

Agent-based Micro-Storage Management for the Smart Grid

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

ABSTRACT

The use of energy storage devices in homes has been advocated as one of the main ways of saving energy and reducing the reliance on fossil fuels in the future Smart Grid. However, if micro-storage devices are all charged at the same time using power from the electricity grid, it means a higher demand and, hence, requires more generation capacity, results in more carbon emissions, and, in the worst case, breaks down the system due to over-demand. To alleviate such issues, in this paper, we present a novel agent-based micro-storage management technique that allows all (individually-owned) storage devices in the system to converge to profitable, efficient behaviour. Specifically, we provide a general framework within which to analyse the Nash equilibrium of an electricity grid and devise new agent-based storage learning strategies that adapt to market conditions. Taken altogether, our solution shows that, specifically, in the UK electricity market, it is possible to achieve savings of up to 13% on average for a consumer on his electricity bill with a storage device of 4 kWh. Moreover, we show that there exists an equilibrium where only 38% of UK households would own storage devices and where social welfare would be also maximised (with an overall annual savings of nearly GBP 1.5B at that equilibrium).

Categories and Subject Descriptors

I.2.11 [Distributed Artificial Intelligence]: Multi-agent Systems

General Terms

Algorithms, Management, Economics

Keywords

Agent-based simulation, Smart Grid, Energy, Micro-storage

1. INTRODUCTION

Energy storage is one of the key underpinnings of the vision of the *Smart Grid* which aims to support sustainable energy provisioning across the world [2, 4, 8]. Given this,

Cite as: Agent-based Micro-Storage Management for the Smart Grid, Vytelingum *et al.*, *Proc. of 9th Int. Conf. on Autonomous Agents and Multiagent Systems (AAMAS 2010)*, van der Hoek, Kaminka, Lespérance, Luck and Sen (eds.), May, 10–14, 2010, Toronto, Canada, pp. 39–46
Copyright © 2010, International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). All rights reserved.

research has been focused on designing new efficient low cost storage devices that would be able to efficiently store electricity for long periods of time and allow a sufficient number of charging/discharging cycles without significant degradation in performance [8].¹ By using such devices, it is expected that energy usage can be improved in a number of ways. If storage devices can be used to supply home devices at peak electricity consumption times (typically in the morning and evening), then it should be possible to lower peak demand such that fewer carbon-intensive and expensive “peaking plant” generators are required, thus reducing both energy costs and carbon emissions. Furthermore, storage devices can be used to compensate for the variability of typical renewable electricity generation (e.g., wind, wave, solar), thus making the integration of such generation facilities into the existing grid more viable in practice [8]. Such energy storage may even take the form of electric vehicles (EVs) or plug-in hybrid electric vehicles (PHEVs); this vision is sometimes referred to as vehicle to grid (V2G).

There are, however, a number of potential challenges in this setting. For example, consider individual homes (among the 26M UK households) storing electricity according to their own needs and all deciding to charge their batteries at the same time (e.g. incentivised by cheaper prices). Now, not only would this cause a higher peak in demand in the electricity market, and hence higher carbon emissions and more costly electricity, but, in the worst case, it could cause blackouts and infrastructure damage if this demand were to exceed network capacity. Moreover, if individuals were only charging their batteries according to the amount they use, they may be paying for electricity at a higher price than if they did not have the device when the cost of the battery is added to the mix. Finally, if most homes in the system start using storage and manage to shave off peak demand, electricity prices may become lower than the price of storing electricity.

To address such issues, the multi-agent systems paradigm has been advocated as both a solution and a framework to analyse the properties of systems in which multiple self-interested parties interact [3, 6, 9]. In particular, with the advent of *smart meters* that can monitor and control devices in the home, it is now possible to envisage that smart software agents could be installed on these devices. These agents would then be able to optimise the usage and storage profile of the house using information from various sources (e.g., weather data to predict energy and hence heating costs or price plan data from suppliers). Now, most of the existing approaches to applying intelligent agents typically study

¹See batteries recently developed by Ceramtec: <http://www.ceramtec.com>

how individual homes could optimise the way they store energy or how storage devices could coordinate with renewable energy generation facilities to maximise energy used from such sources (see Section 2 for more details). However, their approach ignores the individual preferences of each home and does not exactly model the real impact of agents learning to adapt to the constraints that they themselves impose on the system. Thus, an approach that focuses on the system dynamics where all agents in the system are given the freedom to buy electricity whenever they see fit, would address these issues.

In this paper we address this shortcoming and provide a game-theoretic framework for modelling storage devices in large-scale systems where each storage device is owned by a self-interested agent that aims to maximise its monetary profits. Using this framework, under certain assumptions, we are able to predict the behaviour of the system given that each agent behaves rationally (i.e. always adopts a storage profile that minimises its costs) and only reacts to a price signal. Building on this, we then go on to devise intelligent agent-based storage strategies that can learn the best storage profile given the market prices that keep changing as a result of consumers using storage. In more detail, this work advances the state of the art in the following ways:

1. We provide a novel game-theoretic framework to study storage strategies that agents might adopt. Given the normal electricity usage profile of all users in the system, we are then able to compute the Nash equilibria which describe when agents are going to charge their batteries, use their stored electricity, or use electricity from the grid.
2. We provide new agent-based micro-storage strategies that are able to learn the best storage profile to adopt when agents in the system may not have exactly the same storage capacities or efficiencies. Our strategies are shown to converge to the same Nash equilibria as those predicted by our framework.
3. Given our agent-based learning strategies, we are able to show how agents could learn to buy the most profitable storage capacity and using evolutionary game theoretic analysis, we are able to predict the portion of the population that would actually acquire storage capacity to maximise their savings.

In short, this is the first attempt at modelling, predicting equilibria, and building intelligent strategies for the problem of electricity storage on a large scale.

The rest of this paper is structured as follows. In Section 2 we discuss related work in the area of electricity storage and electricity markets. Section 3 discusses the key features of the electricity markets and lays down the general assumptions upon which we build our framework. Section 4 presents our game-theoretic framework and shows how the Nash equilibria of the system can be computed. Building on this, Section 5 describes the dynamics of a market where agents are given the ability to learn their best storage profile and, and Section 6 empirically studies this system through simulations. Finally, Section 7 concludes.

