Decentralised Control of Micro-Storage in the Smart Grid

Thomas D. Voice and **Perukrishnen Vytelingum** and **Sarvapali D. Ramchurn**
Alex Rogers and **Nicholas R. Jennings**
School of Electronics and Computer Science
University of Southampton
{tdv,pv,sdr,acr,nrj}@ecs.soton.ac.uk

Abstract

In this paper, we propose a novel decentralised control mechanism to manage micro-storage in the smart grid. Our approach uses an adaptive pricing scheme that energy suppliers apply to home smart agents controlling micro-storage devices. In particular, we prove that the interaction between a supplier using our pricing scheme and the actions of selfish micro-storage agents forms a globally stable feedback loop that converges to an efficient equilibrium. We further propose a market strategy that allows the supplier to reduce wholesale purchasing costs without increasing the uncertainty and variance for its aggregate consumer demand. Moreover, we empirically evaluate our mechanism (based on the UK grid data) and show that it yields savings of up to 16% in energy cost for consumers using storage devices with average capacity 10 kWh. Furthermore, we show that it is robust against extreme system changes.

1 Introduction

The need for sustainable future energy provision has driven a large research effort into the development of several intelligent electricity network technologies, collectively called the *smart grid* (US Department Of Energy 2003; Galvin and Yeager 2008; UK Department of Energy and Climate Change 2009). A major component of this future vision is that of energy storage. In particular, there is potential seen in the widespread adoption of small scale consumer storage devices (i.e., micro-storage), which would allow consumers to store electricity when demand is low, in order for it to be used during peak loads (Bathurst and Strbac 2003; Ramchurn et al. 2011a; Vytelingum et al. 2010). This technology has the added advantage that it requires no significant change in how home appliances are used, and thus allows consumers to respond to pricing signals with no impact on their own personal comfort. The intention is to lower peak demand, reducing the need to use expensive, carbon intensive “peaking plant” generators and, thus, lowering both carbon emissions and consumer energy costs. This scenario looks increasingly likely given the advent of consumer batteries, either stand-alone or for use in electric vehicles (EVs), able to hold enough energy to satisfy the needs of a home (den Bossche et al. 2006; Gerding et al. 2011). Moreover, the widespread deployment of EVs is likely to improve economies of scale and could potentially act as a source of cheaper (less efficient) used batteries in the long run. The downside to this trend is that a large number of batteries of various capacities, charge rates, and efficiencies are likely to be used and charged on the grid, potentially at the same time. This may result in unpredictable and peaky demand, which could increase carbon emissions, destabilise the grid, and even cause blackouts.

To avoid these pitfalls, it is crucial that widespread micro-storage be controlled so as to make aggregate demand more stable and predictable. Taking a centralised approach, however, would be impractical as it could involve coordinating millions of devices, each with its own individual constraints and consumer preferences. Hence, we turn to the multi-agent systems paradigm, which is a natural fit for managing such large systems in a decentralised fashion. Recent work on applying agents in the smart grid include (Ramchurn et al. 2011b; Kok and Venekamp 2010) where the intention is to locally automate energy management tasks in the home via a *smart meter*. Smart meters are intended to allow suppliers to access detailed energy consumption data and, more importantly, provide network information, such as real-time pricing (RTP) signals, to consumers in an attempt to better control or reduce demand when electricity is expensive or carbon intensive on the grid (Hammerstrom et al. 2008; Smith 2010). Accordingly, we envisage that micro-storage will be controlled by autonomous software agents that will react to RTP signals to minimise their owner’s costs (i.e., they are self-interested). In this vein, we note our recent work (Vytelingum et al. 2010) in which we showed that, when acting purely selfishly, large numbers of micro-storage agents can cause instability in the aggregate demand profile. This is undesirable, as failing to accurately predict consumers needs can be very costly for energy suppliers. Hence, we proposed a stable adaptive learning mechanism for micro-storage agents. However, we only considered agents purchasing energy (and not re-selling) and had no mechanism to ensure participation. In this paper, we take a complimentary approach, exploiting the role an energy supplier can play as an intermediary between consumer and market. The stronger buying power of an energy supplier allows us to use a more detailed and complex model of the wholesale market, and the potential complexity of interaction between supplier and consumer allows us to consider suppliers purchasing energy from consumers and providing explicit incentives for cooperative behaviour. The development of realistic models and efficient micro-storage management mechanisms for suppliers and consumers is essential if the technology is to be widely adopted. This is the challenge we address in this paper.

