Control of Energy Storage with Market Impact: Lagrangian Approach and Horizons

James Cruise

Department of Actuarial Mathematics and Statistics, Heriot-Watt University, Riccarton, Edinburgh, EH14 4AS, UK,
r.cruise@hw.ac.uk

Lisa Flatley

Mathematics Institute, Zeeman Building, University of Warwick, Coventry, CV4 7AL, UK, L.Flatley@warwick.ac.uk

Richard Gibbens

Computer Laboratory, University of Cambridge, William Gates Building, 15 JJ Thomson Avenue, Cambridge, CB3 0FD, UK,
richard.gibbens@cl.cam.ac.uk

Stan Zachary

Department of Actuarial Mathematics and Statistics, Heriot-Watt University, Riccarton, Edinburgh, EH14 4AS, UK,
s.zachary@gmail.com

We study the control of large scale energy storage operating in a market. Re-optimization of deterministic models is a common pragmatic approach when prices are stochastic. We apply Lagrangian theory to develop such a model and to establish decision and forecast horizons when storage trading affects these prices, an important aspect of some energy markets. The determination of these horizons also provides a simple and efficient algorithm for the determination of the optimal control. The forecast horizons vary between one and fifteen days in realistic electricity storage examples. These examples suggest that modelling price impact is important.

Key words: Control, Storage, Arbitrage, Energy

1. Introduction

Electricity storage is likely to play a significant role in balancing future energy systems. Often, much of the value of large-scale storage (e.g. pumped storage and hydro) may be captured in price arbitrage. In the present paper we study the optimal control of storage making its money by buying electricity when it is cheap, and selling it when it is expensive. Our model includes both capacity and rate constraints, and the activities of the store are of a sufficient magnitude as to have market impact, thereby leading to nonlinear convex cost functions associated with buying and selling. (For example, in many European markets there are substantial differences between

day and night electricity prices, while the volume of available storage is sufficient to significantly reduce this differential—see Newbery et al. 2013.) Market prices may be modelled as stochastic. In this case, for some applied storage control problems, an exact stochastic dynamic programming approach may be possible—see, e.g. Secomandi (2010), Bäuerle and Riess (2016). An explicitly stochastic approach to the control of storage with market impact is given by Felix et al. (2012). In practice assumed probability distributions calibrated to data may be incorrect—for a discussion of the potentially significant consequences of this in the context of energy storage, see Secomandi et al. (2015). Thus a common pragmatic approach in a stochastic environment is the use of deterministic models with re-optimisation at successive time steps—see, for example, Lai et al. (2010), Secomandi (2010), Wu et al. (2012) and, for more recent further analysis, Secomandi (2015). However, the literature appears to be missing a deterministic re-optimization method that incorporates market impact. The present paper fills this gap.

Thus we study a deterministic model, in which it is assumed that the above convex cost functions are known in advance. We develop the strong Lagrangian theory associated with the optimal control over some given time period, and use this theory to determine a running *forecast horizon* beyond which it is not necessary to know future buying and selling costs in order to determine the current action. We give a forward algorithm for the determination of both the optimal control and the forecast and associated decision horizons (see Section 4 for formal definitions). This algorithm reduces to the solution of a finite number of instances of the problem in which the storage is not subject to capacity constraints. We use this result to provide a bound on the computational effort involved. In the realistic electricity storage examples we study, the forecast horizon varies between one and fifteen days. These examples suggest that modelling price impact is important.

Our model may be viewed as an instance of the classical wheat trading, or warehouse, model (see, for example, Hartl 1986, Sethi and Thompson 2000 and the references therein). In the present model the storage, or inventory, process is subject to capacity constraints (but not holding costs) and we extend the classical model by allowing market impact. The existence of forecast and decision horizons is here a natural consequence of the bounds on the storage process (as is clear from the role played by the constraints in our algorithm), coupled with the assumed convexity of the cost functions (see also Garcia and Smith 2000). Similar horizons exist for many models in the existing literature in which the level of some stored quantity is controlled. In many cases, e.g. the early production planning model of Modigliani and Hohn (1955), and the wheat trading model of Hartl (1986), the storage bound is one-sided, corresponding only to the requirement that the stored quantity should be nonnegative; however, in such models storage holding costs increase with the quantity stored. As in the above papers, the forecast horizons obtained for such models are typically weak, in the sense that some mild conditions are required on costs beyond the forecast

horizon. For the wheat trading model, an example where storage is subject to a capacity constraint is given by Sethi and Thompson (2000) and a forecast horizon is obtained without the need for such additional conditions. A paper by Hartl (1988) considers the more general wheat trading model of Hartl (1986) with the addition of a capacity constraint, but does not give a general algorithm. A comprehensive review on the identification and use of such forecast horizons is given by Chand et al. (2002).

Similar storage control problems to those of the present paper are studied by Bannister and Kaye (1991), who consider only piecewise linear cost functions, and by Steffen and Weber (2016) who consider the price-taker case of linear cost functions (see also the references in the latter paper). The approaches of these papers also rely on the identification of strong Lagrangian, or Kuhn-Tucker, parameters analogous to those identified in this paper; however, these approaches are otherwise quite different and do not explicitly identify the forecast horizon of the present approach.

Aside from the control problem studied here, there is a considerable literature on the market impact of storage—notably in the context of energy storage (see Newbery et al. 2013 and the references therein)—and on its wider economic consequences, e.g. through price smoothing, within a competitive environment (see in particular Sioshansi 2010, Felix et al. 2012, Chaton and Durand-Viel 2013).

An earlier paper by Cruise et al. (2014) considered the application of strong Lagrangian theory to the present problem, gave sufficient conditions for a control strategy and associated Lagrange multipliers to solve it, and outlined how these might be identified. Here we give a more extensive Lagrangian treatment, enhancing also the model to allow for more general convex cost functions and for time-dependent leakage.