2. BACKGROUND

Very little work exists on the application of agent-based techniques to storage management in electricity grids. Typically, electricity storage has mainly been a concern of the

energy suppliers using large chemical batteries to store energy from intermittent renewable energy sources (e.g., wind or solar) [7]. The effect of such large scale storage on commodity markets in general is a mature area of study (see [5, 11] for some state of the art energy/fuel market specific results). However, the electricity market is unusual in that it has large daily cycles of demand, (see Figure 1 for the average UK daily consumer load profile) which make electricity storage potentially profitable even on an individual household scale. With the advent of new types of batteries that charge up to 20kWh of energy a day (i.e., sufficient to power all the devices in a house for a day), it is now possible to envisage that micro-storage devices will be widely used. Indeed, the energy storage requirements of a typical home are well aligned with the storage required in a feasible EV or PHEV. Moreover, with the advent of smart meters, it will be possible to manage the storage and usage of electricity within a single home using software agents residing on such meters. Thus, decentralised autonomous agent-based approaches are strong candidates for managing energy storage in future electricity networks. In this context, we note the seminal work of Daryanian et al. [1] which illustrated how individual smart meters could optimise, through iterative algorithms, the storage profile of a house. Their approach was, however, limited to considering very basic battery properties and did not consider wider issues for the grid. In the same vein, more recently, [6] provide algorithms for agents to optimise storage using CHPs (Combined Heat and Power) but ignore how populations of such agents would impact on the grid. On the other hand, [3, 9] have studied the application of storage devices on a wider scale. They show that using demand-side management (i.e., *directly* controlling the storage profile of a number of homes) coupled with storage can increase savings made in the system.

If not properly managed, storage systems can be unprofitable [8], so in the setting we consider, it is important to know whether small scale storage can be individually beneficial, and what strategies maximise this profit. It is also important to understand the system-wide effects of such strategies, in particular, quantifying the limits on the usefulness of small scale storage from a social welfare point of view. These are the key open questions that are addressed by this paper.

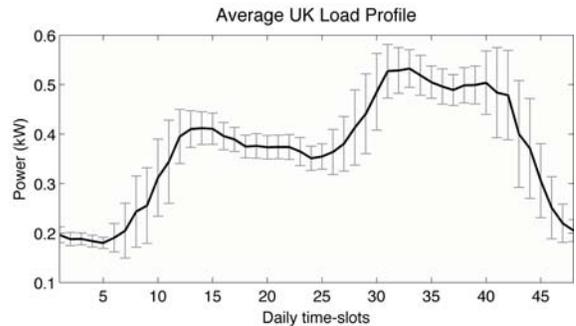


Figure 1: Representative Load Profile in UK (the Domestic Unconstrained profile).

3. MODEL DESCRIPTION

This section describes the models used in this paper. Our analysis considers fixed time interval consisting of single

days, each separated into $T = 48$ settlement periods of half an hour. Each day, agents consume electricity which is bought from suppliers through an electricity market. This market operates for each time interval in the day, so that variations in demand over time can be met. We proceed with a description of our models of behaviour for the agents, followed by a description of the market and a definition of the relevant social welfare metrics we consider.

3.1 Agents

We consider a set of consumers \mathcal{A} which we define as selfish agents that always aim to minimise their individual costs. Each agent $a \in \mathcal{A}$ has a load profile $\ell_i^a \forall i \in \mathcal{I} = \{1, \dots, T\}$, such that ℓ_i^a is the amount of electricity required by agent a for time interval i during each day. The aggregate load profile of the system is given by $\sum_{a \in \mathcal{A}} \ell_i^a = d_i$. We consider this load profile to be fixed over different days (although there are seasonal variations in demand in practice, there is a high degree of consistency from day to day). 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 . Here, the cost c^a may represent ongoing storage costs (for example, some battery devices expend energy through heating while they are in use) or may incorporate a fixed capital investment by a , to be paid off over time. The storage efficiency α^a and cost c^a are modelled to be such that if q amount of energy is stored, then $\alpha^a q$ may be discharged and the storage cost is $c^a q$.

In order to minimise costs, a can attempt to strategise over its storage profile, $b_i^a \forall i \in \mathcal{I}$ where $-b_-^a \leq b_i^a \leq b_+^a$, where b_-^a is the discharging capacity of the storage, and b_+^a , the charging capacity. For all $i \in \mathcal{I}$ we have $b_i^a = b_i^{a+} - b_i^{a-}$, where b_i^{a+} is the charging profile and b_i^{a-} , the discharging profile. Since we are attempting to model the effect of the widespread adoption of small scale household storage devices, we can assume that ℓ_i^a , b_+^a , and b_-^a are small in comparison to d_i . We denote the total storage capacity as $e = \sum_{a \in \mathcal{A}} e^a$, and the net storage profile as $b_i \forall i \in \mathcal{I}$ where $b_i = \sum_{a \in \mathcal{A}} b_i^a$. The net charging and discharging capacities are defined as $b_+ = \sum_{a \in \mathcal{A}} b_+^a$ and $b_- = \sum_{a \in \mathcal{A}} b_-^a$. To supply its load profile and energy charging needs each agent must purchase electricity from the available market. The next subsection contains our market model.

3.2 The Electricity Market

We consider a macro-model of the electricity market; a black box that abstracts the market mechanism and trading, as well as transmission power flow security involved in an actual electricity market mechanism. Given the characteristics of the market, our model gives us the market prices based on the economics of demand and supply (see Figure 2). The supply curve in this case is generated from UK National Grid prices for the period of August and September 2009.

The behaviour of electricity suppliers is specified by the supply curve $s_i(\cdot)$ for every time point $i \in \mathcal{I}$. The supply curve $s_i(\cdot)$ indicates the cost of electricity that generators experience, or minimum price they are willing to sell at, when producing a certain quantity. For our model we assume that $s_i(\cdot)$ is continuous and strictly increasing. As defined above, each time interval $i \in \mathcal{I}$ has an inelastic demand² quantity d_i , representing the total amount of electricity consumed by agents, and a net storage effect, b_i , representing the aggregated effect of storage. Thus, the total amount of electricity bought from suppliers at time interval $i \in \mathcal{I}$ is $q_i = d_i + b_i$.

