Hydroelectric Reservoir Optimization in a Pool Market

G. Pritchard[*], A.B. Philpott[†], P.J. Neame[‡]

Abstract

For a price-taking generator operating a hydro-electric reservoir in a pool electricity market, the optimal stack to offer in each trading period over a planning horizon can be computed using dynamic programming. However, the market trading period (usually 1 hour or less) may be much shorter than the inherent time scale of the reservoir (often many months). We devise a dynamic programming model for such situations in which each stage represents many trading periods. In this model, the decision made at the beginning of each stage consists of a target mean and variance of the water release in the coming stage. This decomposes the problem into inter-stage and intra-stage subproblems.

1 Introduction

In recent years various forms of electricity markets have emerged throughout the world. In this paper we develop a model for determining the optimal policy of a hydroelectric generator operating in an electricity *pool* market such as those implemented in the Nordic countries, Australia and New Zealand, and some parts of North America.

In an electricity pool market, each generator is required to submit a non-decreasing *supply function* $q = S(p)$, indicating how much power it is willing to generate as a function of the price paid. Usually, the supply function is required to be in the form of an *offer stack*, meaning that it must be a step function. That is, the generator offers several blocks (known as *tranches*) of power to the market, with each tranche having a different asking price.

To determine spot prices for energy, a central dispatching authority then clears the market by systematically accepting the least expensive supply offers until demand is met (allowing for transmission network losses and constraints). The marginal cost of supply at any node of the transmission network (i.e. the shadow price of the energy balance constraint) then defines the spot price of energy at this node.

[*]Department of Statistics, University of Auckland

[†]Department of Engineering Science, University of Auckland

[‡]Department of Mathematical Sciences, University of Technology, Sydney

This process is repeated, with new supply offers, for each *trading period*. The duration of a trading period is fixed, and varies with the particular market design. In the New Zealand market a trading period has a duration of 30 minutes.

Specific details on how the central authority makes its dispatching decisions differ with the market design. For a thorough treatment of the issues surrounding the design of electricity markets see e.g. Chao and Huntington [6]. Details of the New Zealand dispatching model can be found in [2].

The problem addressed in this paper is the construction of optimal supply offers for a generator operating a single hydro-electric reservoir. The reservoir has seasonal storage, meaning that the time required to deplete or replenish it (anything from a few weeks to a year or more) is much longer than the market trading period. The generator is assumed to be a price-taker. That is, the price received for all energy generated – the spot price at the local node of the transmission network – is treated as an exogenous random variable, unaffected by the offers made by this generator. Similarly, water inflows to the reservoir are to be treated as random.

The model we develop is similar in style to the dynamic programming model proposed by Pritchard and Zakeri [17]. In their model each trading period is a stage of the dynamic programming model, and the price in each period is a random variable. The optimization at each stage chooses an offer stack to maximize the sum of the expected revenue from the current stage and the expected future revenue obtained from releasing the remaining water optimally over the decision horizon. This dynamic programming model is suitable for reservoirs with small storage volumes that are quick to replenish. These are typically operated so that the reservoir level cycles over a day or at most a week.

For reservoirs with large volumes, the storage levels cycle over seasons. The Pritchard and Zakeri approach is less suitable for such a reservoir, since a reasonable decision horizon must contain a very large number of trading periods. Not only does this imply an enormous computational effort, but it provides a level of modelling detail (for example optimizing a single half-hour offer to be made months from now) that is unwarranted.

One way to address this issue would be to have stages of variable length. The first few stages would each model a single trading period, and then the length of the stages would gradually be increased. Some of the later stages might each represent several days, or even weeks. This approach has been used successfully in other contexts (see [4]). However, the difficulty with this in the pool-market context is that spot electricity prices are highly volatile, with a large intra-day variation. There will thus be a large error in approximating a whole day (for example) as if it were a single trading period, with a single price. Suppose we generate only in a few highly-priced trading periods during the day (perhaps making substantial revenue therefrom); the single-period approximation might well indicate that we do not generate at all, since the price for this stage – which must attempt to be representative of the entire day – is too low.

It is thus clear that in a model where one stage may represent many trading periods, one should attempt to preserve some representation of the price variability within a stage.

The approach followed in this paper is to split the dynamic programming recursion into two subproblems. The first subproblem (called the *intra-stage* problem) computes the single supply function to offer in every trading period of the current stage so as to give the maximum expected revenue in that stage for a specified mean and variance of water release over the stage. This subproblem is solved parametrically for a range of means and variances. The second subproblem (called the *inter-stage* problem) is to use the values computed from the first problem to choose the mean and variance of water release over the stage to maximize the revenue from the current stage plus the expected revenue from future stages.

This approach is quite advantageous from a dynamic programming point of view. The stages can be of a reasonable length (e.g. weekly), yet there are only two decision variables at each stage – the mean and variance of the amount of water that will be released during the stage. Even with the addition of a few auxiliary state variables to model such things as weather and the "state of the market", it is clear that solution will be computationally feasible.

A key feature of the model is that the detailed planning for each stage is hidden inside an intra-stage subproblem which is decoupled from the current reservoir level. This makes it possible to carry out the intra-stage calculations offline, prior to the main SDP recursion. This is a considerable computational advantage, but there is a price to be paid: no dependence of the immediate payoff on the water level is possible; in particular, hydrostatic head effects cannot be modelled. In many (though not all) hydro-electric systems, such head effects are quite small, so this is an acceptable tradeoff.

The application of dynamic programming for hydro-electric reservoir optimization is not new (see e.g. [19] for an early paper in this area, and related stochastic programming contributions by Pereira and Pinto [15] and Jacobs et al. [13]). Most previous models have focused on situations in which uncertainty is restricted to reservoir inflows, and the optimization chooses reservoir releases so as to minimize the cost of displaced thermal generation. Models of price uncertainty are relatively recent. A dynamic programming approach using price as a state variable is described in a series of papers [8],[9] by Gjelsvik and his coauthors, and a stochastic programming model is explored in the paper [7] by Fleten, Wallace and Ziemba who represent price uncertainty by scenarios and maximize risk-adjusted profit within an asset-liability framework. For a survey of other stochastic programming models in energy, see [20].

A related strand of work looks at the equilibrium behaviour of generators offering generation as players in a Cournot oligopoly. Scott and Read [18] construct a dynamic programming recursion to compute equilibrium strategies for this game in the presence of uncertain reservoir inflows. In a recent paper [5], Bushnell shows using a mixed-complementarity model (with deterministic inflows) that in equilibrium generators can behave strategically by withholding water during periods of high-prices, and increasing generation during low-price periods, an effect that is the opposite to that expected in a model with price-taking behaviour.