The organisation of the present paper is as follows. Sections 2 and 3 respectively formulate the mathematical problem for analysis and develop the relevant strong Lagrangian theory. Section 4 proves the existence of running forecast and decision horizons, develops an algorithm for the determination of both horizons and optimal control, and provides a bound on the associated computational effort. Section 5 includes realistic examples, based on real data for UK electricity prices. The e-companion to this paper gives the proofs of Theorems 1 and 2 and other results.

2. Problem formulation

Our model is a modest generalisation of that of Cruise et al. (2014), in that we here construct arguments more carefully to allow for quite general convex cost functions, and we further model leakage over time. It is convenient to think of the available storage as a single store, seeking to maximise the profit which can be made by buying and selling. We assume that the activities of

the store are sufficiently significant as to have market impact, so that the store sees nonlinear cost functions.

We work in discrete time $t = 0, 1, \ldots, T$ where $T \geq 2$ denotes the final time horizon. We assume that the store has a total *capacity* of E (which, in the context of an energy system, would be total energy which could be stored) and input and output *rate constraints* of P_i and P_o respectively. We consider two types of (in)efficiency associated with the store. The first of these (and usually much the more significant in practice) is a *time-independent* round-trip efficiency $\eta \in (0,1]$ which may be defined as the fraction of energy bought which is available to sell. This may be incorporated directly into the cost functions C_t introduced below by suitably rescaling selling prices relative to buying prices. The second type of (in)efficiency may be regarded as *leakage* over time, and is modelled by assuming that at each successive time instant there is lost a fraction $1 - \rho$, where $\rho \in (0,1]$, of whatever is in the store at that time.

Let $X = \{x : -P_o \leq x \leq P_i\}$. Both buying and selling prices at time t may be represented by a single cost function C_t, which we assume to be convex, and is such that $C_t(x)$ is the cost at time t of increasing the level of the store contents (after any leakage—see below) by x, positive or negative. Typically—in a conventional store and with positive prices—we have that each function C_t is increasing and that $C_t(0) = 0$; then, for positive x, $C_t(x)$ is the cost of buying x units (for example of energy) and, for negative x, $C_t(x)$ is the negative of the reward for selling $-x$ units; however, for some applications, the interpretation of the functions C_t may vary slightly from this, and only the convexity condition is required. This convexity assumption corresponds, for each time t, to an increasing cost to the store of buying each additional unit, a decreasing revenue obtained for selling each additional unit, and every unit buying price being at least as great as every unit selling price. Incorporating the time-independent (or "round-trip") efficiency η into the cost functions C_t, as discussed above, automatically preserves convexity whenever these cost functions are increasing. (While the model formally allows the possibility that some of the functions C_t might be decreasing—corresponding to negative prices—the inclusion of round-trip efficiency $\eta < 1$ as above would typically modify such functions so as to violate the convexity assumption. For a discussion of the effect of negative prices on the nature of optimal policies, see Zhou et al. 2016.)

Denote the successive levels of the store by a vector $S = (S_0, \ldots, S_T)$ where S_t is the level of the store at each time t. For each $t \geq 1$, define also

$$x_t(S) = S_t - \rho S_{t-1}. \tag{1}$$

Here ρ is the time-dependent leakage measure defined above, so that $x_t(S)$ represents the addition to the store at time t. It is convenient to assume that both the initial level S_0 and the final level S_T of the store are fixed in advance at $S_0 = S_0^*$ and $S_T = S_T^*$. The optimisation problem of interest may then be expressed as:

P: (given the convex functions C_t) choose S so as to minimise

$$\sum_{t=1}^{T} C_t(x_t(S)) \tag{2}$$

subject to the capacity constraints

$$S_0 = S_0^*, \qquad S_T = S_T^*, \qquad 0 \le S_t \le E, \quad 1 \le t \le T-1, \tag{3}$$

and the rate constraints

$$x_t(S) \in X, \qquad 1 \le t \le T. \tag{4}$$

We shall say that a vector S is *feasible* for the problem **P** if it satisfies both sets of constraints (3) and (4). We assume that the set of feasible vectors S is nonempty. This set is then closed and convex and the function defined by (2) is convex, and strictly so when the functions C_t are strictly convex. Hence a solution to the problem **P** always exists, and is unique when the functions C_t are strictly convex.

In the case where the cost functions C_t are linear, or piecewise linear, the problem **P** may be reformulated as a linear programming problem, and solved by, for example, the use of the minimum cost circulation algorithm (see, e.g., Boyd and Vandenberghe 2004, Ahuja et al. 1993).

3. Lagrangian formulation and characterisation of solution

We apply strong Lagrangian theory (see Boyd and Vandenberghe 2004, Whittle 1971) to the problem **P** defined above. Theorem 1, which is a generalisation of a result given by Cruise et al. (2014) and which is required in the present paper, gives sufficient conditions for a value S^* of S to solve the problem. However, we give in the e-companion to this paper a proof which illustrates the result as essentially an application of the Lagrangian sufficiency theorem (see Whittle 1971 or Courcoubetis and Weber 2003). The Lagrangian theory is here used to manage the capacity constraints only (dealing with the rate constraints in this way does not result in a simpler theory).

THEOREM 1. *Suppose that there exists a vector $\mu^* = (\mu_1^*, \ldots, \mu_T^*)$ and a value $S^* = (S_0^*, \ldots, S_T^*)$ of S such that*

(i) S^* is feasible for the stated problem **P**,

(ii) for each t with $1 \le t \le T$, $x_t(S^*)$ minimises $C_t(x) - \mu_t^* x$ in $x \in X$,

(iii) the pair (S^*, μ^*) satisfies the complementary slackness conditions, for $1 \le t \le T-1$,

$$\begin{cases} \rho\mu_{t+1}^* = \mu_t^* & \text{if } 0 < S_t^* < E, \\ \rho\mu_{t+1}^* \le \mu_t^* & \text{if } S_t^* = 0, \\ \rho\mu_{t+1}^* \ge \mu_t^* & \text{if } S_t^* = E. \end{cases} \tag{5}$$

Then S^* solves the stated problem **P**.