²The market demand is modelled as inelastic to reflect the currently inelastic demand of individual consumers.

Under our model, for each time interval $i \in \mathcal{I}$, the market sets a price for electricity $p_i = s_i(q_i)$. Each agent pays $p_i(\ell_i^a + b_i^a)$ and the total cost for all agents is $p_i q_i$.

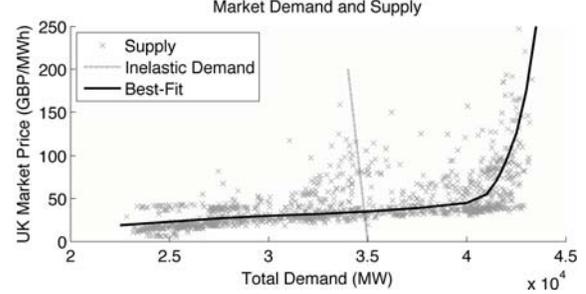


Figure 2: Supply is modelled from actual UK market prices and demand is assumed to be inelastic.

3.3 Social Welfare Metrics

A key aim of this paper is to study the effect of storage on the system and whether the global social welfare of the system improves as agents adopt storage. In more detail, we measure social welfare by considering the following standard metrics of an electricity market:

- Diversity factor (DF) is the ratio of the sum of the individual maximum demands of various consumers of the system to the maximum demand of the complete system. The diversity factor is usually greater than 1.
- The Load Factor (LF) is the average power divided by peak power, over a period of time and, ideally, is 1. A low LF suggests peak demands in the system.
- The Grid Carbon Content intensity is the carbon produced to generate the required electricity. It is expressed as g of CO_2 per kWh and, is ideally as low as possible. The carbon emission from electricity generation in the UK is given in Figure 3 and is calculated from the generation supply mix from the UK National Grid for the period of August and September 2009.

4. A GAME-THEORETIC ANALYSIS

In this section we analyse the models given above from a game-theoretic point of view. For tractability we assume that agents have homogeneous efficiency and running costs, that is $\alpha^a = \alpha$ and $c^a = c$ for all $a \in \mathcal{A}$ for some α and c .

Formally, the game we consider has players which coincide with our agents, $a \in \mathcal{A}$, and the game describes the outcome of a single 24 hour interval. The pay-off an agent receives is equal to minus the total costs that agent experiences when purchasing electricity that day, $-\sum_{i \in \mathcal{I}} p_i(\ell_i^a + b_i^a)$. The strategy space available to each agent is the set of feasible storage profiles, $b_i^a \forall i \in \mathcal{I}$ where $-b_-^a \leq b_i^a \leq b_+^a$. Here we also make two further restrictions on feasible storage profiles. Firstly, for all $a \in \mathcal{A}$ we require that the amount of energy discharged is equal to the amount of energy charged multiplied by the efficiency, that is $\sum_{i \in \mathcal{I}} b_i^{a-} = \sum_{i \in \mathcal{I}} \alpha b_i^{a+}$. Secondly, for all $a \in \mathcal{A}$, we require that $\sum_{i \in \mathcal{I}} b_i^{a+} \leq e^a$, that is the total amount charged is less than the storage capacity. This is a stricter constraint than simply requiring

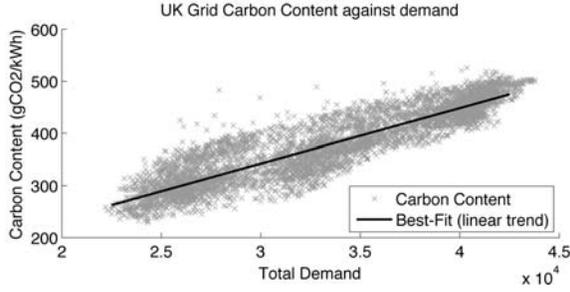


Figure 3: UK Carbon emission from Electricity Generation.

that the capacity is never exceeded at any time. However, it is a reasonable gauge of storage capacity limitations for a day-long time period, where demand typically goes through a single cycle of low to high to low, implying that storage devices would go through a corresponding cycle of charging to discharging to charging. We now proceed to characterise the deterministic Nash equilibria for this game.

4.1 Nash Equilibria as Global Optimisers

Suppose agents have chosen some strategy profiles, and let us consider the effect of a feasible change in strategy for one agent. That is, some $a \in \mathcal{A}$ considers going from $b_i^a \forall i \in \mathcal{I}$ to $b_i^a + \Delta b_i \forall i \in \mathcal{I}$, for some values Δb . The change in payoff for agent a would be:

$$\sum_{i \in \mathcal{I}} ((s_i(q_i) - s_i(q_i + \Delta b_i))(\ell_i^a + b_i^a) - s_i(q_i + \Delta b_i)\Delta b_i).$$

As noted in the previous section, since we are examining widespread micro-storage devices, we can assume that for all $i \in \mathcal{I}$ and $a \in \mathcal{A}$, ℓ_i^a , b_+ , and b_- are small in comparison to d_i and q_i . This means that the $(\ell_i^a + b_i^a)$ term in the above will be small, and $s_i(q_i + \Delta b_i)$ will be close to $s_i(q_i)$. Using these approximations, the deterministic Nash equilibria for this game correspond to strategy profiles which minimise $\sum_{i \in \mathcal{I}} \int_0^{q_i} s_i(x) dx$.

These approximations reduce our search for the Nash equilibrium of a complex multi-player game to a relatively straightforward global optimisation problem — that of minimising global generators costs. We now proceed to find solutions to this optimisation problem.

4.2 Characterisation of Nash Equilibria

We seek to find an aggregate storage profile $b_i \forall i \in \mathcal{I}$ with $-b_- \leq b_i \leq b_+$, $\sum_{i \in \mathcal{I}} (b_i)^+ \leq e$, and $\sum_{i \in \mathcal{I}} \alpha (b_i)^+ = \sum_{i \in \mathcal{I}} (b_i)^-$, which minimises $\sum_{i \in \mathcal{I}} \int_0^{d_i + b_i} s_i(x) dx$. Here, and throughout the paper, we use the notation $(\cdot)^+$ to denote positive part, i.e. $y = (x)^+$ means $y = x$ if $x > 0$, $y = 0$ otherwise. Likewise we use $(x)^-$ to denote $(-x)^+$.

