

Agent-Based Homeostatic Control for Green Energy in the Smart Grid

SARVAPALI D. RAMCHURN, PERUKRISHNEN VYTELINGUM, ALEX ROGERS,
and NICHOLAS R. JENNINGS, University of Southampton

With dwindling nonrenewable energy reserves and the adverse effects of climate change, the development of the smart electricity grid is seen as key to solving global energy security issues and to reducing carbon emissions. In this respect, there is a growing need to integrate renewable (or green) energy sources in the grid. However, the intermittency of these energy sources requires that demand must also be made more responsive to changes in supply, and a number of smart grid technologies are being developed, such as high-capacity batteries and smart meters for the home, to enable consumers to be more responsive to conditions on the grid in real time. Traditional solutions based on these technologies, however, tend to ignore the fact that individual consumers will behave in such a way that best satisfies their *own* preferences to use or store energy (as opposed to that of the supplier or the grid operator). Hence, in practice, it is unclear how these solutions will cope with large numbers of consumers using their devices in this way. Against this background, in this article, we develop novel control mechanisms based on the use of autonomous agents to better incorporate consumer preferences in managing demand. These agents, residing on consumers' smart meters, can both communicate with the grid and optimize their owner's energy consumption to satisfy their preferences. More specifically, we provide a novel control mechanism that models and controls a system comprising of a green energy supplier operating within the grid and a number of individual homes (each possibly owning a storage device). This control mechanism is based on the concept of homeostasis whereby control signals are sent to individual components of a system, based on their continuous feedback, in order to change their state so that the system may reach a stable equilibrium. Thus, we define a new *carbon-based* pricing mechanism for this green energy supplier that takes advantage of carbon-intensity signals available on the Internet in order to provide real-time pricing. The pricing scheme is designed in such a way that it can be readily implemented using existing communication technologies and is easily understandable by consumers. Building upon this, we develop new control signals that the supplier can use to incentivize agents to shift demand (using their storage device) to times when green energy is available. Moreover, we show how these signals can be adapted according to changes in supply and to various degrees of penetration of storage in the system. We empirically evaluate our system and show that, when all homes are equipped with storage devices, the supplier can significantly reduce its reliance on other carbon-emitting power sources to cater for its own shortfalls. By so doing, the supplier reduces the carbon emission of the system by up to 25% while the consumer reduces its costs by up to 14.5%. Finally, we demonstrate that our homeostatic control mechanism is not sensitive to small prediction errors and the supplier is incentivized to accurately predict its green production to minimize costs.

Categories and Subject Descriptors: I.2.11 [Computing Methodologies]: Distributed Artificial Intelligence
General Terms: Algorithms, Economics

Additional Key Words and Phrases: Agents, computational sustainability, electricity, multiagent systems, agentbased control

This work has been done as part of the iDEaS project (<http://www.ideasproject.info>).

Authors' address: S. D. Ramchurn, P. Vytelingum, A. Rogers, and N. R. Jennings, Intelligence, Agents, Multimedia Group, School of Electronics and Computer Science, University of Southampton, Southampton, SO17 1BJ, UK; email: {sdr, pv, acr, nrj}@ecs.soton.ac.uk.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies show this notice on the first page or initial screen of a display along with the full citation. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, to republish, to post on servers, to redistribute to lists, or to use any component of this work in other works requires prior specific permission and/or a fee. Permissions may be requested from Publications Dept., ACM, Inc., 2 Penn Plaza, Suite 701, New York, NY 10121-0701 USA, fax +1 (212) 869-0481, or permissions@acm.org.

© 2011 ACM 2157-6904/2011/07-ART35 \$10.00

DOI 10.1145/1989734.1989739 <http://doi.acm.org/10.1145/1989734.1989739>

ACM Reference Format:

Ramchurn, S. D., Vytelingum, P., Rogers, A., and Jennings, N. R. 2011. Agent-Based homeostatic control for green energy in the smart grid. *ACM Trans. Intell. Syst. Technol.* 2, 4, Article 35 (July 2011), 28 pages. DOI = 10.1145/1989734.1989739 <http://doi.acm.org/10.1145/1989734.1989739>

1. INTRODUCTION

Achieving energy security and reducing carbon emissions have been recognized as two of the most important challenges of this century given the rapid depletion of worldwide oil and gas reserves and the potentially devastating effects of climate change [DECC 2009a; US Department Of Energy 2003]. To this end, the creation of a smart electricity grid has been advocated as a key component in delivering efficient low-carbon energy. In particular, the vision of a smart grid includes technologies that reduce losses in the transmission of electricity, that integrate intermittent renewable energy generators (such as wind turbines and solar panels), and that attempt to reduce or shift demand by providing real-time information about consumption and prices using smart meters in homes.

In this article we focus in particular on the pressing problem of integrating renewable energy generation into the smart grid. This is an important challenge because the integration of renewable energy generators requires a significant shift from the traditional principle of “supply follows demand” whereby generators always have to keep up with demand in the grid. To date, this has been acceptable only because the output from nonrenewable energy sources (such as coal or gas fired power stations) can be turned up or down at will. In contrast, renewable energy generators (i.e., green suppliers) are sensitive to seasons and weather conditions and are therefore variable and intermittent in their output. This means that they cannot be powered up at will to meet the demand, nor is it possible to accurately predict exactly how much energy they will generate. While, from a technical point of view, integrating renewable energy requires the development of technologies to ensure power flows in a controlled manner from renewable energy generators to the grid, from an economic point of view, renewable energy suppliers need to devise new strategies to trade effectively in the electricity markets in which they operate in order to ensure that they can make a profit while still satisfying the demand from their consumers given the intermittency in their supply.

When the supplier is not able to satisfy the demand of its customers from its own renewable generation capacity, it must typically buy additional electricity at short notice (usually incurring a high cost which reduces its potential profits) from the wholesale electricity market within which it operates [Kirschen and Strbac 2004]. Conversely, when the supplier produces more than its customers demand, it must sell this extra electricity at a lower price than what it would typically obtain from its customers.¹ It is therefore crucial for the supplier to predict its production accurately in order to bid effectively in markets, but more importantly, the supplier should be able to *influence the demand* from its customers to mitigate the impact of low production.

Now, while the issue of predicting renewable energy production and bidding in electricity markets has been actively researched in the power systems domain [Kirschen and Strbac 2004; Milligan et al. 2009; Morales et al. 2010], there is relatively little work studying how to influence the demand of residential consumers (see Section 2 for more details). Typical approaches to demand management involve the supplier sending a predicted real-time price, via the consumers’ smart meter, to incentivize

¹In general, a supplier acts as the middle-man between the generators from which it buys electricity and the customers to which it sells electricity. However, we consider the case of the generator and the supplier being the same actor as is often the case in the U.K. or U.S. electricity market, particularly with regards to “green” focused suppliers such as Ecotricity (U.K.), New England GreenStart, or Green Mountain Energy (U.S.).

them to increase or decrease demand [Faruqui and George 2005]. These mechanisms assume that by feeding more information to consumers, the latter will be able to understand the conditions on the grid (e.g., generator outages or a high price for electricity at peak time) and react to such conditions by switching devices on or off, or changing their thermostat settings on their heaters (or air conditioners), and in future, use electricity stored in their home or electric car battery. However, such an assumption is challenged by the fact that the signals from the grid, aside from being rather complex for the average user [Schweppe et al. 1989], require the user to perform complex calculations in order to optimize her consumption and storage of electricity (i.e., to know what is the best time to store or use electricity). Given this, in this article, we advocate the use of autonomous software agents as a decentralized control paradigm whereby each owner would have her own agent, running on her smart meter, advising on the best consumption pattern she could adopt to achieve her goals (e.g., minimizing carbon emissions and/or costs). A key advantage of this approach is that consumers do not need to be experts at energy trading or scheduling in order to maximize their benefits, therefore their adoption allows for a larger impact in implementing demand management strategies based on complex signals from the grid.

