AN ALGORITHM FOR STOCHASTIC MEDIUM-TERM HYDROTHERMAL SCHEDULING UNDER SPOT PRICE UNCERTAINTY

A. Gjelsvik, M. M. Belsnes and A. Haugstad
Sintef Energy Research (formerly EFI)
Trondheim, Norway

Anders.Gjelsvik@energy.sintef.no

Keywords: Energy economics, hydro scheduling, stochastic programming

ABSTRACT

In a deregulated system, the income of a power company will strongly depend on the spot price, which may vary considerably in a system dominated by hydropower, like the Scandinavian. The aim of the present paper is to present a new algorithm that can deal with spot price uncertainty, for use in stochastic mid-term scheduling. The algorithm is a combination of stochastic dual dynamic programming and stochastic dynamic programming. The algorithm may have applications in small to medium-sized power companies.

1. INTRODUCTION

The background for this work is that in a deregulated system, the income of a power company will strongly depend on the spot price in the system, which may vary considerably and cannot be predicted accurately. Therefore price uncertainty should be taken into account in scheduling. In the present paper, we describe an algorithm to take price uncertainty into account in medium-term hydrothermal scheduling. The algorithm is a combination of so-called stochastic dual dynamic programming (SDDP) [1, 2] and ordinary stochastic dynamic programming.

We consider a small or medium-sized company, so that its decisions do not influence the spot price. Thus, we can regard the spot price as an externally given function. This assumption is valid for many Norwegian power companies. It is assumed that the objective of the scheduling is profit maximisation. The scheduling results are mainly used in setting end-point marginal values for short-term scheduling, and also in maintenance scheduling.

The spot price depends on several factors, but in the Scandinavian system a main variable is the total energy stored in water reservoirs throughout the "global" system of Scandinavia. Based on this, forecasts can be obtained.

We take price fluctuations into account by including in the scheduling model a very simple stochastic model for the spot price. The parameters of this model are derived from price forecasts.

The main goal of the present paper is to present the solution algorithm that we have developed. We also give a small numerical example. An application of this algorithm has been described in [3].

2. SYSTEM MODEL

We here give a brief description of the system model used. A finite time horizon is used, usually one to three years ahead. The study period is divided in discrete time steps numbered $1, \ldots, T$, usually of length one week. At the horizon $t = T$ end conditions must be given, for instance as marginal values of stored water.

2.1 Power station model

For a power station we assume that

$$P = f(q)h/h_0 \tag{1}$$

where $f(q)$ is a piecewise linear function, specific for each power station; P is the energy generated from release q, h is the water head, and h_0 a nominal reference head.

The algorithm to be described cannot deal directly with variable head; the head correction factor in (1) must be applied with estimated values of h. In many Norwegian power stations this is a fair approximation. We believe that optimising directly with variable head may lead to nonconvex problems.

Thermal generation is modeled as a set of buying options, each with a fixed marginal cost.

2.2 Reservoirs and inflow

At the end of a given time step the system can be characterised by its state vector x_t, which typically contains the contents of all reservoirs. There may also be other states in x_t, such as inflow states [1, 4], but these will not be explicitly written out here. The price state will be dealt with separately. Further, we let u_t denote a vector of decisions for time step t, typically containing water releases, overflows, thermal generation and transactions outside the spot market. We write the water balances and other state transition equations in the general form:

$$x_t = F_t x_{t-1} + G_t u_t + \xi_t \quad \text{for } t = 1, \ldots, T, \qquad (2)$$

where F_t and G_t are system dependent matrices. The vector ξ_t describes the the inflow.

Since we deal with medium-term scheduling, the inflow must be treated as stochastic, and we approximate the distribution of ξ_t by the discrete probability distribution

$$P(\xi_t = \xi_t^k) = \psi_k, \quad k = 1, \ldots, K. \qquad (3)$$

The parameters in (3) are estimated from historical inflow records. To apply dynamic programming, ξ_t and ξ_{t-1} must be independent for all t. If this is not the case, the dependence must be modeled and included in (2), and ξ_t will then denote the independent stochastic variable of this model. In our implementation we use a multivariable version of the linear autoregressive inflow model in [4].