We begin with a definition.

DEFINITION 1. For a storage system as described, we define the discharging price point, p^d , to be the maximum of the solution to $\sum_{i \in \mathcal{I}} q_i^d(p^d) = \alpha \sum_{i \in \mathcal{I}} q_i^c(\alpha p^d - c)$, and the solution to $\sum_{i \in \mathcal{I}} q_i^d(p^d) = \alpha e$, if one exists.

Here we define, for each interval i , $q_i^d(p) = \max(b_-, (d_i - s_i^{-1}(p))^-)$, and $q_i^c(p) = \min(b_+, (d_i - s_i^{-1}(p))^+)$.

We define the charging price point, p^c , to be the minimum of $\alpha p^d - c$ and the solution to $\sum_{i \in \mathcal{I}} q_i^c(p^c) = e$ if one exists.

This definition is well defined by the Lemma 1 in the Appendix. We can now state the main result of this analysis.

THEOREM 1. For a storage system as described, the set of Nash equilibria for the system is precisely the set of agent strategies where, for all $i \in \mathcal{I}$, $b_i = q_i^d(p^d) - q_i^c(p^c)$.

PROOF. See the Appendix. \square

4.3 Idealised Scenarios

For special idealised scenarios, we have the following two corollaries.

COROLLARY 1. If b_+ and b_- are sufficiently large, then for all i , $p^c \leq s_i(q_i) \leq p^d$. Furthermore, if for any $i \in \mathcal{I}$, $p^c < p_i < p^d$, then $b_i = 0$.

PROOF. If, for all $a \in \mathcal{A}$ we let b_+^a and b_-^a be equal to e^a , then this does not break our smallness assumption, and, furthermore, for all $i \in \mathcal{I}$, we'll have $q_i^d(p^d) < b_-$ and $q_i^c(p^c) < b_+$. Thus, for all i , if b_i is non zero then either $b_i = q_i^d(p^d) = s_i^{-1}(p^d) - d_i < 0$, in which case $q_i = s_i^{-1}(p^d)$ and so $p_i = p^d$, or else $b_i = q_i^d(p^d) = s_i^{-1}(p^c) - d_i > 0$, in which case $q_i = s_i^{-1}(p^c)$ and so $p_i = p^c$. If $b_i = 0$ then $s_i^{-1}(p^c) \leq d_i \leq s_i^{-1}(p^d)$ and so $p^c < p_i < p^d$ as required. \square

So, if the charge and discharge rates are sufficiently high, then we could expect prices to always lie within p^c and p^d .

COROLLARY 2. If b_+ and b_- are sufficiently large, capacity e is sufficiently high, $c = 0$ and $\alpha = 1$, then for all i , $p^c = s_i(q_i) = p^d$.

PROOF. This follows directly from the previous corollary. \square

Hence, in an idealised scenario, with perfectly efficient, cost free, and high capacity storage, we would expect the market prices over time to flatten to a single value. This is because perfect storage capability would allow agents to transport energy from any time interval to any other time interval free of charge. Thus, different suppliers in different time intervals would have to compete with each other, resulting in convergence to a single market price.

4.4 Rationality Assumption

Theorem 1 gives the aggregate storage behaviour when our game is in a deterministic Nash equilibrium. We can use this result to specify limits of the social welfare benefit that can result from adopting small-scale storage. If the actions of such selfish agents are to result in stable aggregate behaviour, then we can do no better than the outcome described above.

However, in using game theory, we have made some implicit assumptions, specifically that agents are rational and have complete information about the market throughout the time period. In reality, information available to those owning storage devices will not be perfect. Furthermore, even with perfect information, it might not be apparent to an automated agent which strategies are preferable. Instead, the agents themselves must adapt over time, to become aware of the repeating daily patterns of supply and demand, and learn which storage strategies are preferable. This is a difficult problem and it is not guaranteed that selfish learning behaviour can converge. For example, if agents over-react to perceived opportunities in the market, cycles of price fluctuations could develop.

In the next section, we provide a novel adaptive storage strategy for agents to maximise their savings. Under this scheme, agents change their storage profiles each day to be closer to their perceived optimal strategy. In Section 6 we show that provided the adaptation is not too fast, it results in the aggregate convergence predicted by Theorem 1.

5. AN ADAPTIVE STORAGE STRATEGY

As discussed above, the next step of our work is to design a novel adaptive storage strategy that an agent can use to decide on when to store energy and when to use the stored energy. Now, because market prices are continuously changing as a result of changing demand (due to consumers using storage devices), we design a learning mechanism that adapts to these changing market prices.

In more detail, our strategy is based on a day-ahead best-response storage. Because the market prices are unknown *a priori*, we can only calculate the storage profile on a day-ahead basis, as a best-response to the predicted market prices. To mitigate prediction errors, the consumer gradually adapts her storage towards the best-response storage. In this section, we first describe how we calculate the day-ahead best response storage profile and, second, we describe our learning mechanism, that is, how the consumer adapts her storage.

5.1 The Day-Ahead Best-Response Storage

The objective of an agent a is to minimise its costs by storing energy when prices are low and using that energy when prices are high. Now, because market prices are unknown until the aggregated load of all consumers, s_i , where $\sum_{a \in \mathcal{A}} \ell_i^a = s_i$, is known, the agent needs to decide on its storage profiles based on the predicted market prices of the following day. Note that in our work, we assume that market prices do not move significantly over days and use a weighted moving average to predict future market prices.³

We compute the storage profile, $b^a = b^{a+} - b^{a-}$ at every time-slot during the day as the solution to an optimisation problem where we minimise the following cost function⁴:

$$\arg \min_{b^a} \left(\sum_{i \in \mathcal{I}} p_i (b_i^{a+} - b_i^{a-} + \ell_i^a) + c^a e^a \right) \quad (1)$$

subject to the following constraints:

Constraint 1: storage efficiency

$$\sum_{i \in \mathcal{I}} b_i^{a-} = \alpha^a b_i^{a+} \quad \forall i \in \mathcal{I}$$