In more detail, agents can be endowed with the ability to learn their owner's preferences from their consumption history and can run optimization algorithms that solve complex scheduling tasks (e.g., heating cycles of boilers or charging profiles of their electric car battery) with very little direct input, if at all, from their owner. In so doing, these agents can minimize their owner's cost, given price signals from the grid, by efficiently scheduling the on/off cycles of deferrable devices (such as washing machines or boilers) or the storage profiles of batteries (i.e., storing when prices are low and using the stored electricity when prices are high). However, even if a real-time price signal is provided by the grid or supplier, agents have no indication as to *how much* they should change their usage (by deferring loads) or store electricity in order to meet the capacity of the renewable energy supplier. This is important because, if, for example, the agents are incentivized to consume more, based on a low predicted price, the aggregate might end up consuming much more than can be provided by the supplier. Conversely, given a high price, the collection of agents might consume much less than what the supplier produces leaving the supplier with unused green energy. Hence, in general, relying solely on the price signal may result in a mismatch between the production capacity of the supplier and the demand from the consumers that might require the supplier to buy extra capacity at a high price or sell its extra production at a loss on the wholesale electricity market. The supplier may, in turn, pass on such costs to its customers and also increase the consumers' carbon emissions as it uses carbon-emitting generators (from which it buys on the wholesale market) to supply their extra demands. Moreover, providing only a price signal ignores the fact that some agents may also be aiming to reduce their owner's carbon emissions (as there is a growing population of such green users [Kockar et al. 2009]) since such a signal ignores the carbon intensity (i.e., the carbon emissions per unit of energy) of the energy being supplied. Given these issues with real-time pricing in this domain, it is crucial that better control mechanisms are developed to mitigate the occurrence of such mismatches in order to make the integration of renewable energy into the smart grid more profitable and more reliable. In particular, given the reaction of the agents to such signals, the controller (i.e., supplier) should be able to provide feedback to the agents in order for them to adjust their consumption profile while still allowing them to minimize their costs and their carbon emissions. The challenge here is to generate such signals and stabilize the system without the supplier being able to fully observe the state of the individual agents in terms of their profile of electricity usage. This challenge is further exacerbated by the fact that the agents may not wish to reveal their battery storage capacity (which determines by how

much they can defer demand to follow supply) as they may not trust the provider to use this information to their benefit.

To meet this challenge, in this article, we introduce the concept of *agent-based homeostatic control* to coordinate a green supplier and its consumers, who may own storage capability, within the smart grid.² In particular, we model a system composed of a green energy supplier (e.g., a wind and/or solar powered generator) providing electricity to a large number of consumers whose aims are to reduce their individual costs and carbon emissions. To capture the consumers' preferences and goals, we develop agents that can optimize their owner's battery storage profile according to signals from the grid. Our model aims to capture real-life settings where, for example, a third-party wind turbine/photovoltaic array provides energy to a small number of houses or buildings, or where a regional/national supplier serves customers across a wide area. We then provide a novel control mechanism to ensure that demand from the agents follows the production from the green supplier. Our approach is inspired from the principle of homeostasis [Marieb and Hoehn 2007]. The latter describes a living organism's adjustment of its internal interdependent components to maintain a stable constant condition and was first used in the context of the electricity grid by Schweppe et al. [1980]. The analogy between our system and a living organism is appropriate given that our system is based on the three underpinning concepts of homeostasis: sensing, sending a control signal, and feedback. First, our system is able to *sense and predict*³ the effect of external conditions (i.e., weather in our case) using widely available weather sensors and satellite imagery. Second, the supplier can compute a *control signal* to be sent to the agents based on the impact of weather conditions on its production of electricity. Third, through the resulting response of the agents in the system (i.e., the combined effect of its load and storage profiles), *feedback* is provided to the supplier which decides how to adapt its control signal to ensure the system reaches the state it desires. Through several iterations of the preceding process, the system aims to converge to a more efficient solution (see Section 6). While the idea of homeostatic control in this domain was pioneered more than 30 years ago, it was not deemed practical at the time due to the lack of communication and distributed computing technologies that would allow large populations of consumers of different types (of different income and usage levels) to participate in such mechanisms. However, with the prevalence of broadband communications⁴ and the development of agent-based technologies for a wide variety of applications including the energy domain [Jennings et al. 1996; Bussmann et al. 2004], it is now possible to revisit this proposition and instantiate homeostatic control mechanisms within the smart grid.

In more detail, this article advances the state-of-the-art in the following ways:

- (1) We implement a novel real-time carbon-based pricing mechanism that builds on the carbon intensity signals provided by the grid (e.g., the National Grid in the U.K. or Eire Grid in Ireland already provide this data in real time) through the Internet.⁵ Using this mechanism, suppliers, using renewable energy sources or not, do not

²Our approach can be generalized to consider loads such as washing machines or boilers that can be deferred to achieve results similar to using storage, but this is beyond the scope of this article.

³Note that we do not develop novel sensing and prediction algorithms as this is beyond the scope of this work. Rather, we assume such predictions are available and focus on building the control mechanism that uses such predictions. However, as we show in Section 6, our mechanism is robust to errors in prediction and hence, as better prediction technology is developed, the benefits accrued from using our mechanism will only improve. Future work will study such prediction algorithms in more detail.

⁴Google PowerMeter (<http://www.google.com/powermeter>) and AlertMe (<http://www.alertme.com>) systems already utilize home broadband links in order to monitor energy consumption over the Web.

⁵See <http://www.bmreports.com> and <http://www.ideasproject.info/software>.

need to install additional hardware to compute and transmit a real-time price to consumers as it can easily be obtained over an existing broadband connection. We show that by setting a price plan to match the carbon intensity of the grid, suppliers can effectively align the carbon-based prices to the real-time prices they effectively trade at on the grid. Such a pricing scheme also means that, should demand be higher than can be met by the green supplier, the greenest form of energy from the grid (where a mix of green and nongreen suppliers also exist) will be bought by the supplier to meet excess demand. By so doing, carbon-sensitive consumers are incentivized to subscribe to the supplier since they can minimize both their costs and carbon emissions at the same time. In turn, this helps the supplier reduce its reliance on other grid energy sources with high carbon emissions (for which the supplier may be taxed according to the market regime it operates within [Kockar et al. 2009; Hobbs et al. 2010]). It is also important to note that our pricing scheme is not restricted to just a green supplier as our prices (except the price of energy strictly obtainable from the green supplier) are only dependent on the conditions on the grid (i.e., grid carbon intensity reflecting generation costs). Thus, our pricing scheme could potentially be used by any supplier trying to appeal to an environmentally conscious consumer aiming to reduce her contribution to the overall carbon emissions of the grid.