2.3 Power balances and objective function

We assume that the power system under consideration is within one area with a single energy balance.

We define, for $t = 1, \ldots, T$:

- y_t^+ – Sale to the spot market
- y_t^- – Purchase from the spot market
- p_t – Spot price (weekly average)
- δ_t – Transmission charge
- d_t – Firm power demand
- c_t – Cost vector associated with u_t.

The power balance equation can be written

$$A_t u_t - y_t^+ + y_t^- = d_t \quad \text{for} \quad t = 1, \ldots, T. \qquad (4)$$

The term $A_t u_t$ represents the hydro and thermal generation, corresponding to the piecewise linear power station models. The firm power demand d_t is considered deterministic; it is zero in the case where all generation is sold in the spot market. For short we define a transaction vector $y_t^T = [y_t^+, y_t^-]$. The cost for one realisation in one time step is then:

$$L_t(u_t, y_t) = c_t^T u_t + (p_t + \delta_t) y_t^- - (p_t - \delta_t) y_t^+ \qquad (5)$$

The value of the water remaining in the reservoirs at the horizon must be subtracted from the cost. Let this value be given by a function $\Phi(x_T)$. We estimate $\Phi(x_T)$ from water value computations on an aggregate model in the long-term scheduling process [5].

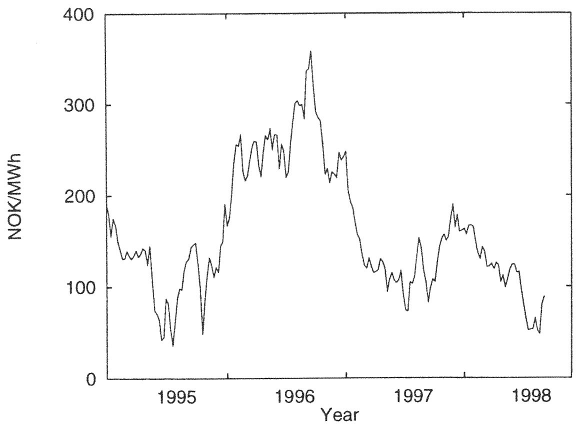

Figure 1: Variation of spot price 1995–1998 (weekly average).

2.4 Summary of model

The model can then be summarised as follows:

$$\min E \left\{ \sum_{t=1}^{T} L_t - \Phi(x_T) \right\} \qquad (6)$$

subject to the constraints

$$x_t = F_t x_{t-1} + G_t u_t + \xi_t \qquad (7)$$

$$A_t u_t + B_t y_t = d_t \qquad (8)$$

$$\underline{x}_t \leq x_t \leq \overline{x}_t \qquad (9)$$

$$\underline{u}_t \leq u_t \leq \overline{u}_t \qquad (10)$$

$$\underline{y}_t \leq y_t \leq \overline{y}_t \qquad (11)$$

for $= 1, \ldots, T$ and x_0 given.

The expectation E is to be taken over both inflow and price. Equation (8) contains the power balance; in general it may also include other constraints that are not coupled in time. A_t and B_t are matrices of suitable dimensions (in the above presentation they would be row vectors). Reservoir limits, equipment ratings etc. are contained in (9) – (11).

3. PRICE MODEL

3.1 Overview

Figure 1 shows the variation of the spot price in Norway during the last few years. In modelling the spot price, we use forecasts obtained by simulations with the so-called EMPS model [6], which is a long-term model covering several years. Analyses of such price series, described in [7] show that:

1. There is a strong serial correlation between the price in a given week and the price in the foregoing week.

2. For a "local" system it is difficult to find a significant correlation between local inflow during a week and the spot price for the same week.

Property 1) means that to apply dynamic programming, which requires that stochastic terms are sequentially uncorrelated, it is necessary to introduce a state space model for the spot price.

Further, it is necessary to establish a joint probability distribution for price and inflow for the next time step, given the present state. Based on property 2) above, we take the stochastic processes for inflow and price as independent of each other, and use the marginal probability distributions for each. A better distribution would be hard to estimate.