The vector μ^* has the interpretation that, for each time t, the quantity μ_t^* may be regarded as a notional reference value per unit volume in storage *at that time*, i.e. the rate at which the residual value of the store, optimally operated to time T, increases with respect to increasing the level of the store at time t. A similar parameter is identified elsewhere in the storage literature—see, for example, Bannister and Kaye (1991). Thus, in the condition (ii) of the theorem, $C_t(x)$ is the cost at time t of increasing the level of the store by x (again positive or negative) and $\mu_t^* x$ may be regarded as a current offsetting measure of value added to the store; the quantity $C_t(x) - \mu_t^* x$ is then to be minimised in $x \in X$.

Theorem 1 does not require the assumed convexity of the cost functions C_t. This convexity is, however, sufficient to ensure the existence of the vector μ^* of that theorem. This follows directly from Theorem 2 in Section 4. It may also be deduced from strong Lagrangian convexity arguments (see Boyd and Vandenberghe 2004 or Whittle 1971).

4. Determination of forecast horizon and optimal control

We show how to determine the running forecast horizon and to use this to construct a pair (S^*, μ^*) satisfying the conditions of Theorem 1, so that in particular S^* is the optimal solution to the problem **P**. Specifically we show how to identify a pair of times $1 \le \tau \le \bar{\tau} \le T$ such that, given the cost functions C_t, $t \le \bar{\tau}$, the initial segment $((S_1^*, \mu_1^*), \ldots, (S_\tau^*, \mu_\tau^*))$ of the pair (S^*, μ^*) is independent of the cost functions C_t for times $t > \bar{\tau}$; the times $\bar{\tau}$ and τ are then respectively initial *forecast* and *decision horizons* (see Chand et al. 2002). We also show how to construct this initial segment, summarising the steps in Algorithm 1. This procedure may then be restarted at the time τ to define the next segment of (S^*, μ^*), and so on, thereby defining an algorithm for the determination of the entire solution $((S_1^*, \mu_1^*), \ldots, (S_T^*, \mu_T^*))$ of **P**.

We assume that the cost functions C_t are strictly convex; we show how to relax this assumption in the e-companion to this paper. For any t such that $1 \le t \le T$ and any scalar μ, define $\hat{x}_t(\mu)$ to be the unique value of x which minimises $C_t(x) - \mu x$ in $x \in X$. Then the function $\hat{x}_t(\cdot)$ is continuous and increasing (though not necessarily strictly so). Again for any scalar μ, define a succession of levels $S(\mu) = (S_0(\mu), \ldots, S_T(\mu))$ of the store (not necessarily satisfying the capacity constraints (3)) by

$$S_0(\mu) = S_0^*, \qquad S_t(\mu) = \rho S_{t-1}(\mu) + \hat{x}_t(\rho^{1-t}\mu), \quad t = 1, \ldots, T. \tag{6}$$

For each t, the function $S_t(\cdot)$ is similarly continuous and increasing and, by the definition of the functions $x_t(\cdot)$, the path $(S_0(\mu), \ldots, S_T(\mu))$ automatically satisfies the rate constraints (4). Allow also $\mu = -\infty$ and $\mu = \infty$ and, for $t = 1, \ldots, T$, define $S_t(-\infty) = -\infty$ and $S_t(\infty) = \infty$. From (6) and the monotonicity of the functions $x_t(\cdot)$, for all $t_1 < t_2$,

$$\mu_1 < \mu_2, \quad S_{t_1}(\mu_1) < S_{t_1}(\mu_2) \qquad \Rightarrow \qquad S_{t_2}(\mu_1) < S_{t_2}(\mu_2). \tag{7}$$

It is convenient to define, for each time $t = 1, \ldots, T$, the quantities a_t and b_t to be respectively the upper and lower bounds on the permissible values of S_t, i.e. $a_t = 0$ and $b_t = E$ for $t = 1, \ldots, T-1$ and $a_T = b_T = S_T^*$. For each time $t = 1, \ldots, T$, let (the scalar quantity) $\mu^{l,t}$ be such that $S_t(\mu^{l,t}) = a_t$; in the event that $\mu^{l,t}$ fails to be thus uniquely defined (as may happen when the functions C_t fail to be differentiable) take $\mu^{l,t}$ to be the maximum value satisfying the above condition; if the rate constraints (defined by X) are such that there is no such $\mu^{l,t}$ define instead $\mu^{l,t} = -\infty$. Similarly, for each $t = 1, \ldots, T$, let $\mu^{u,t}$ be such that $S_t(\mu^{u,t}) = b_t$; in the event that $\mu^{u,t}$ fails to be uniquely defined take $\mu^{u,t}$ to be the minimum value satisfying the above condition; if there is no such $\mu^{u,t}$ define $\mu^{u,t} = \infty$. Since the functions $S_t(\cdot)$ are increasing, it follows from these definitions that, for $t = 1, \ldots, T$,

$$\mu > \mu^{l,t} \;\Leftrightarrow\; S_t(\mu) > a_t \qquad \text{and} \qquad \mu < \mu^{u,t} \;\Leftrightarrow\; S_t(\mu) < b_t. \tag{8}$$

In particular

$$\mu^{l,t} < \mu^{u,t}, \quad 1 \le t \le T-1, \qquad \mu^{l,T} \ge \mu^{u,T}, \tag{9}$$

where, since the feasible region for the problem **P** is assumed to be nonempty and since this would remain the case were the capacity constraints to be dropped, $\mu^{l,T}$ and $\mu^{u,T}$ are necessarily finite.