In more detail, we use agent-based optimisation and control theoretic approaches to design a novel method for an energy supplier to profitably manage widespread home micro-storage in a decentralised fashion, without having full information on the number or capabilities of the storage devices present. Our approach involves using pricing signals that are broadcast to consumers in advance of each daily period, and allowing micro-storage agents to buy and sell electricity at the same price at any given time interval. We argue this makes agent behaviour more predictable as it removes the need for agents to speculate on prices or their owners' load profiles. We also introduce a novel method of charging consumers for *changing* their storage profile from day to day. This incentivises micro-storage agents to adapt to prices slowly, thus improving system stability. Using these key insights, we make the following novel contributions:

1. We propose a novel general adaptive pricing scheme that can be used by suppliers to manage aggregate consumer demand profiles in a decentralised fashion. We prove that the interaction between a supplier using our scheme and the actions of individual selfish micro-storage agents forms a stable feedback loop, under which the aggregate demand profile converges to a unique equilibrium.

2. We provide a specific example of our pricing scheme, with pricing functions that are designed to recover supplier costs. In simulated experiments of realistic scenarios, we empirically show that this pricing scheme stabilises the aggregate micro-storage profile, is robust to shocks (sudden increases or reduction in micro-storage), and provides sufficient consumer revenue to guarantee a profit.

3. We propose a market strategy that allows the supplier to reduce wholesale purchasing costs without increasing uncertainty and variance for its aggregate consumer demand. In simulated experiments, this strategy is shown to reduce suppliers costs over time, while still benefiting the consumer. In particular, using data taken from the UK grid, our approach is show to yield savings of up to 16% in energy cost for consumers using storage devices with average capacity 10 kWh. Moreover, we show that our algorithm is robust against extreme system changes. When taken together our results constitute the first key benchmarks for energy supply management for home micro-storage in the smart grid.

The rest of this paper is structured as follows. In Section 2 we discuss our model of the smart grid agents, supplier, and electricity markets. We then present our novel algorithm in Section 3. Section 4 presents our theoretical results, and gives conditions for the stability of our mechanism. Building on this, in Section 5 we empirically test our algorithm's performance using a simulation of the UK electricity market. Finally, Section 6 concludes.

2 Agents and Supplier Models

In our model, we consider fixed time intervals consisting of single days, each separated into $T = 48$ settlement periods of half an hour. The time of the start of the day is taken to be the beginning of the off-peak hours at night, when electricity is cheapest, so consumers will aim to use all stored energy by time T. We now describe our model of home micro-storage control agents.

Home Storage Control Agents

We consider a set of consumers $\mathcal{A}$ which we define as selfish agents that always minimise their individual costs. Each agent $a \in \mathcal{A}$ has a load profile $l_i^a \ \forall i \in \mathcal{I} = \{1, \ldots, T\}$, such that l_i^a is the amount of electricity used by the owner of agent a during time interval i. The aggregate load profile of the system is given by $l_i = \sum_{a\in\mathcal{A}} l_i^a$. We consider this aggregate load profile to be fixed. Although there are seasonal variations in demand in practice, there is a high degree of consistency from day to day with some variation between weekday and weekend profiles. With that in mind, an energy supplier could operate a concurrent implementation of our pricing scheme for each day of the week, with the daily learning in our algorithm actually occurring between days that are a week apart. Each agent $a \in \mathcal{A}$ may also have some storage available to it, with capacity e^a, efficiency α^a and running costs c^a, such that if q amount of energy is stored, then $\alpha^a q$ may be discharged and the storage cost is $c^a q$. Here, the cost c^a may represent a fixed capital investment divided over the charging cycle lifetime of the storage device.

In order to minimise costs, a can attempt to strategise over its storage profile. A storage profile is a vector of values that represent the amount of energy charged and discharged during individual time intervals. We use $b_i^{a+} \geq 0, b_i^{a-} \geq 0 \ \forall i \in \mathcal{I}$ to denote the storage profile of agent $a \in \mathcal{A}$, where, for each $i \in \mathcal{I}$, b_i^{a+} and b_i^{a-} represent the precise amount of energy charged and discharged during i, respectively. The device cannot charge and discharge at the same time, and must always have between 0 and e^a stored. Thus, a storage profile is feasible if and only if, for all $i \in \mathcal{I}$, $(b_i^{a+}/b_+^a) + (b_i^{a-}/b_-^a) \leq 1$, where b_+^a and b_-^a are the maximum charge and discharge volumes for one time interval, and $0 \leq \sum_{j=1}^{i} \alpha^a b_j^{a+} - b_j^{a-} \leq \alpha^a e^a$, with equality on the left relation at $i = T$. We let $\mathcal{B}^a$ represent the set of valid storage profiles for a, and set $\mathcal{B} = \times_{a\in\mathcal{A}} \mathcal{B}^a$. Here $\times$ denotes the Cartesian product of vector spaces. Since the inequalities that define $\mathcal{B}$ are linear and not strict, and it must be closed and convex. For any $b \in \mathcal{B}$, for all $i \in \mathcal{I}$ we let $b_i^a = b_i^{a+} - b_i^{a-}$ for each $a \in \mathcal{A}$, and we let $b_i = \sum_{a\in\mathcal{A}} b_i^a$.