Our contribution is to place the dynamic optimization problem within the context of pool markets in which participants choose both reservoir releases and prices in the form of an offer stack, and to provide a methodology for solving the dynamic programming model.

The paper is laid out as follows. In the next section we derive the dynamic programming recursion and demonstrate how it can be split into two subproblems.

Section 3 formulates the intra-stage problem, as a convex quadratic programming problem. Section 4 focuses on the inter-stage problem, and demonstrates under mild conditions that the value functions have appropriate convexity properties that will ensure a global solution to the problem. Finally in section 5 we discuss the implications of relaxing some of the assumptions upon which our algorithm rests.

2 Outline of the model

For definiteness, suppose we wish to model the operation of a single hydro-electric reservoir over a time horizon divided into stages one week long. The essential decision to be made at the beginning of each stage is a tradeoff - how much water to release in the current stage, and how much to save for later stages?

This decision is not as simple as it may appear. The pool-market mechanism means that the quantity of water released is dependent (via a supply function) on the market price of electricity, which from the operator's point of view is a random variable. Moreover, the stage can be further subdivided into market trading periods of duration (say) 30 minutes, each of which may have a different electricity price; to model this, some sort of stochastic process is called for. We are thus faced with the choice of a *release policy* (from a large space of possible policies) for the current stage, which will then give rise to the release of a random total quantity of water during the stage.

However, what is important in long-term planning is not the hour-to-hour detail of the release policy, but the probability distribution of the resulting water release for the stage as a whole. At this level, it is helpful to think of this distribution as being the decision variable for the stage; the details of how the distribution is to be achieved can be relegated to an intra-stage sub-problem.

A possible form for the release distribution suggests itself immediately. Since the total water release during a stage is the sum of the random releases during many individual trading periods, a normal distribution might well be a good model. (An oversimplified model would be: if a single tranche were offered in each trading period, and the trading periods had independent and identically distributed prices, the total release for the stage would have a binomial distribution.) Some empirical support for this idea is provided by data such as those shown in Figure 1. Other distributions could also be chosen, without invalidating the essential ideas of this paper.

With this formulation, we have thus reduced the decision space for a single stage to a space of probability distributions characterized by a small number of parameters. For the normal distribution, there are only two parameters: the mean μ and standard deviation σ of the total release for the stage. These parameters have a strong intuitive appeal: μ can be thought of as a "target release" for the coming stage, while σ measures the degree of "risk", or deviation from the target, that the operator is willing to contemplate. To the extent that the planning of a whole week's operations can be specified by only two parameters, μ and σ are natural choices.

Let us now express this formulation of the problem mathematically, as a stochas-

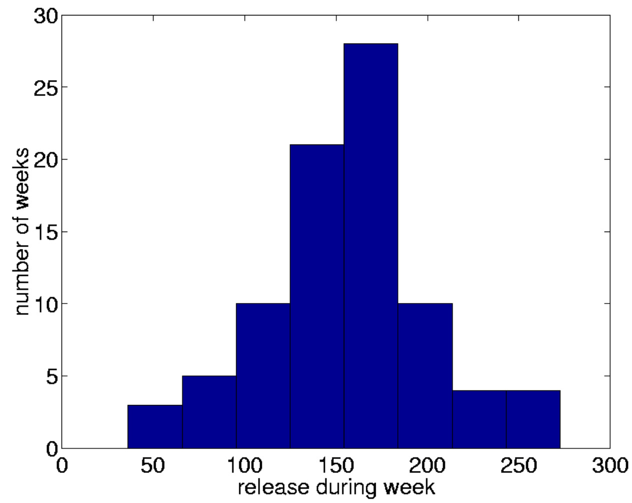

Figure 1: Total weekly release, when the release in each half-hour period is given by $S(p) = \min(1, (1 + ((p/40) - 1)^3)/2)$, where p is the prevailing price in \$/MWh. (S has a shape common among supply functions.) The price data are from the Haywards series in [16], and cover the summer season (November through February) of the five-year period to February 2003.

tic dynamic programming model. The essential relation is:

$$V_t(x) = \max_{\mu,\sigma} \left(g_t(\mu, \sigma, Y_t) + E\left[V_{t+1}(x + W_t - \mu - \sigma Z_t) \,|\, Y_t \right] \right) \qquad (1)$$

where

$V_t(x)$ is the value of having reservoir level x at the beginning of stage t; as it would be determined at that time. In other words, $V_t(x)$ may depend on random events that occurred at earlier stages - and in that sense it is a random variable itself - but not on random events occurring during stage t or later.

W_t is the (random) inflow to the reservoir during stage t.

Z_t is a standard normal random variable, representing the normalized water release. In the simplest formulation, Z_t and W_t would be independent. However, it is possible to assume some other joint distribution for (W_t, Z_t), reflecting a possible correlation between inflows and prices at the same stage.

Y_t is a state variable, representing relevant information (about weather, market conditions, etc.) available at the beginning of stage t. The variable Y_t may be multidimensional, and indeed is likely to be at least two-dimensional, since it must include information relevant to both inflows and prices in future periods. For example, (Y_t) could be a Markov chain with 10 states as indicated by the bullets in the following table.

		weather forecast		
		dry	normal	wet
	glut		•	•
storage conditions:	normal	•	•	•
	shortage	•	•	•
	critical shortage	•	•	

This example models a hydro-dominated electricity system. (The real-world model is New Zealand, where hydro-electricity accounts for approximately two-thirds of all electricity generated.) In such a context, two main kinds of information are available and relevant to the prices and inflows that our generator will receive in future stages: a long-term weather forecast for the next (say) 12 weeks, and current reservoir storage for the whole system. (The latter is not to be confused with the storage in our own reservoir, which will likely be only a small part of the system total.) Note that the weather forecast has relevance to both future prices and future inflows; also that the two least likely combinations have been omitted from the state space.

$g_t(\mu, \sigma, y)$ is the expected payoff during stage t, when the decision made for the stage is (μ, σ) and the auxiliary process Y_t is in state y. The value of $g_t(\mu, \sigma, y)$ is itself the optimal value of an intra-stage subproblem, which has the form

$$\begin{array}{rl} \max_{s \in \mathcal{S}} & E\left[R_t(s) \,|\, Y_t = y\right] \\ \text{subject to} & E\left[X_t(s) \,|\, Y_t = y\right] = \mu, \\ & \mathrm{Var}\left(X_t(s) \,|\, Y_t = y\right) = \sigma^2, \end{array} \tag{2}$$

where

$\mathcal{S}$ is the space of possible release policies for the stage (perhaps a high-dimensional space).