- (2) We develop a novel control signal which ensures that agents are incentivized to store a given quantity of electricity according to changes in green energy generation capacity caused by changing weather conditions. The signal provides an estimate to the agent of *how much* it should aim to change its storage behavior in order to minimize its costs.
- (3) We develop algorithms, based on mathematical programming, for agents to optimize their energy storage, given the control signal they receive, in order to minimize their costs, and in so doing, their carbon emissions.
- (4) We provide a novel adaptive mechanism for the green supplier to learn the best signal to send to the agents in order to maximize its efficiency. In particular, the mechanism works even when the supplier is unaware of the proportion of the consumers in the system that possess enough storage to cope with the variability in supply (which cannot be accurately predicted by the supplier).
- (5) We empirically evaluate our approach and show that our agent-based homeostatic control with our novel adaptive mechanism (compared to a traditional real-time pricing mechanism) results in up to 25% improvement in efficiency (which maximizes the use of green energy by its pool of consumers and hence reduces their carbon emissions) for the supplier and up to 14.5% savings for the agents. Furthermore, we show how errors in the supplier predicting its green production can result in higher costs or carbon emissions. In particular, we show that the supplier is incentivized to accurately predict its green production to minimize its costs and that our homeostatic control mechanism is not sensitive to small prediction errors.

The rest of this article is structured as follows. Section 2 presents the general background of our work and outlines the rationale behind our design choices. Section 3 describes the model of the green supplier and the agents, each endowed with their own preferences and their battery storage capabilities. Given this, in Section 4, we present our novel carbon-based pricing mechanism. Then, Section 5 details our homeostatic control mechanism. Section 6 evaluates our mechanism and Section 7 concludes.

2. BACKGROUND AND RELATED WORK

The issue of integrating renewable energy supplies into the grid to reduce carbon emissions while ensuring their intermittency does not have an adverse impact (e.g., by

requiring more spinning reserves or causing brownouts) has been actively researched for a number of years. The seminal idea of using homeostatic control to account for intermittency (for any energy source type) was initially suggested by Schweppe et al. [1980], who foresaw that improvements in communication technologies would pave the way for new control mechanisms to coordinate energy suppliers with users in order to account for shortfalls (e.g., by asking them to reduce consumption) or excesses (e.g., by letting them consume more) in energy supplies in real time [Schweppe et al. 1980]. Such mechanisms would then allow demand to follow supply. Their work, however, remained preliminary and, as such, untested under simulated settings. In particular, they did not model the individual behavior of agents (in terms of their basic preferences and how they optimized their consumption) and did not consider the issue of carbon emissions to be important at the time. Moreover, they used the frequency of Alternating Current (AC) as a vehicle for real-time pricing. In the absence of other communication mechanisms, this choice was probably the only viable one. However, while frequency is a good indicator of generation failing to meet demand, it does not correspond to the cost of electricity in real-time electricity markets [Schweppe et al. 1989]. In contrast, our more up-to-date approach is able to exploit recent communication and computational advances in order to model individual agent behaviors and develop pricing mechanisms and control signals that can be communicated in real time while still matching the current market conditions.

Since Schweppe's work, most research in the power systems community has focused on facilitating the integration of intermittent sources of energy by providing more reliable predictions and control of energy generation from such sources [Milligan et al. 2009; Morales et al. 2010] or by matching renewable energy generators with high-capacity electricity storage providers (either attached to wind turbines or performing arbitrage in an electricity market) [Burgio et al. 2009; Lubosny and Bialek 2007; Bathurst and Strbac 2003; Makarov et al. 2010; Korpaas et al. 2003]. Hence, most of this work has focused on the use of large "utility-scale" batteries controlled by one user as opposed to small *domestic* batteries (e.g., electric cars or water heaters⁶) where each battery is individually owned and where the owner is only intent on maximizing her own preferences. This is mainly due to the fact that coordinating a large number of individual users (owning batteries) with suppliers relies on robust communications that were previously considered too expensive [Kirschen and Strbac 2004] and that such storage facilities may be too expensive to implement on a large scale (i.e., across 26M households in the U.K.). However, with the advent of high-speed (mobile) broadband and less expensive battery storage technology (e.g., with electric cars such as the Mini E and the Nissan Leaf, and deep-cycle battery technologies from Ceramatec, for example), this argument is no longer valid. Hence, several approaches and field trials (e.g., the GridWise Project, the Energy Demand Research Project [Smith 2010] or PowerMatcher⁷) have emerged which consider the use of demand-side management techniques (DSM) (based on real-time pricing using AC frequency or time of use pricing) to control demand [Infield et al. 2007; Hammerstrom et al. 2008]. However, such approaches do not address the use of storage devices to defer the consumption of electricity and again ignore the individual preferences of the users.⁸

⁶Note that the storage of hot water or air is effectively equivalent to the storage of electrical energy in some settings.

⁷See more details at <http://gridwise.pnl.gov> and <http://www.powermatcher.net>.

⁸The GridWise project trials allowed users to set their preferences for the price at which devices should turn off. However, these preferences were assumed fixed throughout the day. However, it was recognized that users' preferences to use certain devices (e.g., to use heating or hot water) vary across the day and, as a result, they were given the capability to override their settings at any time.

A number of agent-based approaches that do consider the incentives of individual actors on the grid have been developed over the last decade. For example, Ygge initiated work on abstracting electricity markets as multicommodity markets and showed how agents trading energy for different times of the day could generate efficient allocations [Ygge et al. 1999]. Moreover, Jennings et al. [1996] developed coordination mechanisms for different actors on the grid to manage the allocation of transmission capacity [Jennings et al. 1996] and, more recently, Sun and Testfatsion [2007] and Li and Testfatsion [2009] have developed agent-based electricity market simulations which factored in transmission constraints and different types of buyers and sellers in the grid. Kok and Venekamp [2010] take a similar approach and implement an agent-based architecture to run the electricity market where the individual actors represent either generators or devices in the home whereas Gerding et al. [2011] and Chalkiadakis et al. [2010] adopt a game-theoretic approach to managing electric vehicle charging and organizing collectives of wind energy generators respectively. Moreover, Li et al. [2010] recently proposed a mechanism that coordinates a population of agents in optimizing their loads to avoid going over a given limit on demand. Their agents are designed to satisfy both the individual's and the system's goals (in capping demand). However, their approach requires agents to communicate their individual plans (to optimize their resources) to the system, leading it to be scalable to no more than ten thousand agents. In general, none of these works deals specifically with the variability of supply when renewable energy sources are involved. Moreover, these approaches do not consider the daily consumption profile of consumers and how an agent might optimize its consumption to solely maximize its owner's benefit (in cost savings). We believe this is crucial because agents have to be profitable to their users, each with her own specific needs and lifestyle, in order to be commercially viable. Moreover, in our recent work, we showed how such agents can be modelled to individually optimize the storage of electricity given a real-time price and how this led to demand on the grid being flattened (without explicitly exchanging plans among agents) [Vytelingum et al. 2010; Ramchurn et al. 2011]. While we did consider the preferences of the agents, we did not address the issue of external influences (i.e., weather conditions in our case) on the supply of electricity. Hence, this article bridges a major gap in this area by modeling the agents in the system, each with its own preferences, and by providing the control mechanism to ensure demand follows supply.