The Supplier

We consider a single supplier providing energy for its consumers. To this end, as in most electricity markets (e.g., UK or US), the supplier buys electricity from generators either

Baseload	23:00 - 23:00	$p^o \sim 47.3$
Peak	07:00 - 19:00	$p^o \sim 54.01$
Extended peak	07:00 - 23:00	$p^o \sim 51.69$
Off peak	23:00 - 07:00 & 19:00 - 23:00	$p^o \sim 40.58$
4 Hrs block	6 blocks per day from 23:00 to 23:00	$p^o \in [22, 103]$
2 Hrs block	12 blocks per day from 23:00 to 23:00	$p^o \in [21, 139]$
Half hour block	48 blocks per days from 00:00 to 24:00	$p^o \in [17, 157]$

Table 1: Forward market contracts span a specific period of time and charge different prices. The range of prices shown were obtained from APX-ENDEX for January 2010. As can be seen, prices are generally higher for shorter time periods.

directly or through wholesale electricity markets. We focus on the two fundamental types of markets which are the day-ahead *forward market*[1] (where prices are known on a day-ahead basis and where most of the trades occur) and the *balancing market*[2] (where prices are known *a posteriori*).

The forward market runs on a day-ahead basis whereby the supplier can purchase amounts of energy $f_i^o \in \mathbb{R}^+$ from different types of electricity contracts, $\Theta = \{o_1, \cdots, o_{|\Theta|}\}$ defined over the next 24 hours (see Table 2) and valid only at certain times during the day. The supplier purchases a quantity $f_i = \sum_{o\in\Theta} f_i^o \; \forall i \in \mathcal{I}$ at price $C_i^f = \sum_{o\in\Theta} f_i^o p^o \; \forall i \in \mathcal{I}$. We also define $\mathcal{F}$ as the set of all feasible contract purchases, $f_i^o \; \forall i \in \mathcal{I}, \forall o \in \Theta$ (i.e., amounts that the supplier is able to purchase) and $p_f^i = C_i^f / f_i$ as the average contract price. The balancing market then settles any differences in committed supply and actual supply in real-time (these differences occur when consumers or generators behave unexpectedly due to weather effects, outages, or other factors). Thus, any excess purchased, $(f_i - d_i)^+$ (where $d_i = l_i + b_i$), is sold at the balancing sell price p^{sell} and excess electricity used, $(d_i - f_i)^+$, is purchased at the balancing buy price p^{buy}. Note that the balancing buy and sell prices are typically higher and lower respectively compared to forward market prices and hence, it is crucial that the supplier ensures that it can cater for most of its customers demand from forward market contracts rather than leaving it to the balancing market. We next elaborate on our proposed supplier strategy, in particular how to set prices for customers in order to stabilise storage (and, as a result, stabilise aggregate demand), and minimise its costs on the electricity market.

3 Supplier Strategy

The supplier has complex objectives. On the one hand, it needs to manage the behaviour of its customers using only pricing signals, without fully knowing their storage capabilities. On the other hand, it needs to get the best price it can on the electricity market given the high penalties in the balancing market associated with under or over supply. In this section, we propose a novel mechanism for a supplier to stabilise demand while maximising its revenue by optimising its purchases in the forward market. Crucially, in Section 4, we go on to prove that our mechanism is converges to a stable equilibrium.

Pricing Mechanism

We consider the situation where the electricity supplier sets a price for energy p_i for each $i \in \mathcal{I}$. When micro-storage agents react to such pricing signals, the aggregate behaviour can be unstable, as their charging and discharging activities tend to concentrate in those time intervals with extremal prices. This results in agents adopting "all or nothing" behaviour, where they either charge or discharge at maximum rate or do nothing at all. In order to combat this volatility, we propose the following. First, the supplier passes the real-time prices *at the beginning of the day* (instead of communicating them every half-hour as in (Hammerstrom et al. 2008) or at the end of the day as in (Vytelingum et al. 2010)). This allows agents to explicitly optimise their storage profiles without needing to speculate on future price changes. Second, under our mechanism, agents are allowed to both buy and sell electricity *at the same price* p_i for each $i \in \mathcal{I}$. Thus, discharging a quantity of stored electricity generates the same profit to the consumer when sold to the supplier as when used by the consumer, and so the optimal storage profile is completely independent of load at time i. This avoids agents having to speculate against their uncertain future load profile, which can vary greatly from day to day and would be hard for agents to reliably predict. Hence, if all agents store optimally according to prices given, their aggregate behaviour becomes more predictable. Third, as the central part of our control mechanism, we propose that, in order to stabilise the system, agents be charged an additional fee based on *how greatly they change their storage profile*. That is, each agent $a \in \mathcal{A}$ pays a daily fee of $\sum_{i\in\mathcal{I}} \kappa\big(b_i^a - \tilde{b}_i^a\big)^2$, where $\tilde{b}^a$ is the previous days storage profile and $\kappa > 0$ is a parameter given by the supplier. As we show in Theorem 1, if κ is large enough, this introduces enough damping into the system to guarantee stability.