$R_t(s)$ is the (random) revenue that results from adopting policy s during stage t.

$X_t(s)$ is the (random) total water release that results from adopting policy s during stage t.

The essential virtue of this approach is that it decomposes the problem into two parts. The *intra-stage problem* is to compute $g_t(\mu, \sigma, y)$ by solving (2). This can be carried out off-line, to give $g_t(\mu, \sigma, y)$ for a table of values of μ, σ, and y. Note that to do this, one does not need to consider the reservoir level. Observe that the release policy s determines a collection of (possibly different) supply functions to offer in each trading period in a stage; optimizing this policy may present us with a considerable computational challenge.

The *inter-stage problem* is to solve (1), the dynamic programming problem, in the usual backwards-recursive fashion. For each stage, this requires referring to previously computed tables of values of g_t and V_{t+1}, and interpolating where necessary. This inter-stage problem is quite tractable, since the internal structure of each stage has been hidden in the intra-stage problem.

In the next two sections, we will study the intra-stage and inter-stage problems in more detail. Before doing so, however, we will address one question regarding (2). The variance constraint there has been given as $\mathrm{Var}\left(X_t(s) \,|\, Y_t = y\right) = \sigma^2$, but why should it not be $\text{Var}\,(X_t(s)\,|\,Y_t = y) \le \sigma^2$? After all, if a policy can be found with maximal expected payoff, but less risk than the maximum permitted, why not use it? The following result shows that it will usually make no difference which version of the constraint is used.

Lemma 2.1 *Let $g_t(\mu, \sigma, y)$ be given by (2) as above, while $\tilde{g}_t(\mu, \sigma, y)$ is the optimal value of the problem*

$$\begin{array}{rl} max_{s\in\mathcal{S}} & E\,[R_t(s)\,|\,Y_t = y] \\ subject\ to & E\,[X_t(s)\,|\,Y_t = y] = \mu, \\ & Var\,(X_t(s)\,|\,Y_t = y) \le \sigma^2, \end{array} \tag{3}$$

Suppose V_{t+1} is concave. Then $V_t(x)$ and $\tilde{V}_t(x)$ given by

$$\begin{aligned} V_t(x) &= max_{\mu,\sigma}\left(g_t(\mu, \sigma, Y_t) + E\,[V_{t+1}(x + W_t - \mu - \sigma Z_t)\,|\,Y_t]\right) \\ \tilde{V}_t(x) &= max_{\mu,\sigma}\left(\tilde{g}_t(\mu, \sigma, Y_t) + E\,[V_{t+1}(x + W_t - \mu - \sigma Z_t)\,|\,Y_t]\right) \end{aligned}$$

are equal.

Proof. We have $g_t(\mu, \sigma, y) \le \tilde{g}_t(\mu, \sigma, y)$ from their definitions, and so $V_t(x) \le \tilde{V}_t(x)$. To see the reverse inequality, fix y and suppose that when $Y_t = y$, we have $(\mu, \tilde{\sigma})$ achieving the optimum value of $\tilde{V}_t(x)$, and a corresponding s achieving the optimum value of (3) for $(\mu, \tilde{\sigma})$. Of course we have $\text{Var}\,(X_t(s)\,|\,Y_t = y) \le \tilde{\sigma}^2$. Let $\sigma^2 = \text{Var}\,(X_t(s)\,|\,Y_t = y)$. Since s is also feasible for (2), we have $\tilde{g}_t(\mu, \tilde{\sigma}, y) \le g_t(\mu, \sigma, y)$. Furthermore, concavity (and the symmetry of the normal distribution) give

$$E\,[V_{t+1}(x + W_t - \mu - \tilde{\sigma}Z)\,|\,Y_t = y] \le E\,[V_{t+1}(x + W_t - \mu - \sigma Z)\,|\,Y_t = y]$$

It follows that $\tilde{V}_t(x) \le V_t(x)$. ∎

In the next section we study the intra-stage problem, and provide a methodology for computing an approximate solution. We also show that this approximate problem is convex, which allows us to show that $g_t\,(\mu, \sigma, y)$ is concave in both μ and σ. This fact will be important in verifying the tractability of the inter-stage model.

3 The intra-stage model

In this section we describe a more specific version of the intra-stage problem (2). The starting point is the assumption that the generator's "release policy" for a given stage will consist of a single supply function, to be offered to the market in every trading period during the stage. While this may sound simple, it appears to be not so different from the way that at least some hydro-electric reservoirs are operated in practice. Moreover, it makes intuitive sense - since we have already assumed that market prices are exogenous and cannot be influenced, there is little reason to vary the supply function between (say) peak and off-peak periods.

We further assume that the objective that is to be maximized is gross revenue (i.e. price×dispatched quantity). As usual with hydro-electric problems, there is no "fuel cost", other than a water value which is provided for within the model.

A key observation is that the actual price process is not important; the *price duration curve* (giving the proportion of half-hourly prices below any given price p) contains sufficient information for our purposes. Given the price duration curve for a given stage t and state y, and a supply function s used consistently during the stage, it is possible to compute by integration the revenue $R_t(s)$ and water release $X_t(s)$ that result. Instead of seeking stochastic models for random prices, therefore, we will focus on modelling random price duration curves. Then for any random price duration curve π, the intra-stage problem 2 has the form

$$\begin{array}{rl} \max & E\left[\int_{p_{min}}^{p_{max}} pq(p)\, d\pi(p)\right] \\ \text{subject to} & q : [p_{min}, p_{max}] \to [0, q_{\max}] \text{ is a non-decreasing function} \\ & E[R] = \mu, \qquad \text{Var}(R) \le \sigma^2, \\ \text{where} & R = \int_{p_{min}}^{p_{max}} q(p)\, d\pi(p). \end{array} \tag{4}$$

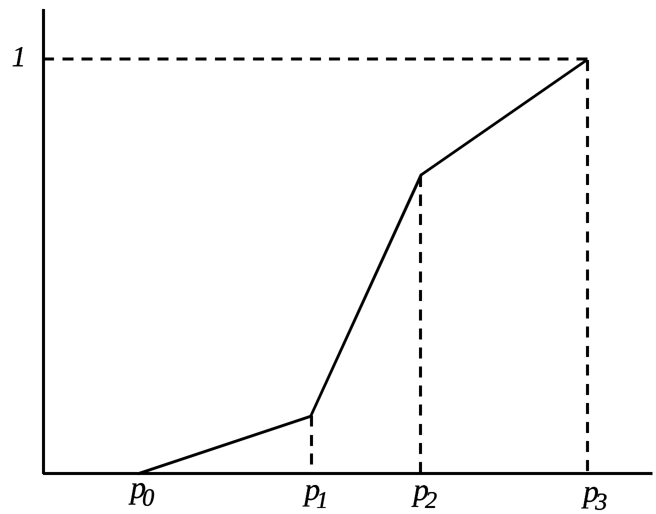

Figure 2: Piecewise linear price duration curve.