Our work is also inspired from the recent upsurge in eco-friendly energy providers such as Ecotricity,⁹ Green Mountain Energy,¹⁰ or New England GreenStart¹¹ that supply their consumers from their own green energy sources. Their success stems from the fact that there is a large portion of consumers intent on minimizing their individual carbon emissions. However, such companies cannot satisfy all their demand from their own renewable generators (e.g., Ecotricity only generated around 41% of the energy they sold in 2009–2010) and they have to buy the rest from other conventional sources. Given this, the carbon-based pricing scheme we propose in this article departs from the fixed pricing scheme provided by these companies, to incentivize users to minimize their carbon emissions in real time while also minimizing their costs. Thus, our work sets the stage for new tariffs that can readily be applied once smart meters are rolled out.¹²

⁹<http://www.ecotricity.co.uk>.

¹⁰<http://www.greenmountainenergy.com>.

¹¹<http://www.massenergy.org/renewable-energy/negs>.

¹²The U.K. government has planned to equip all 26M houses in the country with a smart meter by 2020 while in the U.S. it is expected that about 60M houses will be equipped by the same date. Similar targets are in place in Europe with Italy having already completed its smart meter deployment.

Furthermore, in future energy markets where carbon emissions may be taxed [Kockar et al. 2009; Gajbhiye and Soman 2010], our work on the carbon-based pricing scheme also reduces the suppliers' operational risks. The intuition is that if suppliers can pass on some of their costs, due to carbon taxes, to their consumers, the latter will be incentivized to align their behavior to the needs of the supplier. Thus, if consumers are charged for their carbon emissions, they will optimize their behavior to reduce their emissions, and hence reduce the tax paid. Proposed approaches to mitigating suppliers' carbon costs have focused on balancing the objective of minimizing costs and minimizing carbon emissions as a complex optimization problem [Hobbs et al. 2010; Wong et al. 2010; Kockar et al. 2009; Mazer 2007; Harris 2005]. Alternative approaches considering the use of carbon intensity signals, as we do here, are practically nonexistent since, we believe, such technological solutions involving the use of information and communication technologies, rather than metering devices on the grid, are less popular among power systems engineers. Hence, this article embodies the fundamentals of *computational sustainability* [Gomes 2009] by providing computational solutions to real-time electricity pricing that can be rapidly and reliably deployed in the smart grid to generate greater efficiencies.

3. GREEN SUPPLIER AND AGENT MODELS

In this work, we assume that it will not be possible, at least in the near future, to be fully reliant on green energy sources [DECC 2009b]. Rather, we focus on a more realistic scenario which is more representative of the case today where a green supplier is operating in a market where there are mostly brown (i.e., carbon emitting) sources of electricity (e.g., in markets run by National Grid¹³ or PJM¹⁴ on the U.S. east coast). Specifically, Figure 1 depicts the main elements of the system involving the supplier and the consumers. Thus, the grid (consisting of electricity markets and physical networks) and weather (e.g., wind speed or sunshine) act as external influences on them. Hence, in the following sections, we detail each element of our system and the information flows between them by providing a generic model of a green energy supplier (typically with wind or solar power generation capability) and consumers (represented each by its agent) which may possess electricity storage capacity (e.g., either batteries or electric vehicles) and which are keen to minimize their carbon emissions (hence their contract with the green supplier). In our models, we assume that a day is divided into a set of half-hourly slots $T = \{1, 2, \dots, 48\}$ such that each actor needs to decide its behavior (typically four hours before or a day ahead) for each slot. Smaller slots (e.g., 5 or 15 minutes) could be considered but half-hour slots are used here as this is a common interval used within electricity meters and electricity markets.

3.1. The Green Supplier

We model the green supplier (e.g., with wind turbines and/or solar power generators) as having a pool of consumers that subscribe to it for electricity.¹⁵ The supplier can predict, with a reasonable level of accuracy (usually four hours ahead using sensors and historical data [Milligan et al. 2009]), that it will generate a given amount of electricity $q_i \in \mathbb{R}^+, \forall i \in T$. Now, in line with today's reality we model our supplier as being able to satisfy only part of its consumers' demand, acquiring the remainder

¹³<http://www.nationalgrid.com/uk>.

¹⁴<http://www.pjm.com>.

¹⁵Our model considers only one green supplier operating in the grid where other possibly nongreen suppliers also operate. If we were to consider a grid populated by mostly green suppliers, then each green supplier would have to consider the possibility of not being able to supply its electricity from the grid at times when none of the others is able to generate sufficient electricity. This would mean adding a hard constraint on the consumers' demand and/or price very expensively any consumption above the level provided by the grid.

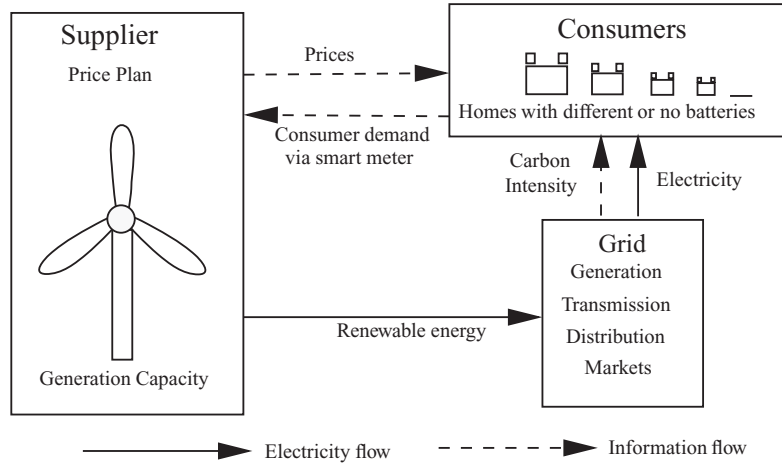
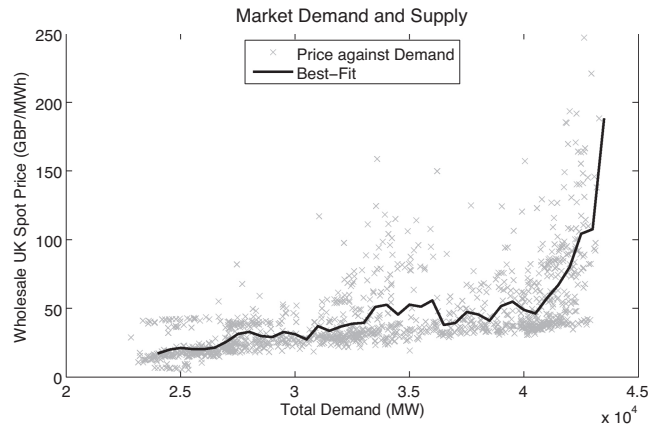


Fig. 1. Here we illustrate the electricity and information flows between the actors in the system we consider. Specifically, our model specifies the price plan used by the supplier, the heterogeneous storage devices used by the consumers, as well as the decision making involved in charging or using their battery. We also assume that the grid provides the carbon intensity signals which we use in our mechanism but do not consider how the electricity is transmitted and delivered to consumers by the grid.


from the wholesale electricity market. While the availability of supply from the grid is guaranteed, its price is generally driven by the market demand and supply (see the relationship between price and demand in Figure 2(a)). In particular, depending on its energy production pattern, the supplier has a number of potential ways of pricing its electricity in real time based on what it expects the wholesale spot market prices to be and any existing future and forward contracts that it holds. Thus, while spot prices determine the supplier's real-time trading strategy, forward and futures contracts aim to hedge the risk of very high peak prices in the spot market. Typically, buying extra capacity on the spot market is more costly than in a forward contract and selling extra capacity is usually done with a very low or possibly negative profit [Harris 2005]. Hence, the green supplier needs to ensure that it minimizes the amount of extra capacity it needs to buy in spot markets by getting its consumers to follow its production pattern. Moreover, a key challenge here is to align the goals of maximizing profit while meeting the goals of its customers to minimize their carbon emissions. Hence, we address these issues through a combination of pricing and control as we show in Sections 4 and 5.