On the other hand, if one looks at a time span somewhat longer than one week, the average price of the interval will certainly depend on the accumulated inflow during this period, and this is not taken into account. Therefore, in the algorithm to be described, values from the inflow-price scenarios are used directly in the forward run of the SDDP part of the algorithm, as a heuristic to take this into account.

3.2 A spot price model

It is assumed that a spot price forecast is available, in the form of price scenarios. When the price scenarios have been obtained by the EMPS model, each scenario corresponds to an historical inflow sequence.

We regard the spot price p_t as a state. The price axis is discretised into a set of M points $\zeta_1, \ldots, \zeta_M$, and we use the following discrete Markov model:

$$Pr(p_t = \zeta_j | p_{t-1} = \zeta_i) = \rho_{ij}(t) \quad \text{for all} \quad i, j \qquad (12)$$

where i and j runs from 1 to M. This means that $\rho_{ij}(t)$ is the probability that $p_t = \zeta_j$, given that p_{t-1} was ζ_i, for all i and j.

Note that the price can take only discrete values. This is done in order to simplify the implementation of dynamic programming. This introduces some discretisation error, but we have not tried to estimate this error.

3.3 Fitting the price model

The numerical values of the transition probabilities for price changes between week $t-1$ and week t are established the following way:

First, the price values within each week are grouped in M groups, and ζ_i is taken as the mean value of the i-th group. In this way, the discrete price points are established for each week in the data period. It is recorded which scenarios go into each group in each week, and ρ_{ij} is then estimated as the fraction of the scenarios from the i-th group at time $t-1$ that belong to group j at time t.

4. SOLUTION METHOD

4.1 Overview

We have chosen to work with a combination of stochastic dual dynamic programming and ordinary stochastic dynamic programming. The ordinary SDP part is introduced to take care of the price process, which is modeled as described above. The reservoir and inflow states are treated as continuous variables and dealt with in a way similar to the ordinary SDDP algorithm.

Another way of dealing with the price process would be to describe it by a scenario tree. This tree would become very large, however, so instead, we consider the discrete price model described above. Due to the term $p_t y_t$ in (5) the cost function is a non-convex function in p_t and y_t, and so it is impossible to model the dependence of the expected future cost functions on p_t using hyperplanes the same way as for the state x_t.

As already discussed, the correlation between inflow and price one week ahead is neglected. However, in this extension of the SDDP algorithm, we use a modified approach, in that on the forward run of the algorithm, we use the "observed" inflow-price scenarios. This heuristic can be debated. It is intended to preserve any coupling between inflow and price when averaged over longer periods, but it may lead to gaps between upper and lower cost estimates. This is because on the backward run, we use a Markov model that has been estimated from the price series, and this model may not give exactly the same average as the "observed" scenarios.

4.2 Dynamic programming approach

We now consider a time interval t, with the initial state given by x_{t-1} and $p_{t-1} = p^i_{t-1}$. There are K realisations of the inflow noise ξ_t and for each of these M possible price values p^j_t. We assume here that we learn ξ_t and p^j_t immediately before the decisions for this time step are to be carried out. Let $\alpha_t(x_t|p^j_t)$ be the expected future cost function at the end of time period t, given the system state x_t and p^j_t, that is, the expected cost in going from the given state at the end of time interval t to an allowed final state using an optimal strategy. Applying the Bellman optimality principle, we obtain the recursive equation

$$\alpha_{t-1}(x_{t-1}|p^i_{t-1}) = \sum_{j=1}^{M}\sum_{k=1}^{K} \rho_{ij}\psi_k \min\left[L_t(u_t, y_t) + \alpha_t(x_t|p^j_t)\right] \quad \text{for all } t \text{ and } i, \qquad (13)$$

where the constraints (2) – (11) must be satisfied for each transition. For each possible outcome (ξ^k_t,p^j_t) separate decisions u^{kj}_t, y^{kj}_t are made, and the final state obtained is x^{kj}_t.

In the description of the ordinary SDDP algorithm [1, 2] it is shown that with a linear model, the expected future cost functions are piecewise linear functions of x and can be represented by hyperplanes in the x state space, which also means that these functions are convex. We now show that this is also so in our slightly more complicated case. *Assume* that $\alpha_t(x_t|p_t^j)$ can be represented by hyperplanes; then we proceed to show that this is also the case for $\alpha_{t-1}(x_{t-1}|p_{t-1}^i)$.