The model we propose for the price duration curve is a piecewise linear function with non-random breakpoints and random slopes. Figure 2 is a price duration curve of this kind, with three linear pieces. In general, there are m fixed price bands, being $[p_0, p_1), [p_1, p_2), \ldots [p_{m-1}, p_m]$, with $p_0 = p_{min}$ and $p_m = p_{max}$. (These price bands might vary with stage t and state y, but for simplicity of notation in this section we will suppress the dependence of the intra-stage problem on t and y.) Within each price band, let the slope of the linear pieces be given by random variables $A_1, \ldots, A_m$ respectively. Note that we must have $\sum_{i=1}^{m} A_i(p_i - p_{i-1}) = N$, the number of trading periods in the stage, and the joint distribution of $A_1, \ldots, A_m$ should be chosen to ensure this. One promising model is the Dirichlet distribution (see [10]).

We seek a supply function $q(p)$ maximizing expected return subject to constraints on mean and variance. We define

$$I_i = \int_{p_{i-1}}^{p_i} q(p)\, dp, \text{ for } i = 1, \ldots, m.$$

Then the intra-stage problem is

$$\begin{array}{rll} g(\mu,\sigma) = \max & \sum_{i=1}^{m} a_i \int_{p_{i-1}}^{p_i} pq(p)\,dp & \\ \text{subject to} & q:[p_0,p_m] \longrightarrow [0,q_{\max}] \text{ is a non-decreasing function} & \\ & I_i = \int_{p_{i-1}}^{p_i} q(p)\,dp, \qquad i=1,\ldots,m, & \\ & \mathbf{a}^\top \mathbf{I} = \mu, & \\ & \mathbf{I}^\top \mathbf{V} \mathbf{I} \le \sigma^2, & \end{array} \tag{5}$$

where $\mathbf{a}$ and $\mathbf{V}$ are, respectively, the mean and variance-covariance matrix of the random vector $(A_1,\ldots,A_m)$. Note that (5) uses the inequality version of the variance constraint.

The problem (5) seeks to maximize a linear functional over a set $\mathcal{Q}$ of feasible supply functions. It is worth noting that the maximum must be attained: $\mathcal{Q}$ is compact under the quadratic norm given by $\|q(\cdot)\|^2 = \int_{p_0}^{p_m} q(p)^2\ dp$, and our objective functional is continuous on $\mathcal{Q}$ with this norm.

Another significant feature of (5) is that it is a convex optimization problem, since $\mathcal{Q}$ is a convex set. (To see that the variance constraint is a convex condition, note that the expression $\sqrt{\mathbf{I}^\top \mathbf{V}\mathbf{I}}$ defines a semi-norm – which is a convex function – on $\mathcal{Q}$.) This convexity will ultimately play an important role in ensuring the tractability of the dynamic programming recursion (1). We begin by using it to establish the concavity of g in each of its arguments.

Lemma 3.1 *$g(\mu,\sigma)$ as defined in (5) is concave in (μ,σ).*

Proof. Let us write the feasible set of (5) as $\mathcal{Q}(\mu,\sigma)$, to emphasize the dependence on the parameters μ and σ. Let $F(q(\cdot)) = \sum_{i=1}^{m} a_i \int_{p_{i-1}}^{p_i} pq(p)\,dp$. Note that if $q_1 \in \mathcal{Q}(\mu_1,\sigma_1)$ and $q_2 \in \mathcal{Q}(\mu_2,\sigma_2)$, then for any $\lambda \in (0,1)$,

$$\lambda q_1 + (1-\lambda)q_2 \in \mathcal{Q}(\lambda\mu_1 + (1-\lambda)\mu_2, \lambda\sigma_1 + (1-\lambda)\sigma_2),$$

and so

$$\begin{aligned} g(\lambda\mu_1 + (1-\lambda)\mu_2, \lambda\sigma_1 + (1-\lambda)\sigma_2) &\ge F(\lambda q_1 + (1-\lambda)q_2) \\ &= \lambda F(q_1) + (1-\lambda)F(q_2). \end{aligned}$$

Taking the supremum over all $q_1 \in \mathcal{Q}(\mu_1,\sigma_1)$ and $q_2 \in \mathcal{Q}(\mu_2,\sigma_2)$ gives

$$g(\lambda\mu_1 + (1-\lambda)\mu_2, \lambda\sigma_1 + (1-\lambda)\sigma_2) \ge \lambda g(\mu_1,\sigma_1) + (1-\lambda)g(\mu_2,\sigma_2),$$

as required. ∎

The problem (5) is an infinite dimensional optimization problem. However, it is not too difficult to see that the optimum $q(\cdot)$ must have a piecewise constant form with m or fewer tranches. In other words, the optimum supply function – conveniently enough – has the form of an offer stack. We formally present this result in the following lemma.

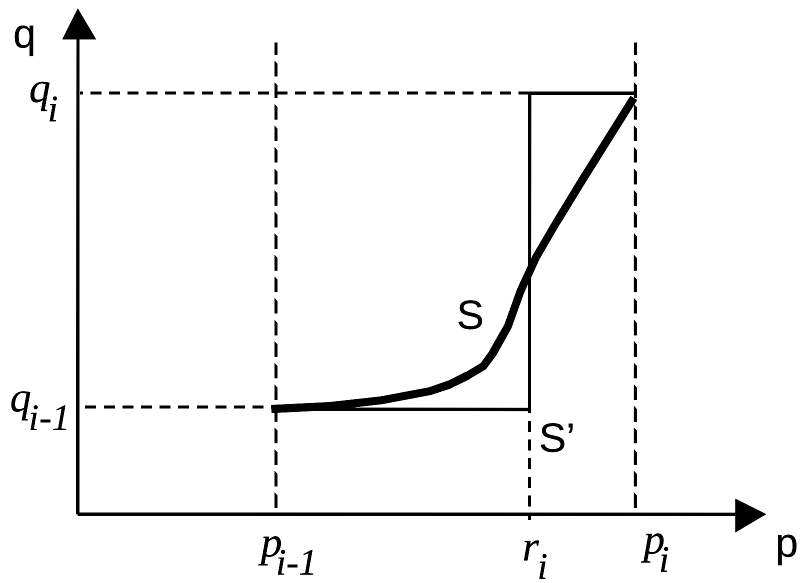

Figure 3: An optimal stack can be taken to be piecewise constant

Lemma 3.2 *The problem (5) has an optimal piecewise constant solution $q(p)$ in which each interval $[p_{i-1}, p_i]$ contains at most one breakpoint.*