3.2. The Consumer Agents

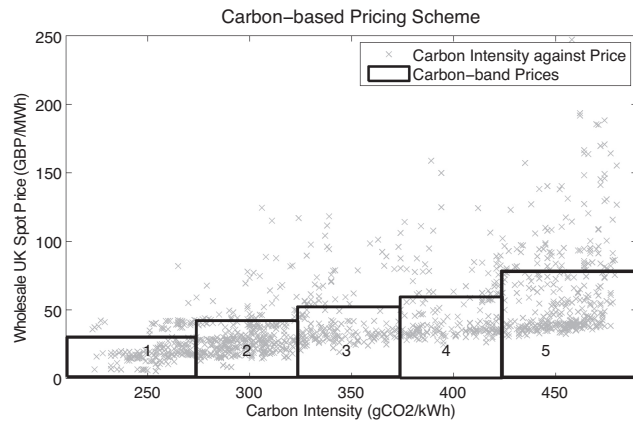
Here, we describe our model of the consumer agent, which builds upon and extends a recent model for homes equipped with smart meters [Vytelingum et al. 2010]. Thus, we consider a set of consumers individually represented by their agent, where the set of agents is \mathcal{A} and where each agent $a \in \mathcal{A}$ has a load profile $\ell_i^a \forall i \in T$, (i.e., its owner's predicted consumption of electricity) such that ℓ_i^a is the amount of electricity required by agent a for time interval i during each day. This electricity can either be bought from the supplier or retrieved from a storage device. Thus, each agent $a \in \mathcal{A}$ may possess an electricity storage device with capacity $e^a \in \mathbb{R}^+$, efficiency $\alpha^a \in [0, 1]$, and running costs $c^a \in \mathbb{R}^+$. Here, the cost c^a represents maintenance costs or the loss in value of the device through use over time. We assume this cost is directly proportional to the size of the device, such that if the size of the battery is e^a then the storage cost is $c^a e^a$. The storage efficiency α^a models the fact that if q_i kWh amount of energy is stored, then only $\alpha^a q_i$ kWh may be subsequently used.



(a) the correlation between UK spot price and grid demand



(b) the correlation between carbon intensity and grid demand



(c) pricing scheme obtained from combining the correlation between price and demand, and between carbon intensity and demand

Fig. 2. Relationships between carbon intensity and the price of electricity against demand on the balancing market (Source: NETA, U.K. August–September 2009).

The objective of each agent a is to minimize its owner's cost and hence, it will optimize 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 (as limited by the efficiency and capacity of the battery), and b_+^a , the charging capacity. Since the discharging efficiency of the storage is regimented by α^a , we separate out the charging and discharging profile such that 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. Then, we denote as $l_i^a = \ell_i^a + b_i^a$ the total consumption of an agent at time i and the aggregate consumption of all agents as $l_i^A = \sum_{a \in \mathcal{A}} l_i^a$.

In the next section we describe how consumers are charged in real time for their electricity and elaborate on the technology to achieve this.

4. THE CARBON-BASED PRICING MECHANISM

To communicate electricity prices, we develop a novel carbon-based pricing scheme that the supplier sends to its customers. The scheme relies on the availability of a real-time carbon intensity signal (measured in gCO_2/kWh and representing the amount of CO_2 emitted for every unit of energy consumed) that is broadcast on the Internet¹⁶ by national grid operators (as in the U.K. and Ireland).¹⁷ The key feature of our pricing scheme is that the users are guaranteed to pay for the *greenest* form of energy at all times if they optimize their consumption pattern (to minimize either costs or carbon emissions). While such a scheme allows the supplier to pass on some of the costs for operating in carbon-taxed electricity markets [Hobbs et al. 2010; Kockar et al. 2009], it also follows from the simple fact that consumers usually subscribe to green suppliers in order to reduce their carbon emissions and it would be simply inconsistent if any extra energy that the supplier cannot provide would be supplied by the worst carbon emitters for the same price. The pricing scheme we develop exploits the correlation between price and grid demand (shown in Figure 2(a)) and between carbon intensity and grid demand (shown in Figure 2(b)) (based on real U.K. data). Using this data, we can infer the correlation between price and carbon intensity as shown in Figure 2(c). Now, as with other examples of pricing mechanisms (e.g., fixed pricing, time-of-use, or real-time pricing), the retail prices must be calculated to reflect the supplier's retail margin profit and its exposure to the risk of peak prices in the wholesale market.¹⁸ However, their complexity must be balanced against the need to provide a simple pricing scheme that would appeal to, and be understood by, the average consumer. Thus, we provide retail rates for different bands of carbon emission (in our case, Band 1 to

¹⁶For example, the GridCarbon iPhone/Android app uses information provided by National Grid in the U.K. to provide real-time updates on the generation mix (i.e., the percentage of energy generated by different types of plant). Using the latter and a function mapping generation source to carbon intensity, the real-time aggregate carbon intensity is provided. For more details see <http://www.ideasproject.info/software>.

¹⁷The U.K. carbon intensity can be computed from the generation mix data provided by Elexon (<http://www.bmreports.com/>) and is provided on a mobile platform through the GridCarbon app. Ireland's real-time carbon intensity can be found at <http://www.eirgrid.com>. The typical carbon intensity of the U.K. grid tends to be around $500\text{gCO}_2/\text{kWh}$ (and range between 550 and 450 across a typical day) and that of Ireland around $225\text{kgCO}_2/\text{kWh}$ (and range between 280 and 200).

¹⁸While we describe the pricing scheme here, the methodology to calculate the retail prices is beyond the scope of this article (and not critical in our work) as it is very dependent on the different formats of futures and forward contracts that different suppliers would settle for. For the purpose of this work, we set the band prices based on a retail margin and considering spot prices of the wholesale market. The supplier would be required to calculate the retail prices for different carbon emissions on a month-ahead or even year-ahead basis, such that these rates would be fixed for long periods at a time and periodically reviewed to match changing grid prices and aggregate consumers' demand.

Band 5)¹⁹ rather than providing the consumer with a complex equation-mapping between carbon and price. More formally, these bands are computed as follows.