Constraint 2: within charging and discharging capacity

$$b_i^{a-} \leq b_-^a \text{ and } b_i^{a+} \leq b_+^a \quad \forall i \in \mathcal{I}$$

Constraint 3: energy that can be stored or used at a time-slot

$$\begin{aligned} b^{a-} &\leq \alpha^a \left(e_0^a + \sum_{j=1}^{i-1} (b_j^{a+} - b_j^{a-}) \right), \quad \forall i \in \mathcal{I} \\ b^{a+} &\leq e^a - e_0^a + \sum_{j=1}^{i-1} (b_j^{a+} - b_j^{a-}), \quad \forall i \in \mathcal{I} \end{aligned}$$

Constraint 4: no-reselling allowed

$$\ell_i^a - b_i^{a-} \geq 0, \quad \forall i \in \mathcal{I}$$

The last constraint can be removed in a system where consumers are allowed to sell power to the grid and that e^a can be *fixed* or *unconstrained*. In line with our model (see Section 3), c^a is the relatively small discounted running cost of using storage, e^a is the storage capacity. α^a is the efficiency of the agent's storage, b_+^a is its maximum charging and b_-^a

³As we will demonstrate later on, this is not very sensitive in our work as the price movements are generally small. However, a number of more sophisticated prediction algorithms, such as Gaussian Processes could be used instead.

⁴We used IBM ILOG CPLEX 9.1 to implement and solve the optimisation problem.

its maximum discharging rates and e_0^a is the storage at the beginning of the day which equals the storage at the end of the day (i.e., charging at the end of a day for the next day).

Because market prices move over trading days, the agent needs to continuously adapt its storage profile to reflect these changes. Now, because of the relatively high cost of storage, it is more sensible and realistic for the agent to gradually change its capacity by analysing the trend of market prices. To this end, we develop a novel learning mechanism to adapt storage profiles in an electricity market.

5.2 Learning in the Market

Our learning mechanism is based on a two-pass approach. Initially, the agent computes the maximum storage capacity, e_U^{a*} it would require to minimise its costs. e_U^{a*} is the cost-minimising capacity by optimising over e^a (see Equation 1).

Now, e_U^{a*} constitutes a *desired capacity* towards which the agent learns its storage capacity, i.e. it adapts its storage capacity progressively to follow the changing market trends. The storage capacity of the agent is defined by Equation 2 as $e^a(t)$ that follows the desired unconstrained storage capacity e_U^{a*} such that:

$$e^a(t+1) = e^a(t) + \beta_1 (e_U^{a*} - e^a(t)) \quad \forall i \in \mathcal{I} \quad (2)$$

where $e^a(0) = 0$ by default and β_1 is the learning rate⁵ of the storage capacity of agent a . Given its storage capacity, the agent then computes its optimal storage profile for the following day by *fixing* e^a at $e^a(t+1)$ in Equation 1.

On the second pass, given its current storage profile, the agent adapts its storage profile (to mitigate any risk of having poorly predicted market prices) as follows:

$$b_i^a(t+1) = b_i^a(t) + \beta_2 (b_i^{a*} - b_i^a(t)) \quad \forall i \in \mathcal{I} \quad (3)$$

where b^{a*} is the *desired storage profile* given as the optimal storage profile subject to a fixed storage capacity of $e^a(t+1)$ and β_2 is the learning rate of the storage profile. Note that we analyse in more detail the sensitivity of the learning parameters as part of the empirical study of the system in the next section.

6. AN EMPIRICAL ANALYSIS OF THE UK MARKET

In this section, we empirically analyse the effect of storage on the UK market. To this end, given our macro-model of the UK market, we setup individual consumers with typical UK load profiles,⁶ different learning rates and different storage types. Learning rates, as are charging and discharging capacities, are uniformly distributed⁷ (to represent consumers with different learning attitudes and storage types) around means that are based on current technologies (see Section 2). Now, because our macro-model is general enough, our framework can be applied to any electricity market around the world, and our results and insights broadly generalise.

Given this setup, we first provide a game-theoretic solution, identifying the Nash equilibrium of the system and, second, we provide a dynamic analysis as to whether or not

⁵As we will empirically demonstrate further on, the choice of the learning rates determines the evolutionary stability of the system and has to be reasonably small.

⁶We do so by adding random noise to the average UK load profile; on the timings using Poisson distribution of demand times and a uniform distribution of noise over demanded quantities.

⁷In all experiments except for when we analyse the effect of learning rates, we use a mean value of 0.05.

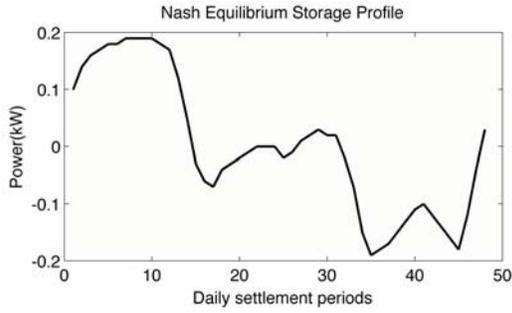


Figure 4: Nash equilibrium (with a storage capacity of 3.55 kWh).

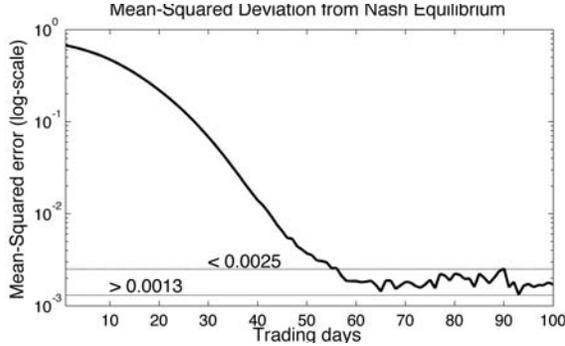


Figure 5: Convergence of the average strategy profile to the Nash equilibrium.

such an equilibrium can be reached if a proportion of the population were to acquire storage devices as well as use our adaptive storage strategy with the aim of maximising their individual savings. Finally, we analyse how the social welfare of the system evolves with a large number of agents adapting their storage and whether the social welfare of the market improves while consumers are able to make a saving.