For each $t = 1, \ldots, T$ define also

$$\bar{\mu}^{l,t} = \max_{1 \le t' \le t} \mu^{l,t'}, \qquad \bar{\mu}^{u,t} = \min_{1 \le t' \le t} \mu^{u,t'}; \tag{10}$$

it is convenient to define also $\bar{\mu}^{l,0} = -\infty$ and $\bar{\mu}^{u,0} = \infty$. The sequence of partial maxima $\{\bar{\mu}^{l,t}\}$ is increasing and the sequence $\{\bar{\mu}^{u,t}\}$ of partial minima is decreasing. We refer to the times $t \ge 1$ such that $\mu^{l,t} = \bar{\mu}^{l,t}$ as *lower record times* and the times $t \ge 1$ such that $\mu^{u,t} = \bar{\mu}^{u,t}$ as *upper record times*.

We now define $\bar{\tau}$ to be the *first* time $t \le T$ such that $\bar{\mu}^{l,t} \ge \bar{\mu}^{u,t}$. It follows from (9) that the time $\bar{\tau}$ is well-defined and that $\bar{\tau} \ge 2$. It further follows (see below) that *exactly one* of the following three conditions holds:

(a) $\bar{\mu}^{u,\bar{\tau}} \le \bar{\mu}^{l,\bar{\tau}-1}$, in which case define also τ to be the *greatest* lower record time $t < \bar{\tau}$, and define the parameter $\hat{\mu} = \mu^{l,\tau} = \bar{\mu}^{l,\tau} = \bar{\mu}^{l,\bar{\tau}-1}$;

(b) $\bar{\mu}^{l,\bar{\tau}} \ge \bar{\mu}^{u,\bar{\tau}-1}$, in which case define also τ to be the *greatest* upper record time $t < \bar{\tau}$, and define the parameter $\hat{\mu} = \mu^{u,\tau} = \bar{\mu}^{u,\tau} = \bar{\mu}^{u,\bar{\tau}-1}$;

(c) neither (a) nor (b) holds, in which case necessarily $\bar{\tau} = T$; here define also $\tau = T$, and define the parameter $\hat{\mu}$ to be such that $S_T(\hat{\mu}) = S_T^*$.

For $\bar{\tau} \le T-1$ that exactly one of the conditions (a) or (b) holds follows from (9) and the definition of the sequences $\{\bar{\mu}^{l,t}\}$ and $\{\bar{\mu}^{u,t}\}$. For $\bar{\tau} = T$, it necessarily follows that $\bar{\mu}^{l,T-1} < \bar{\mu}^{u,T-1}$. Since this implies that at the last time $t \le T-1$ such that $\mu^{l,t} = \bar{\mu}^{l,T-1}$, we have, from (8), that $S_t(\bar{\mu}^{l,T-1}) < S_t(\bar{\mu}^{u,T-1})$, it follows from (7) that also $S_T(\bar{\mu}^{l,T-1}) < S_T(\bar{\mu}^{u,T-1})$. We now have that, for $\bar{\tau} = T$, the

condition (a), (b) or (c) holds according as $S_T^* \leq S_T(\bar{\mu}^{l,T-1})$, $S_T^* \geq S_T(\bar{\mu}^{u,T-1})$, or $S_T(\bar{\mu}^{l,T-1}) < S_T^* < S_T(\bar{\mu}^{u,T-1})$.

We now give the algorithm for the determination of the times $\bar{\tau}$ and τ, and hence also the parameter $\hat{\mu}$. We then define also $((S_1^*, \mu_1^*), \ldots, (S_\tau^*, \mu_\tau^*))$ by

$$\mu_t^* = \rho^{1-t}\hat{\mu}, \qquad S_t^* = S_t(\hat{\mu}), \qquad\qquad 1 \leq t \leq \tau. \tag{11}$$

Iterative application of this algorithm yields a pair of vectors (S^*, μ^*) which we show in Theorem 2 satisfies the conditions of Theorem 1 and so forms the optimal solution of the problem **P**. Since $\bar{\tau}$, τ and $\hat{\mu}$ do not depend on the cost functions C_t for times $t > \bar{\tau}$, it follows that $\bar{\tau}$ and τ are respectively initial forecast and decision horizons.

The algorithm proceeds inductively by considering successive times $1 \leq t < T$ (the "while" loop of Algorithm 1). At each such time t it is checked whether $\tau = t$ (via either the condition (a) or the condition (b) above) and if so the corresponding value of $\bar{\tau} > t$ is identified. Thus suppose it is established that the time $t < T$ is such that $\tau \geq t$. The latter condition implies in particular that $\bar{\tau} > t$, i.e. that $\bar{\mu}^{l,t} < \bar{\mu}^{u,t}$. It now follows from the definitions of the times $\bar{\tau}$ and τ above that, for any $\bar{t} > t$, in order that $\bar{\tau} = \bar{t}$ with $\tau = t$ being defined via the condition (a) above, it is necessary and sufficient that t should be a lower record time, i.e., from (8),

$$S_t(\bar{\mu}^{l,t-1}) \leq a_t, \tag{12}$$

and further that $\mu^{l,t'} < \bar{\mu}^{l,t} < \mu^{u,t'}$ for $t < t' < \bar{t}$ and that $\bar{\mu}^{l,t} \geq \mu^{u,\bar{t}}$. Given that t is a lower record time, it follows from (8) that these latter two conditions are equivalent to

$$a_{t'} < S_{t'}(\bar{\mu}^{l,t}) < b_{t'}, \quad t < t' < \bar{t}, \qquad\qquad S_{\bar{t}}(\bar{\mu}^{l,t}) \geq b_{\bar{t}}. \tag{13}$$

(Thus the steps 3.–8. of Algorithm 1 check, for the current value of t, whether $\tau = t$ via the condition (a), and if so determine also $\bar{\tau}$ and $\hat{\mu}$ as required; in this case the algorithm then stops.)