- (1) We split the users' aggregate consumption in terms of the amount satisfied by the green supplier and the amount drawn from other grid sources such that $l_i^A = l_i^{A,green} + l_i^{A,grid}$ at time $i \in T$. Thus, $l_i^{A,green} \in \mathbb{R}^+$ is the load from the consumers that is completely satisfied by the green supplier from its own renewable source such that $l_i^{A,green} \leq q_i$ (remember q_i is the production capacity of the supplier) while $l_i^{A,grid} \in \mathbb{R}^+$ is supplied from other (possibly carbon-emitting) sources on the grid.
- (2) $l_i^{A,grid}$ is priced by our carbon-based pricing scheme (based on five different bands) as p_n^{grid} where $n \in \{1, 2, 3, 4, 5\}$ and $p_n^{grid} < p_m^{grid}$ if $n < m$. We choose n based on the carbon intensity signal $\rho_i \in \mathbb{R}^+$ (kgCO₂/kWh) at time i , that is, p_n^{grid} is chosen if $\rho_n^{min} \leq \rho_i \leq \rho_n^{max}$ where ρ_n^{min} and ρ_n^{max} are the upper and lower bounds of the band, respectively. For example, according to Figure 2(c), $p_1^{grid} = 29 \text{ £/MWh}$ for $0 < \rho_i \leq 275$ and $p_2^{grid} = 41 \text{ £/MWh}$ for $275 < \rho_i \leq 325$.
- (3) $l_i^{A,green}$ (as a *Band 0*), that is, the amount consumed that is completely satisfied by the green supplier from its own sources, is priced at ϵ less than the greenest band, *Band 1* of the grid as $p^{green} = p_1^{grid} - \epsilon$.

Through the preceding pricing scheme the supplier incentivizes consumers to use the green energy (since it always less expensive) it produces rather than grid energy. It is important to note that our carbon based pricing scheme is not specific to a green supplier and could be used by all suppliers in general as the signals we used apply to the whole grid. Also, note that our approach is different from typical sliding scale or band pricing (where consumers simply pay different rates for different amounts consumed) in that, in our model, the band price (p/kWh) applied to a consumer's usage is independent of the energy consumed by the consumer and, instead, dependent on the real-time carbon intensity of the grid (incentivizing consumers to consume less when the grid has a high carbon intensity).

Building upon our pricing mechanism, we next develop an agent-based homeostatic control mechanism that provides additional signals to the agents to ensure demand follows supply.

5. HOMEOSTATIC CONTROL

Given the variability in weather conditions affecting the generation capacity, the system, comprising of the agents and the supplier, needs to continuously adjust to these conditions to maximize its efficiency. This needs to happen given that both the supplier and the consumers aim to maximize their individual profits. More importantly, they can only do so if they communicate to ensure that the aggregate demand from all the consumers is as close as possible to the real-time supply from the producer, while ensuring that any extra energy needed is bought at times when the carbon intensity of the grid is the lowest. Only by doing so can all of the energy the supplier produces be used (which ensures that both the supplier makes a profit and the users get inexpensive and greener electricity) *and* that any extra energy needed is bought at the least expensive price from the grid.

¹⁹We choose five bands only as a representative set of bands that generally covers the trend we identify in the correlation between carbon intensity and demand. While more bands would result in a more accurate reflection of the costs for the supplier, they would also be more complex for the consumers to understand we believe.

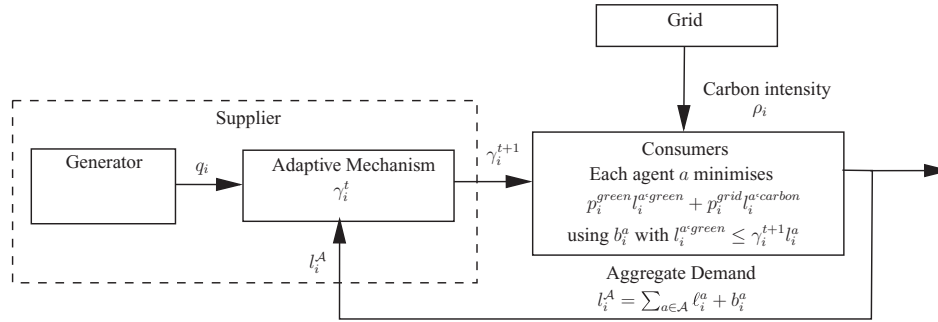


Fig. 3. Flow of information in the system involving the supplier and the consumers connected to the grid. Building on the system presented in Figure 1, consumers receive two signals; a carbon intensity signal ρ_i from the grid (which they use to predict the next day's carbon intensity) and (in addition to the price plan shown in Figure 1) γ_i^{t+1} (for the next day) computed by the supplier using q_i (i.e., its production for the current day) and l_i^A (i.e., the overall consumption for the current day). The consumers then optimize their electricity usage given the signal γ_i^{t+1} and using their battery. The supplier is informed of the amount consumed by all its customers (i.e., l_i^A), hence closing the loop. Note that we do not show the payments the supplier receives based on splitting the consumption of its customers in terms of their green and nongreen components (i.e., $l_i^A = l_i^{A,green} + l_i^{A,grid}$).

While the sensing and prediction of short-term weather conditions can be easily implemented using existing machine learning techniques [Alpaydin 2004], significant challenges lie in determining the appropriate control signal to be sent to the agents to ensure that their behavior in reaction to such a signal meets the aforesaid requirements. Moreover, the supplier may have to adapt the signals it sends based on the feedback (i.e., consumption levels at different times of the day) it receives. In this article, we address these challenges and detail the different elements of our solution (graphically represented in Figure 3) in the following subsections.

5.1. The Signal to the Agents

The information that is typically provided by the supplier to an agent is the price of electricity whether it is a fixed, real-time, or of our carbon-based type. Using this information, an agent can then manage its demand in order to both minimize its costs and also its carbon emissions. While this approach is suitable for most suppliers, it is inefficient for the green supplier with its own renewable energy generation capacity. Specifically, due to the intermittent nature of this generation capacity, the supplier needs to signal to its pool of consumers about its generation pattern so as to maximize the use of its green energy, thereby minimizing the carbon emissions of its consumers.

Now, a price signal would be an ineffective stimulus as it would provide no information about how much the agents need to increase their demand during any particular 30-minute period to maximize the use of green energy (as discussed in Section 1). Thus, we propose two new signals (see Figure 3) that address this shortcoming. First, consumers are given a grid carbon intensity signal. This signal may be provided either by the grid operators (over the Internet) or by the supplier (collecting the data from grid operators) through the meter (the band prices are agreed with the supplier well before). The signal allows the consumer to compute its cost of electricity based on our pricing mechanism outlined earlier on. Second, the supplier provides the consumer with a target $\gamma_i \in \mathbb{R}^+ \forall i \in T$ to increase ($\gamma_i > 1$) or decrease ($\gamma_i < 1$) its consumption with respect to the previous day's consumption at time i such that *demand follows supply* in the day ahead. This target γ_i is based on the aggregate consumption at time interval i given

by l_i^A and the production from the supplier q_i , and is given as $\gamma_i = \frac{q_i}{l_i^A}$. In Section 5.3 we will show how to adapt this signal to consider the feedback that the supplier gets from the population of agents. For now, however, assuming that all consumers in the system react to this signal (by storing electricity when most incentivized to do so), we expect the aggregate demand to approach this target, thus maximizing green energy usage and minimizing the carbon emissions of the consumers. Thus, if agents meet this target exactly, they are expected to pay only p^{green} for every unit of electricity used. However, if an agent overuses (above the target), it increases the expected cost of electricity as the supplier would need to acquire the additional demand at p_i^{grid} prices. On the other hand, if it underuses (below the target), the agent misses out on less expensive and greener electricity.