We define $\alpha_{t-1}^{kj}(x_{t-1}) = \min[L_t(u_t, y_t) + \alpha_t(x_t|p_t^j)]$ in (13). Under the above assumption on hyperplane representation, (13) then decomposes into single-transition subproblems of the following form:

Given x_{t-1}, $p_{t-1} = p_{t-1}^i$, $p_t = p_t^j$ and $\xi_t = \xi_t^k$, find

$$\alpha_{t-1}^{kj}(x_{t-1}) = \min\left[L_t(u_t, y_t) + \alpha\right] \tag{14}$$

subject to

$$x_t - G_t u_t = F_t x_{t-1} + \xi_t \quad | \quad \pi_t^{kj} \tag{15}$$

$$A_t u_t + B_t y_t = d_t \tag{16}$$

$$\underline{x}_t \leq x_t \leq \overline{x}_t \tag{17}$$

$$\underline{u}_t \leq u_t \leq \overline{u}_t \tag{18}$$

$$\underline{y}_t \leq y_t \leq \overline{y}_t \tag{19}$$

$$\left.\begin{array}{rcl} \alpha + (\mu_t^{j1})^T x_t & \geq & \gamma_t^{j1} \\ & \vdots & \\ \alpha + (\mu_t^{jR})^T x_t & \geq & \gamma_t^{jR} \end{array}\right\} \tag{20}$$

In (20) $\mu_t^{j1}, \ldots, \mu_t^{jR}$ and $\gamma_t^{j1}, \ldots, \gamma_t^{jR}$ give the R hyperplanes that define the expected future cost function at the price point p_t^j. For equation (15) we have indicated the corresponding dual variable π_t^{kj}. We then obtain from (13)

$$\alpha_{t-1}(x_{t-1}|p_{t-1}^i) = \sum_{j=1}^{M}\sum_{k=1}^{K} \rho_{ij}\psi_k \alpha_{t-1}^{kj}(x_{t-1}). \tag{21}$$

We assume that the single-transition subproblem described in (14)–(20) has a feasible solution; this can be ensured by artificial variables. The problem is similar to that of the ordinary SDDP algorithm, and it therefore follows from [1, 2] that $\alpha_{t-1}^{kj}(x_{t-1})$, the objective function of the subproblem, is a convex and piecewise linear function of the right-hand side component x_{t-1}.

We have $\sum_{k=1}^{K}\sum_{j=1}^{M} \psi_k \rho_{ij} = 1$, and ψ_k and ρ_{ij} are nonnegative for all k and i, j. Thus, by (21) $\alpha_{t-1}(x_{t-1}|p_{t-1}^i)$ is a convex combination of convex functions, and therefore convex, and in this case piecewise linear, and therefore it can be represented by hyperplanes. At $t = T$, α is identical to $\Phi(x_t)$, which is piecewise linear and convex, then by induction from t to $t-1$ we conclude that all future cost functions can be properly represented by hyperplanes, usually referred to as *cuts*.

The number R of hyperplanes necessary to represent an expected future cost function, will usually be very large. Therefore, the hyperplanes are built iteratively, as in the ordinary SDDP algorithm.

A solution to the single-transition subproblem (14)–(20) with $x_{t-1} = x_{t-1}^s$, say, contributes to a cut for $\alpha_{t-1}(x_{t-1}|p_{t-1}^i)$ at time $t-1$ through its vector of dual variables. From [1] and duality theory for linear programming we can establish that for general x_{t-1}:

$$\alpha_{t-1}^{kj}(x_{t-1}) \geq \alpha_{t-1}^{kj}(x_{t-1}^s) + \pi_t^{kj} F_t(x_{t-1} - x_{t-1}^s). \tag{22}$$

Averaging according to (13) gives the new cut at time $t-1$:

$$\alpha_{t-1}(x_{t-1}|p_{t-1}^i) \geq \alpha_{t-1}^{is}(x_t^s) + \overline{\pi}_t^{is} F_t(x_{t-1} - x_{t-1}^s), \tag{23}$$

where $\overline{\pi}_t^{is} = \sum_{k=1}^{K}\sum_{j=1}^{M} \rho_{ij}\psi_k \pi_t^{kj}$ and $\alpha_{t-1}^{is} = \sum_{k=1}^{K}\sum_{j=1}^{M} \rho_{ij}\psi_k \alpha_{t-1}^{kj}(x_{t-1}^s)$.