Similarly, for any $\bar{t} > t$, in order that $\bar{\tau} = \bar{t}$ with $\tau = t$ being defined via the condition (b) above, it is necessary and sufficient that t should be an upper record time, i.e. that

$$S_t(\bar{\mu}^{u,t-1}) \geq b_t, \tag{14}$$

and further that

$$a_{t'} < S_{t'}(\bar{\mu}^{u,t}) < b_{t'}, \quad t < t' < \bar{t}, \qquad\qquad S_{\bar{t}}(\bar{\mu}^{u,t}) \leq a_{\bar{t}}. \tag{15}$$

(Thus the steps 9.–14. of Algorithm 1 similarly check whether $\tau = t$ via the condition (b), and if so determine again $\bar{\tau}$ and $\hat{\mu}$; in this case the algorithm again then stops.)

Algorithm 1 Algorithm for the determination of the initial forecast and decision horizons $\bar{\tau}$ and τ respectively, and initial segment $((S_1^*, \mu_1^*), \ldots, (S_\tau^*, \mu_\tau^*))$ of the pair (S^*, μ^*) of Theorem 1.

INPUT: S_0^*, S_T^*, E, ρ, functions $\hat{x}_t(\mu)$ for $t = 1, \ldots, T$

1: set $t = 1$

2: **while** $t < T$ **do**

3:      calculate $\bar{\mu}^{l,t}$ as given by (10)

4:      **for** $\bar{t} = t + 1, \ldots, T$ **do**

5:          **if** equations (12) and (13) hold **then**

6:              set $\bar{\tau} = \bar{t}$ and set $\tau = t$

7:              set $\hat{\mu} = \bar{\mu}^{l,t}$ and calculate $((S_1^*, \mu_1^*), \ldots, (S_\tau^*, \mu_\tau^*))$ via (6) and (11)

8:              STOP

9:          **end if**

10:      **end for**

11:      calculate $\bar{\mu}^{u,t}$ as given by (10)

12:      **for** $\bar{t} = t + 1, \ldots, T$ **do**

13:          **if** equations (14) and (15) hold **then**

14:              set $\bar{\tau} = \bar{t}$ and set $\tau = t$

15:              set $\hat{\mu} = \bar{\mu}^{u,t}$ and calculate $((S_1^*, \mu_1^*), \ldots, (S_\tau^*, \mu_\tau^*))$ via (6) and (11)

16:              STOP

17:          **end if**

18:      **end for**

19:      set $t = t + 1$

20: **end while**

21: set $\tau = \bar{\tau} = T$

22: calculate $\hat{\mu}$ such that $S_T(\hat{\mu}) = S_T^*$ where $S_T(\hat{\mu})$ is given via (6)

23: calculate $((S_1^*, \mu_1^*), \ldots, (S_\tau^*, \mu_\tau^*))$ via (6) and (11)

Finally in the event that the algorithm does not find $\tau = t$ for any time $t < T$ (so that the "while" loop of Algorithm 1 terminates with $t = T$), then necessarily $\bar{\tau}$ and τ are defined via the condition (c) above with $\tau = \bar{\tau} = T$. In all cases $((S_1^*, \mu_1^*), \ldots, (S_\tau^*, \mu_\tau^*))$ is then determined as above.

For $\tau < T$ the algorithm may then be restarted at the time τ with the time 0 replaced by the time τ and the initial level S_0^* replaced by the level S_τ^*, and this process repeated at such subsequent times as necessary in order to construct the entire pair $(S^*, \mu^*) = ((S_1^*, \mu_1^*), \ldots, (S_T^*, \mu_T^*))$. We now have the following theorem, the proof of which is given in the e-companion, and from which, as

previously indicated, it follows immediately that the time $\bar{\tau}$ is indeed a forecast horizon for the determination of the solution of $\mathbf{P}$ to the decision horizon τ.

THEOREM 2. *Under the assumed strict convexity of the functions C_t, the pair (S^*, μ^*) constructed by the above algorithm satisfies the conditions of Theorem 1. In particular S^* is the (optimal) solution to the problem $\mathbf{P}$.*

REMARK 1. It follows from the above algorithm, by considering separately the conditions (a)–(c) for the definition of τ and using the continuity of $S(\mu)$ in μ, that the defined scalar $\hat{\mu}$ may be varied by suitable variation of the cost functions C_t for $t \geq \bar{\tau}$. Then also $S_1^* = S_1(\hat{\mu})$ necessarily varies (at least when C_1 is differentiable). In this sense the identified forecast horizon $\bar{\tau}$ is the shortest possible for the identification of even (S_1^*, μ_1^*).

REMARK 2. It should further be clear from the definition of the time $\bar{\tau}$—and is illustrated in the examples of Section 5—that the length of the forecast horizon is typically of the same order as that of the time for the optimally controlled store to empty or fill. In particular when the store is emptying and filling in a completely periodic manner, it is not difficult to see that the length of the forecast horizon does not exceed the period of the cycle.