Next, we present an optimization model for an agent that aims to minimize its costs given its load and storage availability (see Figure 3). In particular, we consider that the agent obtains part of its energy demand from the green supplier described in Section 3.1 and is already contracted to pay p_i^{green} for every kWh provided by its supplier from its own sources (wind or sun). Then, the agent receives two signals from the supplier, that is p_i^{grid} , computed based on the carbon-pricing scheme agreed beforehand and the real-time carbon intensity of the grid, and $\gamma_i^{t+1} \in [0, 1]$. We assume that the agent obtains these signals a day ahead (but this could equally also be a few hours ahead as is commonly the case for wind generators which can predict reasonably accurately how much they are going to be able to provide in such time frames [Milligan et al. 2009]). Thus, the agent would need to reoptimize its behavior whenever it receives updates of γ_i .

5.2. The Agents' Behavior

Since we assume that an agent's daily consumption (load) profile is fixed (i.e., price does not affect the loads in a single time slot but can affect stored or discharged energy), an agent a can only try to minimize its costs by storing energy when prices are low and using as much of that energy as possible when prices are high. Thus, we consider agent a having to plan its storage for the day ahead, that is $t + 1$, if t is the current day. As opposed our previous model [Vytelingum et al. 2010],²⁰ the agent does not need to predict the next day's price for each time slot. Rather, we assume the agent predicts the next day's carbon intensity ρ_i for each half-hourly interval of the day and using the price plan proposed in Section 3.1, computes a price p_i^{grid} (since p_i^{green} is known to all agents a priori). Given the uncertainty in its own consumption for the next day, the agent also needs to predict $\ell_i^{a,t+1}$ which is its own expected total load at time i for the next day and here we assume that $\ell_i^{a,t+1} = \frac{1}{k} \sum_{t-k}^t \ell_i^{a,t}$, that is, a moving average over k previous (similar) days' consumption (we choose $k = 3$).²¹

Now, the agent also has to decide how to adjust its consumption for individual time slots based on the signal γ_i sent by the supplier. Since γ_i^{t+1} implies a factor change in the current day's overall consumption $l_i^{a,t}$, the best an agent can do is to try and bound the total amount of energy it consumes at i by $\gamma_i^{t+1} l_i^{a,t}$. By so doing, the agent

²⁰The model presented here is also different from our previous work in that the aggregate demand of the population of agents does not influence prices or carbon intensity values on the grid. This is due to the fact that the population of agents tied to the green supplier is considered too small to influence the grid's overall performance.

²¹Consumption on average tends to be similar during weekdays but weekday consumption is typically different from weekend consumption. Moreover, load profiles tend to change across seasons and for special days (e.g., Christmas or for special sporting events). Hence, agents would have to take these differences into account in their prediction if such a system were deployed.

$$\text{minimise } \sum_{i \in \mathcal{I}} p_i^{\text{green}} l_i^{\text{a,green}} + p_i^{\text{grid}} l_i^{\text{a,grid}} + c^a e^a \quad (1)$$

where $l_i^{\text{carbon}} \in \mathbb{R}^+$, subject to:

Constraint 1: aggregate load

$$l_i^{\text{a,t+1}} = l_i^{\text{a,t}} + b_i^{\text{a+}} - b_i^{\text{a-}}, \forall i \in \mathcal{I}$$

Constraint 2: separate usage into green and non-green

$$l_i^{\text{a,t+1}} = l_i^{\text{a,green}} + l_i^{\text{a,grid}}$$

Constraint 3: restrict green part of consumption

$$0 \leq l_i^{\text{a,green}} \leq \gamma_i^{\text{t+1}} l_i^{\text{a,t}}$$

Constraint 4: storage efficiency

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

Constraint 5: within charging and discharging capacity

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

Constraint 6: energy that can be stored or used at a time-slot

$$b_i^{\text{a-}} \leq \alpha^a \left(e_0^a + \sum_{j=1}^{i-1} (b_j^{\text{a+}} - b_j^{\text{a-}}) \right), \forall i \in \mathcal{I}$$

$$b_i^{\text{a+}} \leq e^a - e_0^a + \sum_{j=1}^{i-1} (b_j^{\text{a+}} - b_j^{\text{a-}}), \forall i \in \mathcal{I}$$

Constraint 7: no-reselling allowed

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

Fig. 4. Mathematical program used by an agent to find its optimal storage strategy.

can maximize the probability that it will pay the best price p^{green} , as offered by the supplier (as explained in Section 3.1).

Our solution to this optimization problem is formulated as a mathematical program as given in Figure 4. The output is the storage profile, $b^a = b^{\text{a+}} - b^{\text{a-}}$ at every time slot during the day.

The objective of the agent is captured by the minimization function in Eq. (1).²² Then Constraint 1 expresses the fact that the total load of the agent at any point in time is the sum of the amount it stores and the amount it consumes. Constraints 2 and 3 decouple the amount consumed by the agent in terms of the green and nongreen components since the green component (i.e., $l_i^{\text{a,green}}$) is limited by the supplier according to its signal γ_i through $\gamma_i l_i^{\text{a,t}}$ (in Constraint 3). Constraints 4, 5, and 6 restrict the flow in the battery to ensure it respects the battery's efficiency limitations (as limited by α^a),

²²The storage optimization model we present here tries to capture the key features of the storage device and the load profile of the consumers. While other constraints could be added (e.g., a bound on the cost the consumer is willing to pay, or the maximum energy allowed to be charged in any time slot to prevent the battery from overheating), we believe we have captured the main ones and have validated this with our industrial partners.

that the amount (dis)charged in any time slot respects the limits b_i^{a-} and b_i^{a+} , and that the energy that can be stored or used at any point during the day respects the state of the battery at that point in the day. Finally, Constraint 7 prevents the agent from simply storing electricity to sell back to the grid. As can be seen, the mathematical program only has linear constraints and a linear objective function. Hence, linear programming solvers (e.g., CPLEX or glpk) can be directly applied to find the solution. In our implementation, we found that the optimal solution is obtained within a few milliseconds.²³

Note that the last constraint can be removed in a system where consumers are allowed to sell power back 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 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).

Now, not all agents may possess storage capacity (or some may not have enough storage to cover their whole consumption) and this is not known to the supplier a priori. The effect of this can be quite dramatic (as we shall show in Section 6). In particular, with limited storage, agents are less able to react to the variability in supply (communicated through the γ signals). This, in turn, can result in the supplier incorrectly estimating the outcome of aggregate consumers' response to its signals, leading to greater demand from the grid at peak times and hence higher prices. In the next subsection, we present a novel adaptive mechanism that allows a supplier to determine the appropriate γ_i to be sent to a population of agents that may have limited storage.

5.3. The Adaptive Mechanism

The aim of the adaptive mechanism is to learn the optimal signal to send to the agents (see Figure 3) (given their response to a previously sent signal). Now, if the supplier were to have complete and perfect information about the system (i.e., the proportion of consumers with demand-side management capabilities that can use or store energy in reaction to a supplier's signal) and complete control of the system (i.e., control the capabilities owned by its consumers), it would be able to define (and communicate) exactly the storage profile of each consumer to minimize its total cost. The same optimization model defined in Eq. (1) could be used by aggregating the charging and discharging capabilities of all consumers and their loads as

$$\min \sum_{a \in A} \sum_{i \in \mathcal{I}} p_i^{green} \ell_i^{a,green} + p_i^{grid} \ell_i^{a,grid} + c^a e^a, \quad (2)$$

subject to Constraints 1 to 7 as in Figure 4. In so doing, we can find an optimal solution to the behaviors of consumers that will result in the most efficient use of green energy.