6.1 A Game-Theoretic Solution

Given the game-theoretic framework outlined in Section 4, we first calculate the Nash equilibrium given a typical domestic average unconstrained profile (see Figure 4). It is clear that equilibrium behaviour for a consumer is to charge at off-peak hours (at night) and use the stored energy during peak hours (after working hours) when the consumers' load is highest.

6.2 Evaluation of the Adaptive Storage

Given the adaptive storage strategy we designed in the paper, we now analyse how the system evolves as agents are changing their behaviours within a realistic setting and whether the system converges to the Nash equilibrium. As we can see from Figures 4 and 5, our average storage profile indeed converges to the Nash Equilibrium of our game-theoretic analysis. Figure 6 shows how the average storage profile evolved towards an equilibrium as market prices were flattened in the system (see Figure 7).

Given these results (i.e., that we converge to the Nash Equilibrium and hence the optimal solution — see Section 4), we can claim our adaptive strategy sets the benchmark for *any learning strategy* in this system! Now, it is also

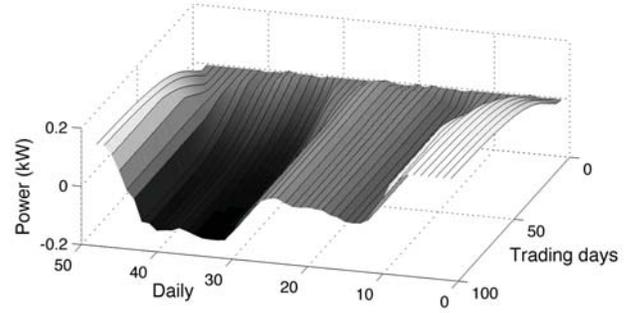


Figure 6: Average Storage Profile converging to Nash Equilibrium.

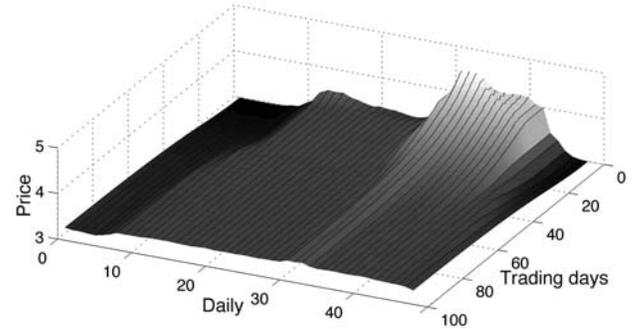


Figure 7: Changing Market Prices (market prices eventually flatten).

important to analyse how the social welfare of the system evolves as the system is evolving to the Nash equilibrium to ensure that agents adopting storage does not break the system (i.e. social welfare does not decrease). To this end, we analyse the market diversity factor, load factor and carbon content reduction (see Section 3).

For the system efficiency given in Figure 8, we considered a population with around 38% with storage capability (our choice of 38% will become clearer further on). As can be clearly observed, the system efficiency improves and gradually converges as agents adapt their storage profiles and market prices are flattened. In more details, the average maximum storage capacity required converges to around 3.55 kWh after several trading days while the market load factor converges to around 0.94 where the load in the market is nearly flattened. Furthermore, the diversity factor increases suggesting that, because of storage, consumers now have less correlated demand requirements from the electricity market (which generally reduces peaks in a system).

A significant benefit of storage at a micro-level is that if a sufficient proportion of the population does adopt storage, the carbon intensity of electricity market would decrease appreciably as peak demands are reduced. Indeed, in Figure 9, we show how the carbon content is reduced by up to 7% for different proportions of population adopting storage (by extrapolating our results to the 26M UK households).

Furthermore, from Figure 10, it is clear that there is a financial incentive for consumers to adopt storage, with a maximum saving of around 13% (based on the current system with no storage). As expected, because storage flattens

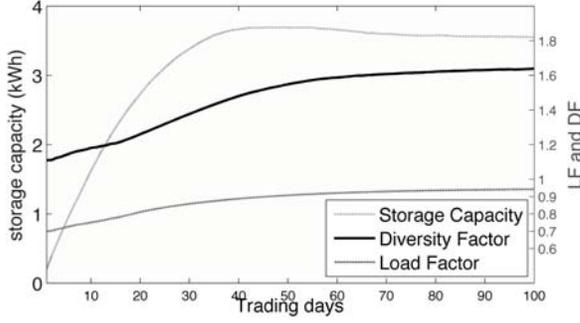


Figure 8: Social Welfare of System.

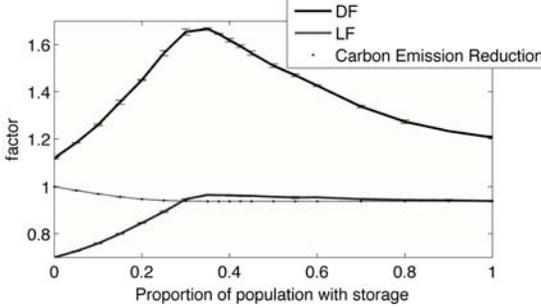


Figure 9: Social welfare for different proportion of the population using storage.

the market prices, other consumers, even without storage, also indirectly benefit. Now, as storage becomes more and more popular (as consumers become aware that they can save on their electricity bills), we observe a decrease in their savings, reaching a point where a consumer can save more by not having storage than having storage (see average savings in Figure 10). In the next subsection, we analyse in more detail this social trend.

6.3 When to Adopt Storage

Here, we formulate the problem as a game where agents have a mixed strategy $x_r \in (0,1)$, i.e. a probability that they have storage capability and are only motivated by financial gains. By analysing how x_r evolves as the payoffs change for different x_r , we want to study how the proportion of the population using storage evolves. To this end, we use the classical evolutionary game-theory (EGT) [10] based on the following equations:

$$\dot{x}_r = [u(e^r, x) - u(x, x)]x_r \text{ where } u(x, x) = \sum_{r \in S} u(e^r, x)x_r$$

$$x_{nash} = \arg \min_{x \in (0,1)} \sum_S (\max[u(e^j, x) - u(x, x), 0])^2$$

First, we compute a heuristic payoff table (based on simulations) to calculate the payoffs for using and not using storage for different x_r . The replicator dynamics \dot{x}_r describes the dynamics of the population, i.e. how x_r is evolving, and, whether it converges to any Nash Equilibrium x_{nash} .