Since the behaviour of the agent will be to maximise the profit of its user, we can now predict that each $a \in \mathcal{A}$ will adopt the following optimal storage profile,

$$b^a = \arg\min_{b^a \in \mathcal{B}^a} \sum_{i\in\mathcal{I}} p_i b_i^a + c^a b_i^{a+} + \kappa\big(b_i^a - \tilde{b}_i^a\big)^2 \qquad (1)$$

This characterisation of agent behaviour is sufficient for us to model the control loop between supplier and consumer agents and come up with a pricing strategy that allows aggregate micro-storage management without requiring the supplier to know fully the number or parameters of the storage devices belonging to its consumers.

From day to day, we propose that the electricity supplier updates its prices according to the current loads, following some strictly increasing, differentiable pricing function $p_i(\cdot)$, with $|p_i'(\cdot)| < K$ for some K. At the end of each day the new price profile for the next p^{new} is set as $p_i(\cdot)$ of the

[1] We use real UK market data from APX-ENDEX energy exchange (`http://www.apxendex.com`) for the period from January to March 2010. Note that our model can be extended to consider contracts over longer terms, such as weeks, seasons or years. However, this is beyond the scope of this work.

[2] We use real data obtained from the UK balancing market (`http://www.bmreports.com`).

previous days total demand:

$$p_i^{\text{new}} = p_i(l_i + b_i). \tag{2}$$

As we will see later in Section 4, Theorem 1 shows that if $\kappa > K|\mathcal{A}|/3$, then (1) and (2) form a stable system. This means the supplier only needs to know an upper bound on the number of storage devices in operation in order to safely apply this mechanism.

If the supplier has particular targets for the equilibrium storage profile (e.g., it has information about its customers' storage capacity that would allow it to perform arbitrage in the market or reduce its carbon emissions), it can adjust its price functions to directly steer the consumers. For example, by making the price functions steeper, the supplier can encourage storage use, and by making them less steep, the supplier can motivate consumers to use storage less. Otherwise, the supplier can choose price function to reflect its market costs.

In our experiments (see Section 5) we sought to use pricing functions that reduce forward market costs, and ensure sufficient customer revenue to make a profit. We chose the following price functions (which satisfy the criterion that the price function is strictly increasing and differentiable), for each $i \in \mathcal{I}$,

$$p_i(x) = \begin{cases} p_i^f \frac{f_i}{x} + (p_i^{\text{buy}} + 2\Delta p^f)\frac{x-f_i}{x} & , x > f_i \\ p_i^f - \frac{\Delta p^f}{x}(f_i - x) & , f_i/2 \le x \le f_i \\ p_i^{\text{sell}} & , x < f_i/2 \end{cases} \tag{3}$$

where we set Δp^f to be the maximum of $p_i^f - p_i^{\text{sell}}$. In this case a suitable value for K would be:

$$K = \max_{i \in \mathcal{I}} \max\Big(\frac{4\Delta p^f}{f_i}, \frac{1}{f_i}(p_i^f + p_i^{\text{buy}} + 2\Delta p^f)\Big).$$

With these choices, Proposition 1, in Section 4, shows that under (1) – (3), when storage profiles converge, the supplier is guaranteed a profit. We achieve this in part by ensuring that for each $i \in \mathcal{I}$, it sets $p_i(f_i) = p_i^f$, where p_i^f is the average forward contract price for time interval i. This has the added benefit that since the $p_i(\cdot)$ functions are strictly increasing, if the agents react to the price signals by changing their storage allocation, they will do so in a way that would be profitable to the supplier if they were charged according to p_i^f. Hence, if the aggregate demand deviates from f_i, the supplier can and will reduce costs by re-optimising their forward contracts (as we show later in Section 5). We next discuss how the supplier can optimise its contracts given realistic uncertainty in its customers' demand.

Market Strategy

Since the aggregate load profile gradually changes as per our mechanism (as we show later, it eventually reaches an equilibrium), we can expect that it will deviate from the forward contract amounts on a daily basis and the supplier may have to trade on the balancing market to cater for under/over usage. Since the balancing market is more expensive than the forward market, the supplier can reduce costs by re-optimising its forward contracts daily. Note, if using the pricing functions (3), these must be re-computed every time the forward market contracts are changed. In what follows, we present an optimisation model that allows a supplier to optimise its contracts against uncertain demand.