Proof. Suppose $q = S(p)$ is an optimal solution for (5) that is not constant on some interval $[p_{i-1}, p_i]$. Let $q_{i-1} = S(p_{i-1})$ and $q_i = S(p_i)$, and define

$$r_i = \frac{(q_i p_i - q_{i-1} p_{i-1}) - \int_{p_{i-1}}^{p_i} S(p)\, dp}{q_i - q_{i-1}}.$$

Then it is easy to show that $p_{i-1} \le r_i \le p_i$. If $r_i = p_{i-1}$, then let

$$S'(p) = \begin{cases} q_i, & p_{i-1} < p \le p_i, \\ S(p), & \text{otherwise,} \end{cases}$$

and if $r_i = p_i$, then let

$$S'(p) = \begin{cases} q_{i-1}, & p_{i-1} \le p < p_i, \\ S(p), & \text{otherwise.} \end{cases}$$

If $p_{i-1} < r_i < p_i$, then let

$$S'(p) = \begin{cases} q_{i-1}, & p_{i-1} \le p < r_i, \\ q_i, & r_i \le p < p_i, \\ S(p), & \text{otherwise} \end{cases}$$

(see Figure 3). Observe that $S'(p)$ is piecewise constant with at most one breakpoint in $[p_{i-1}, p_i]$. Furthermore by virtue of the definition of r_i we have

$$\int_{p_{i-1}}^{p_i} S'(p)\, dp = \int_{p_{i-1}}^{p_i} S(p)\, dp,$$

so S and S' contribute equally to the constraints on mean and variance.

Let $U(x)=\int_{p_{i-1}}^{x} S(p)\,dp$, and $U'(x)=\int_{p_{i-1}}^{x} S'(p)\,dp$. Then

$$\frac{d(U(x)-U'(x))}{dx}=S(x)-S'(x)$$

which is nonnegative for $p_{i-1}\le x\le r_i$, and non positive for $r_i\le x\le p_i$. So $U(x)-U'(x)\ge 0$ for all x in $[p_{i-1},p_i]$, with equality at p_{i-1} and p_i. We thus have

$$\begin{aligned}\int_{p_{i-1}}^{p_i} p(S'(p)-S(p))dp &= [p(U'(p)-U(p))]_{p_{i-1}}^{p_i}-\int_{p_{i-1}}^{p_i}(U'(p)-U(p))\,dp\\ &= -\int_{p_{i-1}}^{p_i}(U'(p)-U(p))\,dp\\ &\ge 0.\end{aligned}$$

It follows that S' is optimal for (5). Repeating the argument for all intervals $[p_{i-1},p_i]$ yields the result. ∎

By virtue of lemma 3.2 we can reformulate (5) so as to include the breakpoints of $q(p)$ as variables. We have shown that there exists a piecewise constant optimal solution with m or fewer breakpoints, so we label these as r_i, $i=1,\dots,m$, and let q_i, $i=1,\dots,m$, be the (constant) offer quantities at these prices (we denote by q_0 the amount offered at price p_0). This gives the following optimization problem.

$$\begin{array}{ll}\max_{q_0,\{r_i,q_i,I_i\}:i=1,\dots,m} & \sum_{i=1}^{m} a_i\left(\int_{p_{i-1}}^{r_i} q_{i-1}p\,dp+\int_{r_i}^{p_i} q_i p\,dp\right)\\ \text{subject to} & I_i=\int_{p_{i-1}}^{r_i} q_{i-1}dp+\int_{r_i}^{p_i} q_i dp,\quad i=1,2,\dots,m,\\ & 0\le q_0\le q_1\le\dots q_m\le q_{Max},\\ & p_{i-1}\le r_i\le p_i,\quad i=1,\dots,m,\\ & \mathbf{a}^{\top}\mathbf{I}=\mu,\\ & \mathbf{I}^{\top}\mathbf{V}\mathbf{I}\le\sigma^2.\end{array}\tag{6}$$

A disadvantage with this approach is that (6) is no longer a convex mathematical program. However it can be approximated by a convex quadratic program, by discretizing the prices that can be chosen. We do this by defining for each $j=1,\dots,J$, a price interval $[d_j,d_{j+1})$ of fixed length $u=(p_{Max}-p_{Min})/J$, and choosing a quantity q_j so that $q(p)$ is constant (and equal to q_j) over the interval $[d_j,d_{j+1})$. For each interval $[p_{i-1},p_i)$ in the price duration curve it is helpful to define the set $R_i=\{j:p_{i-1}\le d_j<p_i\}$.

The objective function is then approximated by

$$\begin{aligned}&\sum_{i=1}^{m} a_i\sum_{j\in R_i}\frac{1}{2}\left(d_{j+1}^2-d_j^2\right)q_j\\ =\;& u\sum_{i=1}^{m} a_i\sum_{j\in R_i}\left(d_j+\frac{u}{2}\right)q_j,\end{aligned}$$

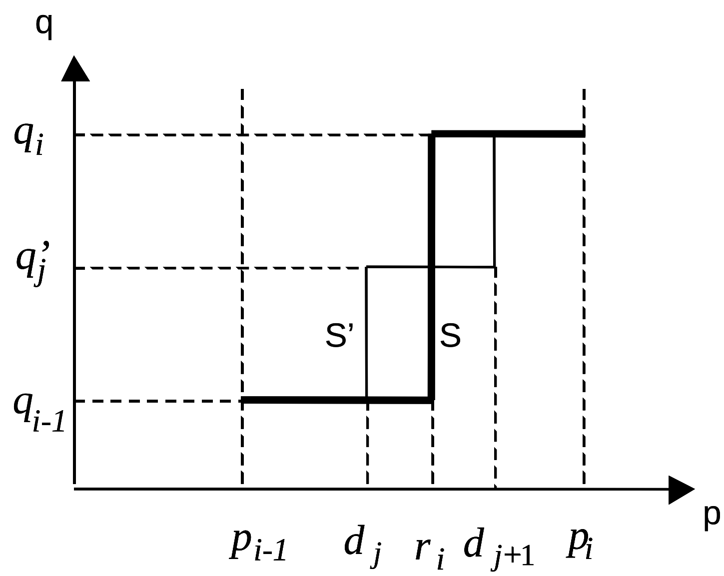

Figure 4: Optimal offer S (bold line) and discretized offer S' (solid line)