We now provide a simple bound for the work involved in the implementation of the above algorithm. We assume an ability to evaluate as necessary the quantities $S_t(\mu)$ defined by (6). This requires only an ability to evaluate, again as necessary, the quantities $\hat{x}_t(\mu)$ minimising $C_t(x) - \mu x$ in $x \in X$; depending on the cost functions C_t the functions $\hat{x}_t$ may be available analytically or numerically (in particular when C_t is differentiable $\hat{x}_t(\mu)$ is simply the value of $x \in X$ whose marginal cost $C_t'(x)$ is closest to μ). In the determination of the initial forecast horizon $\bar{\tau}$ and decision horizon τ, the algorithm is driven by the determination of the sequences $\{\bar{\mu}^{l,t}\}$ and $\{\bar{\mu}^{u,t}\}$ up to the time τ. For the former, for each time t we have $\bar{\mu}^{l,t} = \bar{\mu}^{l,t-1}$ except perhaps where t is a lower record time, i.e. the condition (12) holds. At such a time t we have $\bar{\mu}^{l,t} = \mu^{l,t}$; the latter quantity is the defined as solution of $S_t(\mu^{l,t}) = a_t$ and $S_t(\mu)$ is (continuous and) increasing in μ. Similar remarks apply to the determination of the sequence $\{\bar{\mu}^{u,t}\}$. It follows from these observations and from the specification of the algorithm as summarised in Algorithm 1 that, in the determination of $\bar{\tau}$ and τ and so also of the initial segment $((S_1^*, \mu_1^*), \ldots, (S_\tau^*, \mu_\tau^*))$ of the (optimal) solution of $\mathbf{P}$, the computation involved consists of at most 2τ *one-dimensional* searches for the zero of a monotonic function—together with a finite number of simple evaluations of $S_t(\mu)$ for given values of t and μ and a finite number of binary comparisons. The algorithm is now restarted at the time τ, and so ultimately a maximum of $2T$ such one-dimensional searches are required in order to determine the optimal control to the time T. In the case where the cost functions C_t have

an appropriate analytical form, e.g. are linear or quadratic (see Section 5 for a justification of the latter as first approximation to market impact), the quantities $\mu^{l,t}$ and $\mu^{u,t}$ may be determined analytically and so only a finitely terminating calculation is required in the operation of the entire algorithm.

5. Examples

We illustrate some of our results with an example storage facility which has market impact. We use half-hourly time units and a cost series $(p_1, \ldots, p_T)$ corresponding to the real half-hourly spot market wholesale electricity prices in Great Britain for the year 2011, as supplied, along with corresponding total GB demand data, by National Grid plc—see National Grid (2012). (Spot prices are readily available and used for convenience; ideally one might use forward prices or forecasted prices.) These prices show a strong daily cyclical behaviour. We assume that the store is large enough to have market impact on prices, but small enough in relation to the rest of the network that, at each time t, the unit price at which the store buys sufficient energy to increase its level by $x > 0$ units may be approximated by a linear function $p_t + p'_t x$, where $p'_t \geq 0$ is a measure of the market impact of the store on the price at that time. This linearised dependence of price on modest variations in overall traded volumes on energy seems a reasonable first approximation to market impact and is consistent with the existing energy economics literature—see, for example, Sioshansi (2010, 2014) and the references therein. It follows that the corresponding cost $C_t(x)$ is quadratic in x. We assume the same linear dependence of price on quantity sold for $x < 0$; however, since the round-trip efficiency η of the store means that it only sells back to the market a fraction η of what it buys, the complete cost function C_t is assumed to be given by

$$C_t(x) = \begin{cases} (p_t + p'_t x)x & \text{if } x \geq 0, \\ (p_t + \eta p'_t x)\eta x & \text{if } x < 0. \end{cases} \tag{16}$$

In the following examples, we assume further that each p'_t is proportional to the wholesale price p_t at that time, so that $p'_t = \lambda p_t$ for some $\lambda \geq 0$. This reflects the intuition that the market becomes more price-responsive when prices are high. We assume a common input and output rate constraint $P_i = P_o = P$ and, as before, denote by E the capacity of the store. Finally, we assume throughout that there is no leakage from the store over time, i.e. that $\rho = 1$. This assumption is consistent with the existing literature on energy storage, where round-trip inefficiency (which we do model) is significant, but where gradual leakage over time is much less so and not usually modelled—see, for example, Felix et al. (2012), Sioshansi (2010, 2014).

The optimal strategy associated with the cost function (16) is shown in Figure 1 (the upper plot in each of the four panels) for various choices of parameters. We present the behaviour of the store over the month of December. The optimisation is started at a point in time sufficiently prior
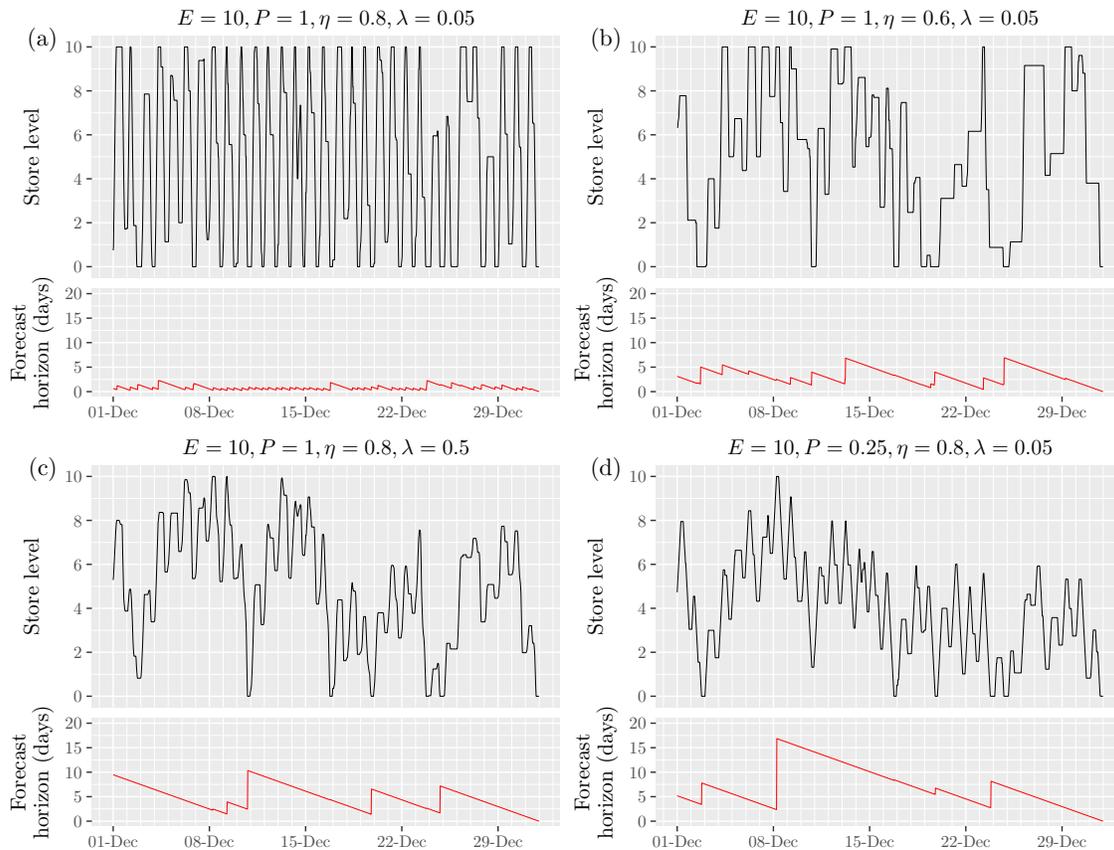
Figure 1 Examples in which the parameters associated with the store are varied. In each case, the upper plot shows the optimal level of storage and the lower plot shows the forecast horizon required at each stage of the optimisation.

