of Power Systems I*, Springer Berlin Heidelberg, 2010, pp. 3-32.

[3] Anders Gjelsvik, Birger Mo, Arne 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*, Springer Berlin Heidelberg, 2010, pp. 33-55.

[4] Nils Flatabø, Arne Haugstad, Birger Mo, Olav Fosso, *Short-term and Medium-term Generation Scheduling in the Norwegian Hydro System under a Competitive Power Market Structure*, in VIII SEPOPE, Brazilia, 2002.

[5] Arild Helseth, G. Warland, Birger Mo, *Long-term hydro-thermal scheduling including network constraints*, in *7th International Conference on the European Energy Market, EEM 2010*, 2010, pp. 1-6.

[6] Ove Wolfgang, Hans Ivar Skjelbred, Magnus Korpås, *Evaluating North Sea grid alternatives under EU's RES-E targets for 2020*, SINTEF Energy Research, TR A7325, 2013.

[7] Keith Bell, Ian Staffell, *Audit of the BID3 Pan European Market Model for National Grid*, University of Strathclyde, Warwick, 2016.

[8] John Perkins, Gareth Davies, *Benefits of Long Duration Storage: Methodological Annex*, BEIS Research Paper Number 2022/019, 2022.

[9] M.V.F. Pereira, *Optimal stochastic operations scheduling of large hydroelectric systems*, International Journal of Electrical Power & Energy Systems, vol. 11, no. 3, pp. 161-169, 1989.