which is linear in q. Under the discretization the first m constraints of (6) change to

$$I_i = u \sum_{j \in R_i} q_j, \; i = 1, \ldots, m,$$

which are linear in q and I. The stack must be increasing, so we require

$$0 \leq q_1 \leq \ldots \leq q_J \leq q_{Max}.$$

This yields the quadratic program

$$\begin{array}{rl} g(\mu, \sigma) = \max_{q_j : j=1,\ldots,J} & u \sum_{i=1}^{m} a_i \sum_{j \in R_i} \left(d_j + \frac{u}{2}\right) q_j \\ \text{subject to} & I_i = u \sum_{j \in R_i} q_j, \; i = 1, \ldots, m, \\ & \mathbf{a}^\top \mathbf{I} = \mu, \\ & 0 \leq q_1 \leq \ldots \leq q_J \leq q_{Max}, \\ & \mathbf{I}^\top \mathbf{V} \mathbf{I} \leq \sigma^2, \end{array} \tag{7}$$

a convex problem since $\mathbf{V}$ is positive semidefinite.

By choosing J large enough we can obtain an approximate solution to (5). This is made precise as follows. Let $q = S(p)$ be a globally (piecewise constant) optimal solution for (5). We depict part of this solution by bold lines in Figure 4. If we discretize the price axis into J equal sections, then for each of the m pieces, the optimal solution S can be approximated by a supply function $q' = S'(p)$, which is a feasible solution of (7). This approximation will be exact except in any interval $[d_j, d_{j+1})$ containing a breakpoint r_i of S, as shown in Figure 4.

To ensure that S' is a feasible solution for (5), we may choose q_j' to give

$$\int_{p_{i-1}}^{r_i} q_{i-1} dp + \int_{r_i}^{p_i} q_i dp = \int_{p_{i-1}}^{d_j} q_{i-1} dp + \int_{d_j}^{d_{j+1}} q_j' dp + \int_{d_{j+1}}^{p_i} q_i dp,$$

so

$$q_{i-1}\left(r_i - p_{i-1}\right) + q_i\left(p_i - r_i\right) = q_{i-1}\left(d_j - p_{i-1}\right) + q_j'\left(d_{j+1} - d_j\right) + q_i\left(p_i - d_{j+1}\right),$$

whereby

$$q_j' = \frac{q_{i-1}r_i - q_i r_i - q_{i-1}d_j + q_i d_{j+1}}{d_{j+1} - d_j}. \tag{8}$$

Then the difference in objective between the two stacks over the range of prices $[p_{i-1}, p_i]$ is

$$\begin{aligned}
\Delta &= \int_{d_j}^{r_i} p q_{i-1} dp + \int_{r_i}^{d_{j+1}} p q_i dp - \int_{d_j}^{d_{j+1}} p q_j' dp \\
&= \frac{1}{2}\left[q_{i-1}\left(r_i^2 - d_j^2\right) + q_i\left(d_{j+1}^2 - r_i^2\right) - q_j'\left(d_{j+1}^2 - d_j^2\right)\right].
\end{aligned}$$

Using (8), we get

$$\begin{aligned}
\Delta &= \frac{1}{2}\left[q_{i-1}\left(r_i^2 - d_j^2\right) + q_i\left(d_{j+1}^2 - r_i^2\right) - \left(q_{i-1}r_i - q_i r_i - q_{i-1}d_j + q_i d_{j+1}\right)\left(d_{j+1} + d_j\right)\right] \\
&= \frac{1}{2}\left(q_i - q_{i-1}\right)\left(d_{j+1} - r_i\right)\left(r_i - d_j\right) \\
&= \frac{1}{2}\left(q_i - q_{i-1}\right)\left[\lambda\left(d_{j+1} - d_j\right)\left(1 - \lambda\right)\left(d_{j+1} - d_j\right)\right], \text{ for some } 0 \le \lambda \le 1, \\
&\le \frac{1}{8}\left(q_i - q_{i-1}\right)\left(d_{j+1} - d_j\right)^2 \\
&= \frac{1}{8}\left(q_i - q_{i-1}\right)\left(\frac{p_{Max} - p_{Min}}{J}\right)^2.
\end{aligned}$$

So, summing over the entire range of S, the difference in objective between the optimal stack and the feasible solution to the discretized problem is bounded above by

$$\sum_{i=1}^{m} \frac{1}{8}\left(q_i - q_{i-1}\right)\left(\frac{p_{Max} - p_{Min}}{J}\right)^2 \le \frac{1}{8} q_{Max}\left(\frac{p_{Max} - p_{Min}}{J}\right)^2.$$

4 The inter-stage problem

We now return to the stochastic dynamic programming problem (1):

$$V_t(x) = \max_{\mu,\sigma}\left(g_t(\mu, \sigma, Y_t) + E\left[V_{t+1}(x + W_t - \mu - \sigma Z_t) \,|\, Y_t\right]\right). \tag{9}$$

The random variable Y_t represents the information available at the beginning of stage t and relevant to inflows and prices during stages t and later. It is envisaged in this paper that the stochastic process (Y_t) will be a fairly simple one; perhaps a finite-state Markov chain, as suggested in section 2. For each possible value y of Y_t, there will be a distribution of price duration curves for use in stage t (described by parameters $\mathbf{a}_{y,t}$ and $\mathbf{V}_{y,t}$ as discussed in section 3); a distribution for the inflow W_t at stage t; and a distribution for Y_{t+1}.

The first step in solving (9) is to solve the intra-stage problem to find values of $g_t(\mu, \sigma, y)$ for each t,y and (a table of possible values of) μ, σ. These values can be computed offline and re-used in many different versions of the overall problem. For example, the decision horizon, inflow distribution, reservoir size, and transition probabilities for (Y_t) can all be varied without needing to re-solve the intra-stage problem. The value of $g_t(\mu, \sigma, y)$ is decoupled from all of these parameters, being only a measure of how much payoff can be extracted from a given set of market conditions during a stage where water use is constrained by the planning parameters μ, σ. Note that although we have allowed for dependence on the stage t, it is possible that the computations for this step will be identical for several stages, or even for all stages; this offers a careful implementation some potential gains in execution time.