to the beginning of that month that the optimal behaviour of the store throughout that month is independent of the store level at the earlier starting time. (Since the problem **P** is invariant under time reversal, this "lead" interval is of the same order as the running forecast horizon for the solution of the problem.)

The plot in panel (a) corresponds to a "base" case, with the parameter choices $E = 10$, $P = 1$, $\eta = 0.8$ and $\lambda = 0.05$. The time $E/P = 10$ half-hours units for the store to completely fill or empty and the round-trip efficiency of 0.8 correspond approximately to the Dinorwig pumped storage facility in Snowdonia in North Wales. On the assumption that the observed relationship between price and total GB demand throughout the period of the example is approximately that which would also obtain at any point in time as demand was varied, the available price-demand data appear reasonably compatible with the modelling assumption $p_t' = \lambda p_t$. In particular the prices p_t at their daily peak are approximately twice those at their nightly minimum, and correspond to an approximate 25 GW variation in demand. The choice of market impact factor $\lambda = 0.05$ then corresponds to the power units of our example being a little less than 2 GW; since the parameter

$P = 1$, this implicit 2 GW power unit is also the rate constraint for our example, and again corresponds closely to that of the Dinorwig facility. The upper portion of the plot in panel (a) shows the variation of the store level with time t, while the lower portion shows, for each time t, the corresponding forecast horizon at that time as defined in Section 4. It is seen that, under the optimal strategy, the store usually completely empties and fills on a daily cycle, with some lull in activity over the Christmas period. As might then be expected (again see Section 4) the forecast horizon necessary for an optimal decision is of the order of a day or so.

The plots in the remaining three panels of Figure 1 are each formed by varying one of the parameters of the base case example, in each case in such a way that the store is less active. The plot in the panel (b) corresponds to a reduction in the round-trip efficiency of the store from $\eta = 0.8$ to $\eta = 0.6$. (The latter figure is something of a lower bound: the round-trip efficiencies of nearly all forms of storage technologies in significant current use are in the region 0.70–0.85 (see Manuel 2014). Here it is seen that the store level cycles less frequently and tends to remain at the same value for longer periods of time than in the base case; further the forecast horizons necessary for optimal decision making are significantly longer than in the base case. The plot in panel (c) corresponds to an increase in the "market impact" factor from $\lambda = 0.05$ to $\lambda = 0.5$, while that in panel (d) corresponds to a tightening of the rate constraint from $P = 1$ to $P = 0.25$, the resulting ratio $E/P = 40$ half-hours corresponding to the Cruachan and Foyers pumped storage facilities in Scotland. In both cases the store is almost continuously active but trades at lower volumes than in the base case; consequently forecast horizons are considerably longer. The broad similarity of the behaviour in these two examples may be explained by noting that an increased market impact factor acts to slow down the activity rate of the store in much the same way as a tightening of the rate constraint.

Finally we consider the effect of failing to account for market impact when the latter is present. For the example here, and for the base case parameter choices $E = 10$, $P = 1$, $\eta = 0.8$, Figure 2 shows the total annual profit (negative cost) of the store as a function of the market impact factor λ, the behaviour of the store being optimised over the entire year 2011. If the units of power of the example are gigawatts, then the example corresponds to a store of the approximate size of Dinorwig, and the units in which the profit is recorded are millions of pounds. Figure 2 further shows, again as a function of λ, the corresponding profit when the behaviour of the store is optimised on the assumption $\lambda = 0$ but in which the profit of the store is then calculated according to the actual value of λ. The latter profit decreases linearly in λ and becomes negative at around $\lambda = 0.12$—a value which, as argued above, is not at all unrealistic in the presence of significant storage. As noted above, forecast horizons increase with increasing λ: mean forecast horizons for $\lambda = 0, 0.05$, 0.10, and 0.15 are respectively 0.87, 1.40, 2.50, and 3.26 days.
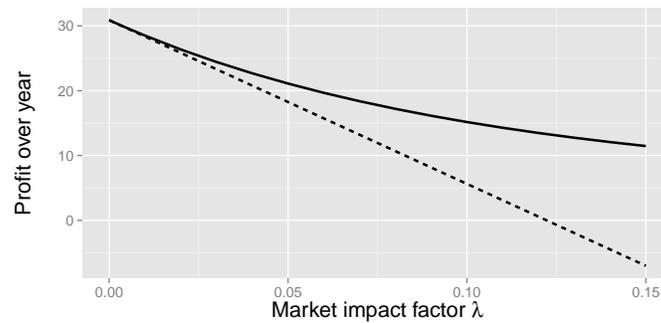
Figure 2   Optimised annual profit as function of market impact factor λ (solid line) and corresponding annual profit when market impact has been ignored in performing the optimisation (dashed line).