Figure 10 shows our EGT analysis, with x_{nash} at 0.382 adopting storage, and the replicator dynamics all converging towards that equilibrium. Surprisingly enough, this means that the population will gradually settle at an equilibrium where only 38% of the population use storage. At that equilibrium, all consumers make an average savings of 8.54%

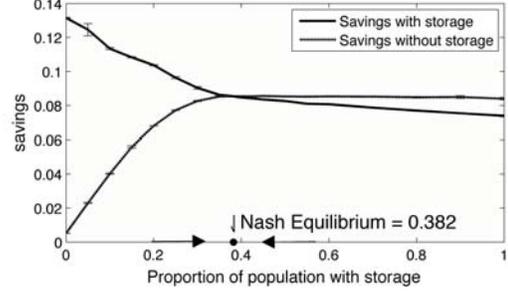


Figure 10: Savings with and without storage. Replicator dynamics (arrows on the x-axis) converge to the Nash Equilibrium at 0.382 with a saving of 8.5% for using and not using storage.

(i.e. an annual saving of GBP60 per household – based on an average annual electricity bill of GBP675). Now, the equilibrium suggests that too many consumers storing can be counterproductive for the system. This is because there is a point beyond which additional storage adds more volatility to already flattened market prices (seen from a decrease in the load factor in Figure 8), and those agents that store are more exposed to this volatility. Finally, around the equilibrium point, we also observe that the social welfare of the system peaks such that the individual goals of the agents (to save on the electricity bills) is well aligned with maximisation of the social welfare, with the diversity factor DF decreasing as too many households start storing energy.

6.4 Sensitivity of the Learning Mechanism

Finally, we analyse the sensitivity of the learning mechanism against the social welfare and the agent’s self-interested objectives. Figure 11 shows, as expected, that the smaller the learning rate, the more efficient the system (with a higher load factor) and the better the average savings of the individual agents. Now, because an infinitely small learning rate is infeasible as it implies an infinitely long time to reach the equilibrium, a trade-off is required. Specifically, because the learning parameters are not very sensitive when they are small, a value of 0.05 to 0.20 would be reasonable. A high learning rate, on the other hand, would result in agents adopting their optimal storage profile immediately rather than adapting gradually, which clearly results in poor savings and poor system efficiency. Finally, there is no convergence, implied by the load factor dropping below 0.7 – without any storage – to 0.59 such that there are now more peaks in the system (as everyone is charging at the same time).

7. CONCLUSIONS

In this paper, we developed a framework to analyse agent-based micro-storage management for the smart grid. Specifically, we designed a storage strategy (with an adaptive mechanism based on predicted market prices) for consumers and empirically demonstrated that the average storage profile converges towards a Nash equilibrium. At that point, peak demands are reduced, reducing the requirements for more costly and carbon-intensive generation plant. Moreover, in our analysis of the social welfare at this equilibrium we show that, while being stable, it results in reduced costs and carbon emissions. This also shows that the objective of buying

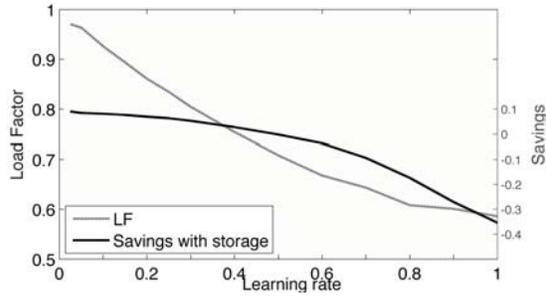


Figure 11: The effect of the learning rate.

storage to save on electricity bills is aligned with maximising social welfare. Finally, we show that the population would adopt storage until an equilibrium of 38% is reached, around which the social welfare is maximised.

For future work, we intend to integrate a more accurate model of the electricity market mechanism in our work as well as models of deferrable loads and how an agent can control such loads in parallel with its storage for more efficient cost-saving behaviours. Furthermore, we would like to explore how consumers with preferences for low carbon electricity, not just low cost electricity, could interact within this model.

8. REFERENCES

- [1] B. Daryanian, R. Bohn, and R. Tabors. Optimal demand-side response to electricity spot prices for storage-type customers. *Power Systems, IEEE Transactions on*, 4(3):897–903, 1989.
- [2] U. S. Department-Of-Energy. Grid 2030: A National Vision For Electricity’s Second 100 Years, 2003.
- [3] L. Exarchakos, M. Leach, and G. Exarchakos. Modelling electricity storage systems management under the influence of demand-side management programmes. *International Journal of Energy Research*, 33(1):62–76, 2009.
- [4] R. Galvin and K. Yeager. *Perfect Power: How the MicroGrid Revolution Will Unleash Cleaner, Greener, More Abundant Energy*. McGraw-Hill Professional, 2008.
- [5] A. Holland. Welfare losses in commodity storage games. In *Proceedings of The 8th International Conference on Autonomous Agents and Multiagent Systems*, pages 1253–1254, Budapest, 2009.
- [6] M. Houwing, R. R. Negenborn, P. W. Heijnen, B. D. Schutter, and H. Hellendoorn. Least-cost model predictive control of residential energy resources when applying μ chp. In *Power Tech*, pages 425–430, London, UK, 2007.
- [7] M. Korpaas, A. T. Holen, and R. Hildrum. Operation and sizing of energy storage for wind power plants in a market system. *International Journal of Electrical Power & Energy Systems*, 25(8):599–606, October 2003.
- [8] D. MacKay. *Sustainable energy: without the hot air*. UIT, Cambridge, 2009.
- [9] K. H. van Dam, M. Houwing, and I. Bouwmans. Agent-based control of distributed electricity generation with micro combined heat and power—cross-sectoral learning for process and infrastructure engineers. *Computers & Chemical Engineering*, 32(1-2):205 – 217, 2008.
- [10] J. W. Weibull. *Evolutionary Game Theory*. MIT Press, Cambridge, MA, 1995.
- [11] J. Williams and B. Wright. *Storage and Commodity Markets*. UIT, Cambridge, 1991.