Since μ and σ are continuous variables, it would be possible to improve the accuracy of the representation of the function $g_t(\cdot, \cdot, y)$ by using an interpolation scheme (e.g. a spline) for values of (μ, σ) in between those for which the intra-stage problem has been solved. However, as we shall see shortly, this may prove to be unnecessary.

Now, with the function g_t in hand, let us consider the other parts of (9). For computational purposes, we replace the random variable $V_t(x)$ with the more explicit expression $v_t(x, Y_t)$, where v_t is a deterministic function. The problem (9) is to be solved in the usual backwards-recursive fashion: we choose a final water value function $v_T(x, y)$ at the time horizon T and then for $t = T-1, T-2, \ldots, 0$ (where $t = 0$ represents the present time):

1. Extend v_{t+1} by setting

$$\begin{aligned} v_{t+1}(x,y) &= v_{t+1}(x_{max}, y) && \text{for } x > x_{max} \\ v_{t+1}(x,y) &= v_{t+1}(0,y) - \rho(|x|) && \text{for } x < 0, \end{aligned}$$

   where x_{max} is the maximum capacity of the reservoir, and ρ is a penalty function (see below).

2. Compute

$$v_t(x,y) = \max_{\mu,\sigma} \left(g_t(\mu, \sigma, y) + h_t(\mu, \sigma, x, y)\right)$$

   for each possible y and a finite collection of values $0 \le x_1 \le \cdots \le x_n \le x_{max}$ of x. Here

$$\begin{aligned} h_t(\mu, \sigma, x, y) &= E\left[V_{t+1}(x + W_t - \mu - \sigma Z_t) \,|\, Y_t = y\right] \\ &= \sum_{y'} p_{y,y'} \int_0^\infty \int_{-\infty}^\infty v_{t+1}(x + w - \mu - \sigma z, y') f_y(w,z)\, dz\, dw, \end{aligned}$$

   where $p_{y,y'}$ is the transition probability $P\left(Y_{t+1} = y' \,|\, Y_t = y\right)$, and f_y is the joint probability density of (W_t, Z_t), conditional on $Y_t = y$.

Step 1 above addresses an issue which we have not previously considered – the range of permissible values of the reservoir level x. In any realistic problem there must be both an upper and a lower bound on this level. Without loss of generality, we may consider these bounds to take the form $0 \leq x \leq x_{max}$. But in our model, the decision made (release policy formulated) at the beginning of each stage does not determine how much water will remain at the end of the stage. (Nor could it, in any model taking realistic account of the uncertainty in prices and inflows.) It is therefore not possible strictly to enforce the bounds on reservoir level.

The upper bound x_{max} is the easiest to deal with. It is reasonable to assume that if the reservoir overflows, the excess water will be spilled and lost, thus having no value. This assumption can be incorporated into the model by extending the function v_{t+1} to values of $x > x_{max}$ as indicated above when computing v_t.

The lower bound is more troublesome. In reality, if the reservoir began to run out of water during a stage, the operator would modify its release policy for the remainder of the stage. But we cannot incorporate this into our model without substantially redesigning it; our paradigm is that any decision can be made only at the beginning of a stage. Instead, we extend v_{t+1} to negative values of x by adding a penalty term $\rho(|x|)$; a linear penalty $\rho(u) = Cu$ seems to be sufficient. The effect of this is to preclude the choice at stage t of (μ, σ) pairs for which there is an appreciable chance of offering to supply more energy than the reservoir can generate during stage t. The magnitude of the penalty constant C indicates the degree of aversion to this particular risk.

So far we have said nothing about the form of the distribution of W_t. One convenient (though not particularly realistic) assumption is that W_t is normally distributed; this means that $W_t - \mu - \sigma Z_t$ is also normally distributed, and it is possible to compute the integral in step 2. above by Hermite-Gauss quadrature (see Hildebrand [11] for an explanation of the procedure and Abramowitz and Stegun [1] p. 890 for numerical values used in the quadrature). Other distributions for W_t can be handled equally well by a similar technique. Note that the numerical evaluation of this integral also requires interpolation to obtain values of v_{t+1}.

Step 2 requires the solution of an optimization problem. We will show in a moment that this is a convex problem, i.e. that the function $g_t(\mu, \sigma, y) + h_t(\mu, \sigma, x, y)$ to be maximized is concave in (μ, σ). But the value of this function is known only at the finitely many points at which $g_t(\cdot, \cdot, y)$ has been computed; we are reliant on interpolation for the values at other points. It would be natural to use an interpolation scheme which preserves convexity. Linear interpolation is one particularly simple possibility, as it ensures that the maximum of the interpolated function will occur at one of the points where the value of $g_t(\cdot, \cdot, y)$ is known, so that the optimization reduces to evaluating $g_t(\cdot, \cdot, y) + h_t(\cdot, \cdot, x, y)$ at these points and identifying the largest value.

Theorem 4.1 *If $V_{t+1}(x)$ is concave in x with probability 1, then so is $V_t(x)$.*

It follows from this that if $v_T(x, y)$ is concave in x for each y, then the same is true of each $v_t(x, y)$.

Proof. First note that

$$h_t(\mu, \sigma, x, y) = E\left[V_{t+1}(x + W_t - \mu - \sigma Z_t) \,|\, Y_t = y\right]$$

is concave in x, μ, and σ. To see this, note that $x - \mu - \sigma Z_t$ is linear in x, μ, and σ with probability 1, and so $V_{t+1}(x - \mu - \sigma Z_t + W_t)$ is concave in those variables with probability 1. Since the expectation of concave functions is concave, h_t is concave in x, μ, and σ.

We also have from Lemma 3.1 that $g_t(\mu, \sigma, y)$ is concave in μ and σ. Hence $g_t(\mu, \sigma, y) + h_t(\mu, \sigma, x, y)$ is concave in x, μ, and σ.

It now remains to be shown that the concavity in x is preserved by the taking of the maximum in Step 3 above. For fixed y, suppose μ_1^* and σ_1^* maximize $g_t(\mu, \sigma, y) + h_t(\mu, \sigma, x_1, y)$, and suppose μ_2^* and σ_2^* maximize $g_t(\mu, \sigma, y) + h_t(\mu, \sigma, x_2, y)$.