To decide on the contracts to acquire, the supplier simply computes the optimal quantity for each contract[3] $f_i^o, \forall o \in \Theta, \forall i \in \mathcal{I}$ that minimises its costs. However, because the consumers' aggregate demand d_i as well as balancing buy and sell prices are not known *a priori*, the supplier needs to compute the amounts for each contract that will minimise its expected costs subject to uncertainties in prices and demand (see Figures 1 and 2 respectively).

We address this problem by modelling the aggregate load and prices using a standard ARMA (auto-regressive moving-average) statistical model[4] that can deal with the seasonal trends of loads and prices (Weron 2006). In so doing, we can build a distribution of the following day's balancing prices and the aggregate load. The supplier's problem is now to minimise its cost based on distributions of the demand and buy and sell prices. To deal with the optimisation under uncertainty, we use Monte Carlo simulations, drawing a large number of samples $\mathcal{S}$ from the distributions and expanding the objective function to factor in the samples as follows:

$$\arg\min_{f \in \mathcal{F}} \sum_{s \in \mathcal{S}} \sum_{i \in \mathcal{I}} \Big(C_i^f + p_{i,s}^{\text{buy}}(d_{i,s} - f_i)^+ + p_{i,s}^{\text{sell}}(d_{i,s} - f_i)^-\Big) \tag{4}$$

where $f_i = \sum_{o \in \Theta} f_i^o$ for amounts f_i^o selected. The above optimisation returns a profile of optimal quantity $\hat{f}_i^o$ to purchase from each contract at time i which are, in turn, used to formulate our pricing functions as in the previous section. It is important to note that this method does not rely on detailed information about the consumer storage devices present, but instead *adapts* to changes in the aggregate demand profile in order to reduce costs. We next present the key analytical properties of our proposed mechanism.

4 Theoretical Results

In this section we give the theoretical results which support our proposed algorithm. We begin by showing that (1) – (2) is stable. Our proof uses a novel application of a Lyapunov function (Slotine and Weiping 1991), where the function is chosen so that it is always less than or equal to total agent costs with equality if agents do not change their storage profile. As agents attempt to reduce costs, they cause the Lyapunov function to decrease monotonically, and the system converges to its minimising aggregate profile. This

[3]The supplier can purchase any set of contracts to cover any quantity. For example, a baseload contract will have a lower price (as it can be provided by cheap nuclear generation) than the average daily balancing buy prices. Generally, contract prices tend to be marginally lower that the average balancing prices over the same period. However, forward contracts reduce market risk as forward prices are known *a priori* and are thus desirable.

[4]We chose a standard model, as forecasting of load and prices is beyond the scope of this paper, and not central to the supplier's adaptive mechanism (Weron 2006). More complex quantitative models (e.g., GARCH or Jump-Diffusion models) can be used for marginal improvements.

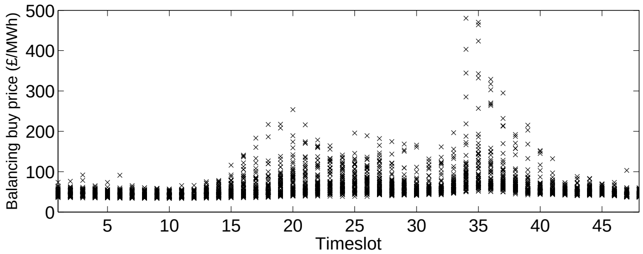

Figure 1: Distribution of balancing buy prices in the UK (October 2009 to March 2010).

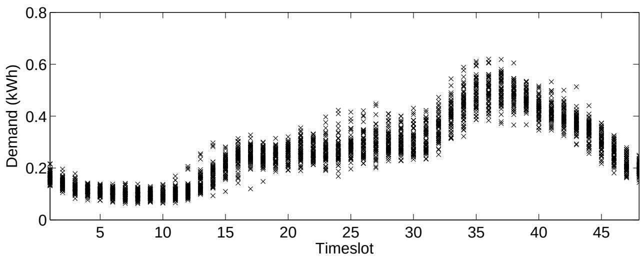

Figure 2: Load distribution (with seasonal trend removed) in the UK (January to March 2010).

method of proof allows the supplier to manage agent responses without having any information about the capabilities of the micro-storage devices they control. All that is required is an upper bound on the number of active micro-storage devices in the system.