We note that the price of the previous week, p_{t-1}^i, does not enter the subproblem. Therefore, in (13) it is not necessary to solve the subproblem for all M^2 combinations of i and j. One solves for the M different p_t^j; then p_t^i comes in through averaging with the transition probabilities ρ_{ij}.

To visualise the approach, the price dimension used in the new algorithm is shown schematically in Figure 2. At each discrete price point there is a set of cuts representing the expected future cost functions, instead of a single number in table-based dynamic programming.

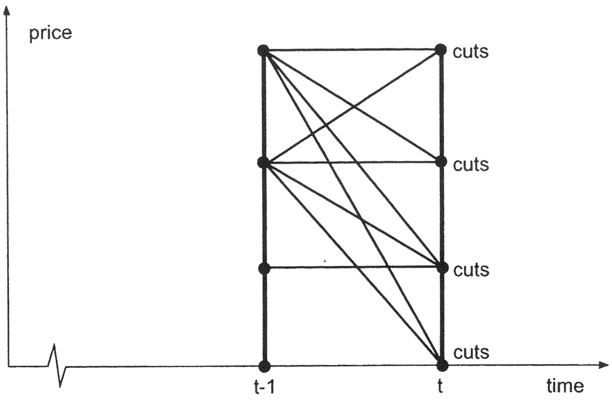

Figure 2: View of the dynamic programming part of the combined approach, in the time-price plane

4.3 Algorithm

One main iteration of the algorithm consists of a forward simulation with the strategy developed in the previous iterations and a backward recursion based on (13) where the strategy is updated by generating more cuts. Since we now use only a subset of all cuts, we shall write $\widehat{\alpha}$ instead of α to show that we deal with approximate future cost functions only. The algorithm is:

1) Initialisation

Set $\overline{J} := \infty$ (a large value). Create a set S^* of scenarios for inflow and price. This can be done by sampling randomly from the inflow distribution. As a heuristic, we here use "observed" scenarios of inflow and price. ξ_0^s and p_0^s are assumed given and common to all scenarios (this means that the price and inflow for the first time interval is considered known). Initialise $R_t = 0$, $t = 1, \ldots, T$.

2) Forward simulation

For each scenario the system is simulated forward in time:

Repeat for $t = 1, \ldots, T$
  Repeat for all $s \in S^*$
    Let $x_{t-1} = x_{t-1}^s$ and $\xi_t = \xi_t^s$.
    Chose p_t as the value p_t^j closest to p_t^s and solve the single-transition problem (14)–(20) giving u_t^s and y_t^s.
    Store x_t^s and $L_t(u_t^s, y_t^s)$.

Compute the operating cost

$$J^s = \sum_{t=1}^{T} L_t(u_t^s, y_t^s) - \Phi(x_T^s) \quad \text{for all } s \in S^*$$

Let p^s be the probability for scenario s, and compute

$$E\{J\} = \sum_{s \in S^*} p^s J^s$$

Update the upper limit for average operating costs:

$$\overline{J} := \min(\overline{J}, E\{J\})$$

3) Backward recursion

Repeat for each time step, $t = T, T-1, \ldots, 1$
  Repeat for $j = 1, \ldots, M$ (loop over p_t^j)
    Let $p_t = p_t^j$
    Repeat for each $x_{t-1}^s, s \in S^*$
      Repeat for each noise value $k = 1, \ldots, K$
        Let $\xi_t = \xi_t^k$ and solve the single-transition problem (14)–(20). Save the optimal objective value $\widehat{\alpha}_{t-1}^{kj}(x_{t-1}^s)$ and dual variables π_t^{kj} for the transition equation (15).
  Repeat for $i = 1, \ldots, M$ (loop over p_{t-1}^i)
    Repeat for each $x_{t-1}^s, s \in S^*$
      Compute