Let

$$\hat{x} = \lambda x_1 + (1 - \lambda) x_2, \qquad \hat{\mu} = \lambda \mu_1^* + (1 - \lambda) \mu_2^*, \qquad \text{and } \hat{\sigma} = \lambda \sigma_1^* + (1 - \lambda) \sigma_2^*,$$

for some $\lambda \in (0, 1)$. Then

$$\begin{aligned} g_t(\hat{\mu}, \hat{\sigma}, y) + h_t(\hat{\mu}, \hat{\sigma}, \hat{x}, y) &= g_t\left(\lambda \mu_1^* + (1 - \lambda) \mu_2^*, \lambda \sigma_1^* + (1 - \lambda) \sigma_2^*, y\right) \\ &\quad + h_t\left(\lambda \mu_1^* + (1 - \lambda) \mu_2^*, \lambda \sigma_1^* + (1 - \lambda) \sigma_2^*, \lambda x_1 + (1 - \lambda) x_2, y\right) \\ &\geq \lambda g_t(\mu_1^*, \sigma_1^*, y) + (1 - \lambda) g_t(\mu_2^*, \sigma_2^*, y) + \lambda h_t(\mu_1^*, \sigma_1^*, x_1, y) \\ &\quad + (1 - \lambda) h_t(\mu_2^*, \sigma_2^*, x_2, y). \end{aligned}$$

But since this is just one feasible μ, σ pair, the left-hand-side of this expression is a lower bound for $v_t(\lambda x_1 + (1 - \lambda) x_2, y)$, and the result follows. ∎

5 Conclusions and Further Work

In this paper we have described a model based on parameterized price duration curves to represent the prices over the course of a week, and derived a convex optimization model to compute an optimal supply function for this week. Under mild assumptions on the form of our price duration curve, we have shown that the optimal stack is piecewise constant.

We have developed a model that aggregates trading periods into stages with the same stack offered at each trading period within a stage. Using a normal distribution to represent the total amount of water used in each week leaves us with three decisions to make for each week — the mean, the standard deviation and an offering strategy yielding these properties.

A possible weakness in our approach is the assumption that generators will not affect the price by their choice of offer. For large generators, this may not be the case. Note that to improve the model to allow for such "price-making" behaviour, only the intra-stage value functions $g_t(\mu, \sigma, y)$ need to be modified; the rest of the model can be left unchanged.

One approach to handling price-making is given by Anderson and Philpott [3], who define a *market distribution function* $\psi(q, p)$ that gives the probability of an offer of q units of electricity at price p not being fully dispatched.

One might hope to estimate and use $\psi(q, p)$ to calculate weekly returns $g(\mu, \sigma)$ in an intra-stage model with price-making behaviour. However, it is very likely that ψ will vary over the trading periods in the stage in response to the stack we offer to the market in the first trading period, so assuming a single offer stack for every trading period in the stage would be unrealistic, even if it were easy to compute. Furthermore, the inter-stage model requires some methodology of estimating market distribution functions for future stages by relating the behaviour of other generators to observable state variables like regional water storage and weather forecasts. This is a difficult undertaking.

Another challenge is the extension of the model from one to several (hydrologically coupled) reservoirs. This would require additional state variables for the reservoir levels, and a corresponding exponential increase in computational effort (the well-known "curse of dimensionality"). The handling of underflow and overflow constraints would also become more difficult.

Acknowledgement

The authors are grateful to the anonymous referees whose comments have improved the exposition of this paper. The authors also acknowledge the support of FRST under contract UOAX0203.

References

[1] Abramowitz, M., Stegun, I. A. (eds) (1972): *Handbook of Mathematical Functions with Formulas, Graphs and Mathematical Tables*, Wiley New York.

[2] Alvey, T., Goodwin, D. Ma, X., Streiffert, D., Sun D. (1998): A security-constrained bid-clearing system for the New Zealand wholesale electricity market. *IEEE Transactions on Power Systems* 13, 340–346.

[3] Anderson, E. J., Philpott A. B. (2002): Optimal offer construction in electricity markets. *Mathematics of Operations Research* 27, 82–100.

[4] Birge, J.R. (1985): Aggregation bounds in stochastic linear programming. *Mathematical Programming* 31, 25–41.

[5] Bushnell, J. (2003): A mixed-complementarity model of hydrothermal electricity competition in the Western United States. *Operations Research* 51, 80–93.

[6] Chao, H-P., Huntington H. (eds) (1998): *Designing Competitive Electricity Markets.* Kluwer Academic, Boston.

[7] Fleten S.-E., Wallace, S.W., Ziemba, W.T. (2002): Hedging electricity portfolios via stochastic programming. In: *Decision Making Under Uncertainty: Energy and Power,* Greengard, C., and A. Ruszczynski (ed.), 71–94.

[8] Gjelsvik, A., Wallace, S.W. (1996): Methods for stochastic medium-term scheduling in hydrodominated power systems. Norwegian Electric Power Research Institute, Technical report A4438.

[9] Gjelsvik, A., Belsnes, M.M., Haugstad, A. (1999): An algorithm for stochastic medium-term hydrothermal scheduling under spot price uncertainty. In: *Proceedings, 13th Power Systems Computation Conference,* Trondheim, Norway, 1079-1085.

[10] Guan, Z. (2003): Optimal intra-week policy for electricity generation. M.Sc. thesis, University of Auckland.

[11] Hildebrand, F. B. (1974): *Introduction to Numerical Analysis.* McGraw-Hill, New York.

[12] Hu, P.-T. (2002): Single reservoir long-term hydro scheduling in an electricity pool market. M.O.R. Thesis, Department of Engineering Science, University of Auckland.

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

[14] Neame, P. J., Philpott, A. B., Pritchard, G. (2003): Offer stack optimization in electricity pool markets. *Operations Research* 51, 397-408.

[15] Pereira, M. V. F., Pinto, L. M. V. G. (1991): Multi-stage stochastic optimization applied to energy planning. *Mathematical Programming* 52, 359–375.

[16] Pritchard, G. New Zealand electricity prices *http://www.stat.auckland.ac.nz/~geoff/elecprices*

[17] Pritchard, G., Zakeri, G. (2003): Market offering strategies for hydro-electric generators. *Operations Research* 51, 602-612.

[18] Scott, T. J., Read, E. G. (1996): Modelling hydro reservoir operation in a deregulated electricity market. *International Transactions in Operations Research* 3, 243–253.

[19] Turgeon, A (1980): Optimal operation of multireservoir power systems with stochastic inflows. *Water Resources Research* 16, 275-283.

[20] Wallace, S. W., Fleten, S.-E. (2003): Stochastic programming models in energy. In: *Stochastic Programming: Handbooks in OR & MS, Vol. 10* (A. Rusczynski, A. Shapiro eds.), Elsevier Science .