Theorem 1. *Given a fixed set of forward contracts, the system* (1) – (2) *is stable, in that the vector* b *converges to an equilibrium, provided:*

$$\kappa \geq \frac{KN}{3},$$

where N *is the number of agents with active storage devices and* K *is an upper bound on* $|p_i'(\cdot)|$ *for all* $i \in \mathcal{I}$ *over the relevant range of demand quantities. In this case, the aggregate storage profile at equilibrium is unique.*

Proof. Let us define the Lyapunov function $V(\cdot)$ to be,

$$V(b) = \sum_{i \in \mathcal{I}} \int_{f_i}^{l_i + b_i} p_i(l_i + b_i) + \sum_{a \in \mathcal{A}} c^a b_i^{a+}.$$

Since the $p_i(\cdot)$ are increasing, $V(\cdot)$ is convex, with a convex region of optimality over $\mathcal{B}$, and no other local minima. Furthermore, as the $p_i(\cdot)$ are strictly increasing, the aggregate storage profile must be the same for all optimal profiles. For any two storage profiles b, b', let us also define $V_{b'}(\cdot)$,

$$V_{b'}(b) = \sum_{i \in \mathcal{I}} p_i(l_i + b_i')(l_i + b_i) + \sum_{a \in \mathcal{A}} c^a b_i^{a+}.$$

Now let us consider a particular day. Suppose the previous days storage profile is given by $b \in \mathcal{B}$ and the new storage profile is $b + \Delta b$. From (1) – (2) we can deduce that Δb is chosen so that $b + \Delta b$ minimises $W_b(b + \Delta b)$ over $\mathcal{B}$ where we define $W_b(\cdot)$ to be:

$$W_b(b+\Delta b) = V_b(b+\Delta b) + \kappa \sum_{i \in \mathcal{I}} \sum_{a \in \mathcal{A}} \left(\Delta b_i^{a+} + -\Delta b_i^{a-}\right)^2.$$

Now, if we use ∇ to denote gradient, then, by definition, $\nabla V_b(b) = \nabla V(b)$, and so, $\nabla W_b(b) = \nabla V(b)$. So $W_b(b + \Delta b) - W_b(b) \leq 0$, with equality only if $\nabla V(b) \cdot (b' - b) \geq 0$ for all feasible $b' \in \mathcal{B}$.

We calculate that:

$$V(b + \Delta b) - V(b) = \int_0^1 t \Delta b \cdot \nabla V(b + t\Delta b) dt,$$

and

$$V_b(b + \Delta b) - V_b(b) = \int_0^1 t \Delta b \cdot \nabla V(b) dt.$$

The difference between these two is:

$$\int_0^1 t \Delta b \cdot (\nabla V(b) - \nabla V(b + t \Delta b)) dt.$$

However, since the price functions $p_i(\cdot)$ for $i \in \mathcal{I}$ are the only non-linear component of $V(\cdot)$, we have that the rate of change of each coordinate of $\nabla V(\cdot)$ is bounded by K. Thus, $V(b + \Delta b) - V(b)$ is less than or equal to:

$$V_b(b + \Delta b) - V_b(b) + \int_0^1 t^2 K \sum_{i \in \mathcal{I}} (\Delta b_i)^2 dt,$$

$$= V_b(b + \Delta b) - V_b(b) + \frac{K}{3} \sum_{i \in \mathcal{I}} (\Delta b_i)^2,$$

$$\leq V_b(b + \Delta b) - V_b(b) + \frac{KN}{3} \sum_{i \in \mathcal{I}} \sum_{a \in \mathcal{A}} \left(\Delta b_i^{a+} - \Delta b_i^{a-}\right)^2$$

$$\leq W_b(b + \Delta b) - W_b(b) \leq 0,$$

by choice of κ. Thus, for each day after the forward market strategy is computed, $V(b)$ is non-increasing. Since it has a global minimum, the size of these decreases must tend to zero. However, from the above, this can only happen if $\nabla V(b) \cdot (b' - b)$ tends to a value greater than or equal to 0 for all $b' \in \mathcal{B}$. By the continuity of $\nabla V(b)$, we must have that b converges to the optimal region for $V(b)$ within $\mathcal{B}$. □

Note, although N will be potentially very large, K should be very small, and should be roughly $O(1/|\mathcal{A}|)$, since K depends on how much prices vary when a single user varies their storage profile. In our experiments we use $\kappa = |\mathcal{A}|K/3$ and calculate K as given in Section 3.

We can also show that the prices we give in (3) are suited to recovering the costs of the energy supplier.

Proposition 1. *Provided the total daily electricity bought in the forward market is at most the total daily load then if* b *is sufficiently close to the previous days storage profile, then the supplier recovers their costs from consumers. More precisely,*

$$\sum_{i \in \mathcal{I}} p_i(l_i + b_i)(l_i + b_i) \geq C_i(l_i + b_i).$$

Proof. For all $i \in \mathcal{I}$ let $d_i = l_i + b_i$. If $d_i < f_i$ then, from (3),

$$p_i(d_i)d_i \geq C_i^f - (f_i - d_i)(p_i^f + \Delta p^f),$$

and if $d_i \geq f_i$,

$$p_i(d_i)d_i \geq C_i^f + (d_i - f_i)(p^{\text{buy}} + 2\Delta p^f).$$

Thus, if the total market costs are

$$C_i^f + p_i^{\text{buy}}(d_i - f_i)^+ + p_i^{\text{sell}}(d_i - f_i)^-,$$

then the total revenue minus total costs is at least:

$$2d^+\Delta p^f - d^-\Delta p^f - \sum_{i \in \mathcal{I}}(f_i - d_i)^+(p_i^f - p_i^{\text{sell}}),$$

where d^+ is the sum of $(d_i - f_i)^+$ and d^- is the sum of $(f_i - d_i)^+$. However, since the total sum of f_i is at most the total sum of l_i, and all storage devices have efficiency less than or equal to 1, the amount discharged to cause a surplus must be less than the amount charged. Therefore, we must have $d^+ \leq d^-$, and so the total revenue is at least as big as the total costs to the supplier. □

This means that, once (1) – (3) has converged, the energy supplier is guaranteed to make a profit each day.

5 Experiments

In this section, we evaluate the performance of our decentralised control mechanism. Specifically, we set up a simulation based on 1000 consumers, each with a different battery,[5] using real load data from January to March 2010 (from a set of anonymised UK consumers – see Figure 2), over 100 runs.[6] We used real market data over the same period with an additional historical data of 3 months of the balancing market prices (October 2009 to March 2010 – see Figure 1) to calibrate our model of price distribution.

As discussed in Section 1, an important objective of the pricing mechanism is to ensure stable aggregate consumer behaviours. To this end, we analyse the stability of the system over a number of days to ascertain the effectiveness of the mechanism. From Figure 3, we can observe that the root mean squared of the difference between the forward-contracted demand and the aggregate consumer demand decreases from 31.5 (when there is no storage) and converges to an equilibrium around 14.1. As it does so, the expected wholesale cost of the system also converges to an equilibrium. Figure 3 also shows how the expected wholesale cost gradually decreases (by 16% from 33.47p to 28.22p) and converges after a number of trading days. The consumer cost (ignoring the supplier's profit and operational margin) decreases by 9.28p (from 42.42p to 33.14p). Note that a supplier typically inflates the retail price to include a profit and an operational margin which the consumers pays on top of the consumers' cost (or retail revenue) we describe here. Thus, given the 16% decrease in its wholesale cost and based on its margin, the supplier can markup on the, now reduced, consumer's cost or can simply reduce the retail price to incentivise storage. Furthermore, we observe from Figure 4 that the volatility (measured as the normalised standard deviation of prices) of the expected wholesale cost is noticeably lower (decreasing from 6.6% to 5.3%). Lower volatility is be desirable for the supplier, especially with the high volatility of prices in electricity markets.

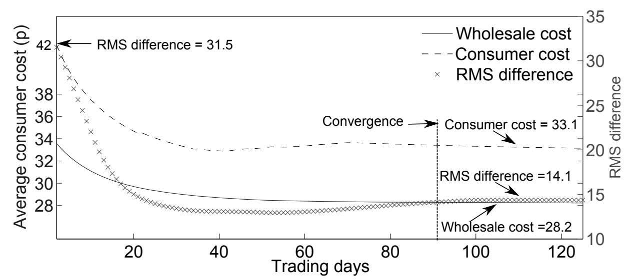

Figure 3: Expected daily wholesale cost of electricity per consumer (starting at 33.19p with no storage) and consumer cost (without supplier's profit and operational margin).

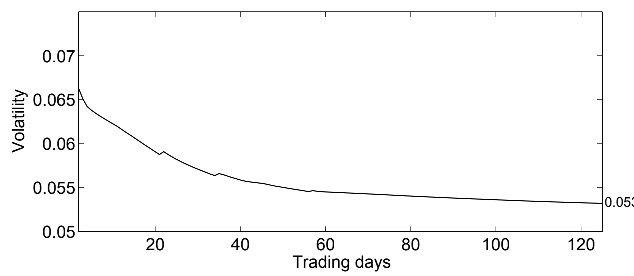

Figure 4: Volatility of expected wholesale cost.

Finally, we analyse the stability of our mechanism when faced with system shocks (i.e., a sudden and drastic change in micro-storage demand response capability). Specifically, we simulate 50% of the batteries failing (e.g., as a result of a glitch in batteries or network outage) on Day 100 and all *repaired* and reinstated on Day 200. Figure 5 shows the effect of such shocks on the system by looking at how the root mean squared difference between total daily demand and forward-contracted demand changes. In particular, we observe that from Day 100, there is a considerable difference between the contracted demand and the actual demand (as a result of the micro-storage failure). However, the system quickly re-adapts, converging to a new equilibrium. On the other hand, when all the failed batteries are repaired on Day 200, the agents with newly repaired batteries *gradually* build up their storage profile due to dampening by our pricing mechanism, while the supplier re-optimises its forward contracts subject to the increasing demand. The rate of increase of the storage is sufficiently low that the supplier can optimise its forward contracts and ensure the difference

[5] For each battery, its capacity is defined by a normal distribution $\mathcal{N}(10, 3)$, its efficiency by $\min(1.0, \mathcal{N}(0.8, 0.2))$, its charging capacity by $N(0.8, 0.4)$ and, finally, its discharging capacity by $\mathcal{N}(0.8, 0.4)$. These values were based on typical batteries.