$$\overline{\pi}_t^{is} = \sum_{k=1}^{K} \sum_{j=1}^{M} \rho_{ij} \psi_k \pi_t^{kj}$$

and

$$\widehat{\alpha}_{t-1}^{is} = \sum_{k=1}^{K} \sum_{j=1}^{M} \rho_{ij} \psi_k \widehat{\alpha}_{t-1}^{kj}(x_{t-1}^s)$$

and create a new cut for $\widehat{\alpha}_{t-1}(x_{t-1}|p_{t-1}^i)$ of the kind (20) with

$$\mu_{t-1}^{ir} = -(\overline{\pi}_t^{is})^T F_t$$

$$\gamma_{t-1}^{ir} = \widehat{\alpha}_{t-1}^{is} - (\overline{\pi}_t^{is})^T F_t x_{t-1}^s$$

Let $R_t = R_t + 1$.

At $t = T$ one has instead of (20) a similar representation for $\Phi(x_T)$.

An estimate of the future cost at the beginning of the first interval becomes

$$\widehat{\alpha}_0 = \min(L_t(u_t, y_t) + \widehat{\alpha}_1)$$

where $\widehat{\alpha}_1$ refers to the value at the end of the first interval (in the first interval a single realisation was assumed). This is a lower bound since only a subset of the cuts are created at each stage.

4) Control of convergence.

$\underline{J} = \widehat{\alpha}_0$ is a lower limit for expected operating costs J. An estimate of the upper limit is given by $\overline{J}$ from the forward run, and in the end $\underline{J}$ and $\overline{J}$ should meet. Due to the sampling of scenarios, $\overline{J}$ will have a sampling variation, so that $\overline{J} - \underline{J}$ does not reach zero exactly; it is even possible that $\overline{J} < \underline{J}$. A possible criterion is therefore to stop when $| \overline{J} - \underline{J} |$ is comparable to the standard deviation of $\overline{J}$, ensuring that a minimum number of main iterations in the algorithm has been carried out. In practice, though, we often carry out a specified number of iterations.

4.4 Computational issues

Apart from the "outer" dynamic programming treatment of the spot price state, the approach is similar to that of the ordinary SDDP algorithm, and the same computational approach can be used for this part. To solve the single-transition subproblem, we use a relaxation approach, as in [4]. Thus, the LP problems actually solved are quite small. There is a limit to the number of cuts allowed for each of the M price values; after that, cuts are overwritten.

A set of initial scenarios are assumed to be available at the start of the solution process, so that we can start with the backward recursion step of the algorithm (step 3). The inflow loop is put innermost in the algorithm, since this only changes the right-hand side of the single-transition problem. Each problem with a new inflow is started from the basis of the previous one. When the price p_t changes, both the cost row and the set of cuts changes; in this case, an all slack basis is used for start.

5. A SMALL EXAMPLE

Figure 3 shows a small test system. The power station ratings are given in MW and inflows and storage capacities are in million cubic meters. For the case

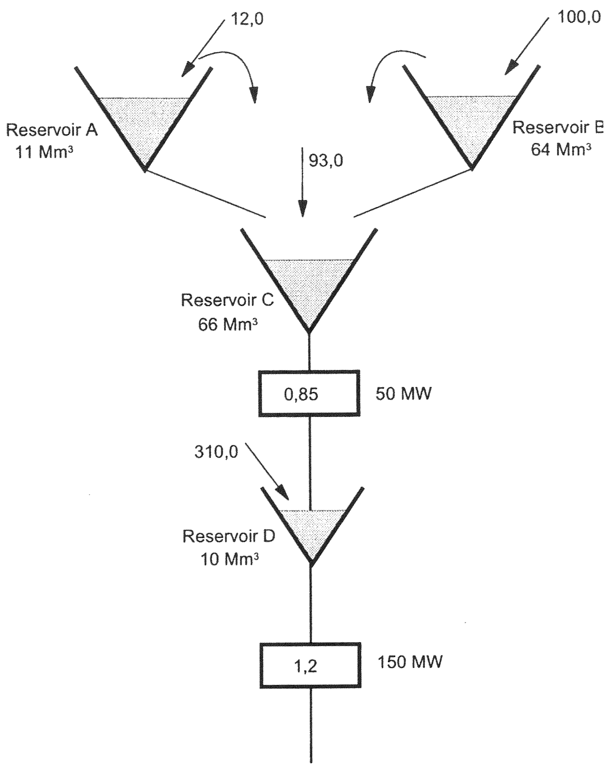

Figure 3: System for test case.