6. Conclusion

In the present paper we have developed the strong Lagrangian theory of the optimal control of energy storage which is used for arbitrage and whose activities are sufficiently significant as to have market impact. We have further shown how this theory may be used to determine a simple and efficient algorithm for the identification of that control and of the associated forecast and decision horizons. We have given examples based on real GB electricity price data and realistic storage parameters. These show the relevance of modelling market impart to the optimal control.

Acknowledgments

References

Ahuja RK, Magnanti TL, Orlin JB (1993) *Network Flows: Theory, Algorithms and Applications* (Prentice Hall).

Bannister CH, Kaye RJ (1991) A rapid method for optimization of linear systems with storage. *Operations Research* 39(2):220–232.

Bäuerle N, Riess V (2016) Gas storage valuation with regime switching. *Energy Systems* 7(3):499–528.

Boyd S, Vandenberghe L (2004) *Convex Optimization* (Cambridge University Press).

Chand S, Hsu VN, Sethi S (2002) Forecast, solution, and rolling horizons in operations management problems: A classified bibliograph. *Manufacturing & Service Operations Management* 4(1):25–43.

Chaton C, Durand-Viel L (2013) Real asset valuation under imperfect competition: Can we forget about market fundamentals? *Journal of Economics & Management Strategy* 22(1):125–139.

Courcoubetis C, Weber R (2003) *Pricing Communication Networks: Economics, Technology and Modelling* (Wiley).

Cruise JR, Gibbens RJ, Zachary S (2014) Optimal control of storage for arbitrage, with applications to energy systems. *48th Annual Conference on Information Sciences and Systems (CISS)*, 1–6.

Felix B, Woll O, Weber C (2012) Gas storage valuation under limited market liquidity: an application in Germany. *The European Journal of Finance* 19(7–8):715–733.

Garcia A, Smith R (2000) Solving nonstationary infinite horizon dynamic optimization problems. *Journal of Mathematical Analysis and Applications* 244(2):304–317.

Hartl RF (1986) A forward algorithm for a generalized wheat trading model. *Zeitschrift für Operations Research* 30(3):A135–A144.

Hartl RF (1988) On forward algorithms for a generalized wheat trading model. *Engineering Costs and Production Economics* 15:367–370.

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

Manuel W (2014) Turlock Irrigation District: Energy Storage Study 2014, at http://www.energy.ca.gov/assessments/ab2514_reports/Turlock_Irrigation_District/2014-10-28_Turlock_Irrigation_District_Energy_Storage_Study.pdf.

Modigliani F, Hohn FE (1955) Production planning over time and the nature of the expectation and planning horizon. *Econometrika* 23(1):46–66.

National Grid (2012) Personal communication.

Newbery DM, Strbac G, Pudjianto D, Noël P, Co B, LeighFisher (2013) Benefits of an integrated european energy market: Final report for DG ENER, at https://ec.europa.eu/energy/sites/ener/files/documents/20130902_energy_integration_benefits.pdf.

Secomandi N (2010) Optimal commodity trading with a capacitated storage asset. *Management Science* 56(3):449–467.

Secomandi N (2015) Merchant commodity storage practice revisited. *Operations Research* 63(5):1131–1143.

Secomandi N, Lai G, Margot F, Scheller-Wolf A, Seppi D (2015) Merchant commodity storage and term-structure model error. *Manufacturing & Service Operations Management* 17(3):302–320.

Sethi SP, Thompson GL (2000) *Optimal Control Theory: Applications to Management Science and Economics, 2nd ed.* (Springer).

Sioshansi R (2010) Welfare impacts of electricity storage and the implications of ownership structure. *The Energy Journal* 31(2):173–198.

Sioshansi R (2014) When energy storage reduces social welfare. *Energy Economics* 41:106–116.

Steffen B, Weber C (2016) Optimal operation of pumped-hydro storage plants with continuous time-varying power prices. *European Journal of Operational Research* 252(1):308–321.

Whittle P (1971) *Optimization Under Constraints: Theory and Applications of Nonlinear Programming* (Wiley).

Wu OQ, Wang DD, Qin Z (2012) Seasonal energy storage operations with limited flexibility: the price-adjusted rolling intrinsic policy. *Manufacturing & Service Operations Management* 14(3):455–471.

Zhou Y, Scheller-Wolf A, Secomandi N, Smith S (2016) Electricity trading and negative prices: storage vs. disposal. *Management Science* 62(3):880–898.

James R Cruise is is an assistant professor of applied probability in the School of Mathematical and Computer Sciences at Heriot-Watt University. His research interests include the use of probabilistic methods to model energy markets, power systems and communication networks and the development of the associated probability theory.

Lisa Flatley is a research fellow in the Department of Mathematics at the University of Warwick. Her research interests are in the optimisation and control of complex energy networks.

Richard Gibbens is a reader in the Computer Laboratory at the University of Cambridge. His research interests are in the applications of probability and statistics to the control of complex transportation, communication and energy networks.

Stan Zachary is an honorary research fellow in the School of Mathematical and Computer Sciences at Heriot-Watt University and senior researcher in the School of Mathematics at the University of Edinburgh. His research interests are in probability and statistics with applications to the optimisation of complex systems.