[6] Using a t-test, we validated our results with $\alpha = 0.05$ to ensure that they were significant.

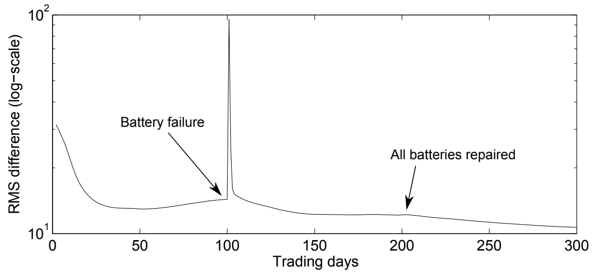

Figure 5: Root mean squared difference between contracted and actual demand, with a shock on Day 100 when 50% of the batteries fail and on Day 200 when they are repaired.

between the contracted demand and actual demand remains small. The effective response of our mechanism to extreme changes implies that our mechanism is indeed robust against system changes.

6 Conclusion

In this paper we proposed a novel algorithm for the decentralised control of widespread micro-storage in the smart grid. We proved theoretical results that showed the stability and profitability of the algorithm and then conducted experimental simulations to verify those results. We empirically showed that, in a realistic scenario our mechanism reduced consumer costs by 16%, and further, we showed that it is stable against dramatic short term changes in the system.

We see this as an important step to showing that the adoption of widespread, supplier managed home energy micro-storage is a practical, desirable technology to develop. Using the techniques described in this paper, we can envisage energy suppliers utilising large numbers of affordable small scale storage devices in order to manage aggregate load profiles, improve efficiency and reduce carbon output. Future work should involve integrating these models and simulations into further important elements of the future smart grid. This could include modelling the interaction between micro-storage and intermittent or unreliable renewable generation, and analysing the impact of vehicle to grid schemes on aggregate load profiles.

Acknowledgements

We thank the reviewers for their valuable comments. This work was carried out as part of the industrially funded (by an energy company) IDEAS project.

References

Bathurst, G. N., and Strbac, G. 2003. Value of combining energy storage and wind in short-term energy and balancing markets. *Electric Power Systems Research* 67(1):1–8.

den Bossche, P. V.; Vergels, F.; Mierlo, J. V.; Matheys, J.; and Autenboer, W. V. 2006. Subat: An assessment of sustainable battery technology. *Journal of Power Sources* 162(2):913–919.

Galvin, R., and Yeager, K. 2008. *Perfect Power: How the MicroGrid Revolution Will Unleash Cleaner, Greener, More Abundant Energy*. McGraw-Hill Professional.

Hammerstrom, D.; Ambrosio, R.; Carlon, T.; DeSteese, J.; Horst, G.; Kajfasz, R.; Kiesling, L.; Michie, P.; Pratt, R.; Yao, M.; and et al. 2008. Pacific northwest gridwise testbed demonstration projects; part i. olympic peninsula project. PNNL-17167.

Kok, K., and Venekamp, G. 2010. Market-based control in decentralized power systems. In *Proceedings of the First International Workshop on Agent Technology for Energy Systems*, 61–66.

Ramchurn, S. D.; Vytelingum, P.; Rogers, A.; and Jennings, N. R. 2011a. Agent-based homeostatic control for green energy in the smart grid. *ACM Transactions on Intelligent Systems and Technology* (in press).

Ramchurn, S.; Vytelingum, P.; Rogers, A.; and Jennings, N. R. 2011b. Agent-based control for decentralised demand side management in the smart grid. In *Proc. of the Tenth International Conference on Autonomous Agents And Multi-Agent Systems – to appear*.

Slotine, J. E., and Weiping, L. 1991. *Applied Nonlinear Control*. Prentice Hall.

Smith, K. 2010. Energy demand research project - fifth progress report, OFGEM UK.

UK Department of Energy and Climate Change. 2009. Smarter grids: The opportunity.

US Department Of Energy. 2003. Grid 2030: A National Vision For Electricity's Second 100 Years.

Vytelingum, P.; Voice, T. D.; Ramchurn, S. D.; Rogers, A.; and Jennings, N. R. 2010. Agent-based micro-storage management for the smart grid. In *Proc. of the Ninth International Conference on Autonomous Agents and Multiagent Systems*, 39–46.

Weron, R. 2006. *Modelling and forecasting electricity loads and prices*. Wiley Finance.