that we show here, the study period is almost three years, beginning in week number 14 in the first year, with reservoirs 25 per cent full. All power is sold in the spot market. 50 price scenarios are used; they are obtained from a run with an EMPS model of the Scandinavian system. The transmission charge δ_t is one per cent of the price p_t for all t. We have carried out computations for $M = 5$ price levels and for $M = 1$.

The results to be used from such a computation are the expected future cost functions at the end of the first week, that is, the cuts describing $\widehat{\alpha}_1(x_1|p_1^j)$ for all j. They can be transferred directly to a short-term scheduling model.

Although it is not used directly, it is interesting to look at the strategy that the algorithm uses over the rest of the study period. We first consider the case $M = 5$. Total storage in GWh from the final forward simulation in step 3 of the algorithm is shown in Figure 4. The curves give the probability distribution of total stored energy in the form of percentiles (10,25,50,75,90). This can be regarded as a "parallel" simulation of the various inflow/price scenarios from the given starting point. The 10 per cent curve in the figure gives rather low storage compared to the mean (note that the curves give percentage points for each week and not trajectories), and this is below what is considered ordinary practise. This is a common experience with several test systems; the reservoirs run relatively low in some cases where prices are high over a longer period.

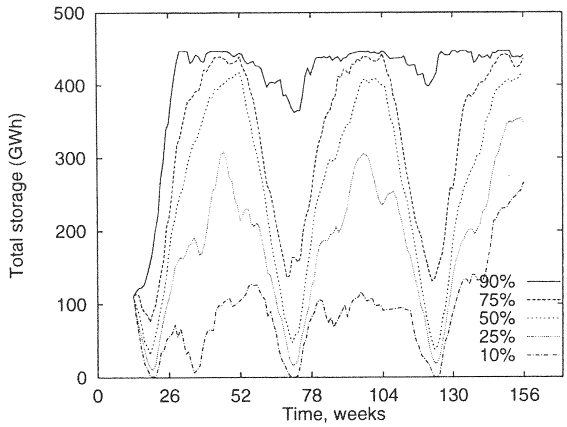

Figure 4: Percentiles for total stored energy, $M = 5$

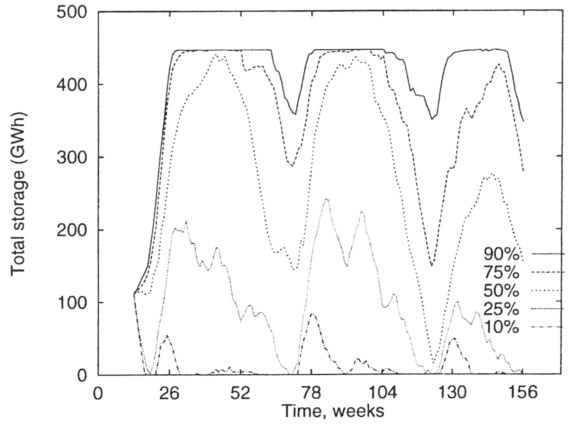

Figure 5: Percentiles for total stored energy, $M = 1$

To indicate what would happen without a price model, we have run the system with a price model with $M = 1$, corresponding to the mean value. The same percentiles as above for this case are displayed in Figure 5. We see that the variation in reservoir content is larger on the low side, and we also note that the operation is significantly different the last year. Also, we find that there is more spillage. Of these two cases, one would say that the use of a price model ($M = 5$) gives the most reasonable operation. However, simulations of this kind are not fully realistic. One reason for this is that the horizon is fixed; we simulate the strategy as obtained standing at the initial week, not considering that the horizon moves and the updating of the price forecast and the strategy that would take place in real operation. Especially in the case of scheduling against a mean price forecast ($M = 1$), updating would be important

6. DISCUSSION

Use of the above algorithm has been implemented as an option for mid-term scheduling in the long-term scheduling model EOPS [8]. It is intended for small to medium-sized systems. The main advantage of such an algorithm (as with the ordinary SDDP algorithm) is that one can optimise with a detailed model of the hydropower system, so that one obtains incremental water values for each reservoir. The algorithm has been tried on several test systems, with sizes ranging from 4 to over 50 reservoirs. Schedules may differ from those of the long-term model, as was exemplified in the case in [3].