Appendix

This appendix contains the lemmas and proofs left out of the main body of the paper. We begin with a lemma which justifies the definition of charging and discharging price points given in Section 4.

LEMMA 1. *There always exists a solution to, $\sum_{i \in \mathcal{I}} q_i^d(p) = \alpha \sum_{i \in \mathcal{I}} q_i^c(\alpha p - c)$. Furthermore, if p is the solution then $p^d = p$ and $p^c = \alpha p - c$ unless $\sum_{i \in \mathcal{I}} q_i^d(p) > \alpha e$, in which case, p^d is the solution to $\sum_{i \in \mathcal{I}} q_i^d(p^d) = \alpha e$, and p^c is the solution to $\sum_{i \in \mathcal{I}} q_i^c(p^c) = e$.*

PROOF. If p is sufficiently small then for all $i \in \mathcal{I}$ we will have $s_i^{-1}(p) < d_i$ and hence $q_i^d(p)$ will be strictly positive and, since $\alpha p - c < p$, $q_i^c(\alpha p - c)$ will be zero. Likewise if p is sufficiently large then for all $i \in \mathcal{I}$ we will have $s_i^{-1}(p) > d_i$ and so $q_i^c(p)$ will be strictly positive and $q_i^d((p + c)/\alpha)$ will be zero. Since the functions $q_i^d(\cdot)$ are decreasing for all i and $q_i^c(\cdot)$ are increasing for all i , we can conclude that $\sum_{i \in \mathcal{I}} q_i^d(p^d) - \alpha \sum_{i \in \mathcal{I}} q_i^c(\alpha p^d - c)$ is a continuous decreasing function in p which is negative for sufficiently small p and positive for sufficiently large p . This implies the existence of some solution \hat{p} such that $\sum_{i \in \mathcal{I}} q_i^d(\hat{p}) = \alpha \sum_{i \in \mathcal{I}} q_i^c(\alpha \hat{p} - c)$.

Now if $\sum_{i=1}^n q_i^d(\hat{p}) \leq \alpha e$ then $p^d = \hat{p}$ and $\sum_{i \in \mathcal{I}} q_i^c(\alpha \hat{p} - c) \leq e$, hence $p^c = \alpha \hat{p} - c$. If $\sum_{i=1}^n q_i^d(\hat{p}) > \alpha e$ then, since we know that $\sum_{i=1}^n q_i^d(p) = 0$ for large enough p , there must exist some $\tilde{p} \geq \hat{p}$ such that $\sum_{i=1}^n q_i^d(\tilde{p}) = \alpha e$. By definition, p^d must be equal to this \tilde{p} . Similarly, since $\sum_{i \in \mathcal{I}} q_i^c(\alpha \hat{p} - c) > e$, we can deduce that $p^c \leq \alpha \hat{p} - c$ and $\sum_{i=1}^n q_i^c(p^c) = e$, as required. \square

We now prove Theorem 1 from Section 4.

PROOF THEOREM 1. We seek to find an aggregate storage profile $b = \{b_i\}_{i \in \mathcal{I}}$ which minimises $f(b)$ where $f(b) = \sum_{i \in \mathcal{I}} \int_0^{d_i + b_i} s_i(x) dx$. If, for all $i \in \mathcal{I}$ we extend the definition of $s_i(x)$ to be 0 for negative x , then we can see that $f(\cdot)$ tends to infinity as for large feasible b . Thus, $f(\cdot)$ must have at least one local minimum over the feasible domain, one of which has to be the global minimum. To do find these allocations we seek feasible b for which the derivative of $f(b)$ is non-negative in every direction that leads to another feasible allocation. The gradient of $f(b)$ is $\{p_i\}_{i \in \mathcal{I}}$, thus it remains to characterise all b such that $\sum_{i \in \mathcal{I}} p_i \Delta b_i \geq 0$ for every Δb where $b + \Delta b$ is feasible.

Now suppose we have some b which locally maximises $f(\cdot)$. If there is i, j with $b_+ > b_i > 0$ and $b_+ > b_j > 0$, then it would be feasible to increase b_i and decrease b_j by an equal quantity, (or vice versa), hence we must have $p_i = p_j$. From this we can deduce that if for some i, j , $p_i < p_j$, $b_i > 0$ and $b_j > 0$, then we must have $b_i = b_+$. This means that there will be some price, \hat{p}^c such that if $b_i > 0$, for any i , then $p_i \leq \hat{p}^c$, with equality if $b_i < b_+$. Similarly, we can show that there will be some price \hat{p}^d such that if $b_i < 0$ for i , then $p_i \geq \hat{p}^d$, with equality if $b_i > -b_-$. Furthermore, there cannot be i, j such that $p_i + c > \alpha p_j$ and $b_i > 0$ and $b_j < 0$, for then it would be feasible to decrease b_i by some Δb_i and increase b_j by $\Delta b_j = \alpha \Delta b_i$. Hence, we must have $\hat{p}^c + c \leq \alpha \hat{p}^d$. This implies that for all i , $b_i = q_i^d(\hat{p}^d) - q_i^c(\hat{p}^c)$.

Furthermore, if for some i, j , $b_i > 0$ and $b_j < 0$ $p_i + c < \alpha p_j$, then it would be profitable to increase b_i by some Δb_i and increase b_j by $\Delta b_j = \alpha \Delta b_i$. So, either $\hat{p}^c + c = \alpha \hat{p}^d$ and this does not happen, or else $\hat{p}^c + c < \alpha \hat{p}^d$ but this change is never feasible due to the capacity constraint. In which case $\sum_{i \in \mathcal{I}} (b_i)^+ = e$ and $\sum_{i \in \mathcal{I}} (b_i)^- = \alpha e$. Thus, we have that $\hat{p}^c = p^c$ and $\hat{p}^d = p^d$, and so all local maximisers of $f(\cdot)$ must be as in the statement of the theorem. However, this precisely specifies the storage profile, and so there can only be one local minimum of $f(\cdot)$, which is also the global minimum, and is given by the statement of the theorem, as required. \square