The computing time depends on the number of time steps, the number of inflow scenarios, the number K of inflow cases, and the number of discrete points in the price model. For the test case, the computation time was about 15 minutes on a machine running Unix, but for larger systems a typical time is one to two hours.

A special problem is that of constructing the final value function $\Phi(x_T)$. We have used results from an aggregated long-term model of the system for this purpose, so that Φ is a function of total storage. However, this may underestimate the value of water in well regulated reservoirs and overestimate the value in the others. Difficulties arise when there are long-term reservoirs that store water for several years. This could be partially corrected for by imposing more constraints; another way would be to move the horizon far enough away, at the expense of computing time.

Another difficulty is with the fitting of the price model. From "observed" series created from the long-term model only about 50 scenarios are available, and this gives rough estimates of the transition probabilities. This could be improved by generating a large number of synthetic inflow scenarios and obtaining price forecasts from the long-term model for all of them.

The scheduling algorithm is intended for use from a given initial state at a given time. It would be very interesting to evaluate the economic gain resulting from the new algorithm with the stochastic price model; for this purpose its use over a long period (50 years, say) might be simulated, so that results with and without the stochastic price model could be compared. However, this is not possible at present. One reason for that is the long computation time.

7. CONCLUSION

In this paper, we have given an algorithm whereby spot price stochasticity can be taken into account in stochastic scheduling. We believe that the combination of ordinary stochastic dynamic programming and stochastic dual dynamic programming is new. A disadvantage with the method is that the computing time is long; on the other hand, the medium-term scheduling is an application where this can be tolerated.

It is not computationally feasible to carry out the computer simulations necessary to quantitatively give a measure of the improvement in operation that the new algorithm may give, but it seems that the results are positive.

Among the areas for future work are improved price modelling; also a reduction in computing time would be desirable.

ACKNOWLEDGEMENTS

This work was supported by the Norwegian Electricity Federation (EnFO) and the Research Council of Norway (NFR).

REFERENCES

[1] M. V. F. Pereira, "Optimal stochastic operations scheduling of large hydroelectric systems", Electrical Power & Energy Systems, vol. 11, no. 3, pp. 161–169, July 1989.

[2] M. V. F. Pereira and L. M. V. G. Pinto, "Multi-stage stochastic optimization applied to energy planning", Mathematical Programming, vol. 52, pp. 359–375, 1991.

[3] A. Gjelsvik, M. M. Belsnes, and M. Håland, "A case of hydro scheduling with a stochastic price model", in Hydropower '97, E. Broch, D. K. Lysne, N. Flatabø, and E. Helland-Hansen, Eds., Rotterdam, 1997, pp. 211–218, A. A. Balkema.

[4] T. A. Røtting and A. Gjelsvik, "Stochastic dual dynamic programming for seasonal scheduling in the Norwegian power system", IEEE Transactions on Power Systems, vol. 7, no. 1, pp. 273–279, February 1992.

[5] A. Johannesen and N. Flatabø, "Scheduling methods in operation planning of a hydro-dominated power production system", International Journal of Electrical Power & Energy Systems, vol. 11, no. 3, pp. 189–199, July 1989.

[6] O. J. Botnen et al., "Modelling of hydropower scheduling in a national/international context", in Hydropower '92, E. Broch and D. K. Lysne, Eds., Rotterdam, 1992, A. A. Balkema.

[7] O. B. Fosso et al., "Generation scheduling in a deregulated system. The Norwegian case", Presented at the IEEE PES Winter meeting 1998, paper PE-229-PWRS-0-11-1997.

[8] N. Flatabø et al., "Short-term and medium-term generation scheduling in the Norwegian hydro system under a competitive power market structure", Presented at EPSOM '98, Zurich, September 23–25, 1998.