Now, in a more realistic setting, the only feedback that the supplier obtains from its consumers is their total load on the grid, that is $\ell_i + b_i$ (as shown in Figure 3). Assuming that the supplier can accurately estimate the average load of each consumer ℓ_i^a , a simplistic approach to set the control signal γ_i^{t+1} would be

$$\gamma_i^{t+1} = \frac{q_i^t}{\sum_{a \in A} \ell_i^a}. \quad (3)$$

²³We used IBM's ILOG CPLEX 12.2 to implement the linear programming model.

$$\gamma_i^{t+1} = \begin{cases} \gamma_i^t + \beta_S \frac{q_i^t}{l_i^{A,t}} & , \text{ if } q_i^t > l_i^{A,t} \\ \max\left(0, \gamma_i^t - \beta_S \frac{q_i^t}{l_i^{A,t}}\right) & , \text{ otherwise} \end{cases}$$

where $\beta_S \in [0, 1]$ is the learning rate

Fig. 5. The adaptive mechanism given a chosen learning rate β_S and inputs q_i^t and $l_i^{A,t}$.

This is simply the ratio between the total green energy from the supplier q_i^t and the total load from the agents at time slot i . The aim here would be to tell the agents to adjust their loads (using the battery b_i^a) so that γ_i ideally converges to 1 (i.e., demand following supply). However, this approach assumes that all consumers have storage capability, as the supplier is signaling to each one of its consumers to increase its demand by a portion of the green energy produced. Ideally, the supplier should do so only for those consumers that are able to react to such a signal (i.e., those with storage capability). However, such information is not readily available to the supplier as consumers do not have to and may not want to reveal such information to their supplier (because they may not trust it to use this information to their benefit). Hence, the challenge is to adapt γ_i using only the feedback from the agents. We present our solution to this problem in Figure 5.

In more detail, we propose a novel adaptive version of γ_i based on current aggregate demand and the supplier's green energy production. As can be seen, using the adaptive signal, the homeostatic control mechanism tries to minimize the difference between green energy produced and current total demand based on a learning rate β_S which determines the rate at which it will change γ_i to reflect the feedback it receives. This rate is critical in determining whether the mechanism converges to the optimal solution (as we will see in Section 6). Thus, when the green energy produced is higher than the total demand, it implies an opportunity to maximize the use of green energy and reduce the carbon emissions of the consumers is lost and therefore γ_i is compensated positively (i.e., top condition in Figure 5) to increase the total demand of all the consumers. However when the total demand is too high, the supplier signals its consumers to reduce their use of green energy (i.e., bottom condition in Figure 5) at time i (and ensuring that γ_i is never negative). As we empirically demonstrate in Section 6, the daily total demand gradually changes such that the efficiency of green energy usage converges to the optimal (defined by the optimal solution of Eq. (2)).

In the next section, we evaluate our homeostatic control mechanism for different proportions of the population with storage capability and demonstrate how our adaptive system remains robust (i.e., keeps costs low and efficiency of green energy usage high) even when the error with which the supplier predicts his own supply of green energy (i.e., q_i) increases.

6. EVALUATION

The aim of this evaluation is to empirically demonstrate the effectiveness of the homeostatic control mechanism presented in the previous section in inducing demand to follow supply such that the carbon emissions of the system are reduced. We adopt an empirical approach here, as opposed to an analytical approach seeking equilibrium characterisations, in order to evaluate the system under a wider variety of settings than

would be feasible within a theoretical framework,²⁴ and predict equilibrium outcomes (for differently parameterized versions of our mechanism) without overly restrictive assumptions. As we will see, despite the utility-maximizing behavior of consumers with storage capability, the system quickly converges to the optimal behavior. In this context, we define optimal behavior as the solution to Eq. (2) as discussed in Section 5.3.

In what follows, we first detail our experimental setup and then go on to evaluate the efficiency of our homeostatic control mechanism (against the optimal behavior) in terms of green energy use and costs of electricity, in particular for different proportions of the population with storage capability. We focus on the proportions with storage capacity to identify the critical amount of storage needed in the population to achieve the maximum returns on investment in batteries for individual users and for the population as a whole. We also demonstrate how the supplier's adaptive behavior significantly improves the system. Finally, we analyse the sensitivity of our mechanism against prediction (of green energy production) errors.²⁵

6.1. Experimental Setup

We simulate a pool of 5000 consumers²⁶ subscribed to a green supplier that is able to supply 50% of its demand from its own renewable sources and the remainder bought from the grid. Our choice of the proportion of green energy was designed to model one of the most popular green suppliers in UK, namely *Ecotricity* which produced 41% of the electricity it supplies to its consumers from its own renewable sources in 2009–2010 and expects to reach 50% within the year.

Electricity is priced using the carbon-based pricing scheme detailed in Section 4. Furthermore, we model the generation capacity as values drawn from a uniform distribution (between 0 and 1) for each of the 30-min time slots. Each value is averaged using a discrete Gaussian filter (i.e., using normally distributed weights with mean 0 and standard deviation of 1 to model the effect of gradually changing wind patterns by removing any high-frequency changes) over the vector of 48 time slots across the day. The values of the time slots are then linearly scaled such that the sum total of these values equals 50% of the daily demand of the consumers. Furthermore, we set the limit for the prices bands at $\rho^{min} = \langle 0, 275, 325, 375, 425 \rangle$, $\rho^{max} = \langle 275, 325, 375, 425, \infty \rangle$ and $p^{grid} = \langle 29, 41, 51, 59, 71 \rangle$ assuming a 10% profit margin for the supplier. Finally, the demand of each consumer is modeled on the real UK data profile of the average user on the *Domestic Unconstrained Tariff* (i.e., the typical UK average profile before any demand-side management technologies, such as storage, are used [Vytelingum et al. 2010]) and the simulation is run 500 times, each run lasting 100 days. These values were chosen to ensure the statistical significance of our analysis.²⁷

²⁴We acknowledge that a theoretical analysis of our mechanism would allow us to ascertain the properties of our control mechanism and, since this goes well beyond the scope of this article, we aim to do such an analysis in future work.

²⁵Apart from the model presented by Vytelingum et al. [2010], we are not aware of any other micro-storage optimization models that have been devised to respond to real-time prices. Since Vytelingum et al.'s model did not consider the control signal we use in our case, we benchmark against it as a system where homeostatic control is not used.

²⁶Larger population sizes do not show any significant differences to the insights we provide here. Specifically, the results with different population sizes showed that the system becomes more efficient as the population size increases. We also experimented with real data from a population of about 1000 users in the south west of England and confirmed that similar results were obtained to the simulation provided the optimization carried out by agents was carried out more frequently (on a half-hourly basis) to take into account the spiky nature of their consumption (as opposed to the smooth cycles simulated in our experiments which allow a daily optimization to work). However, due to U.K. data protection issues, such results cannot be reproduced here.

²⁷Using multiple samples of the results, we validated our conclusions using a t-test with $\alpha = 0.05$ and found that there was no significant difference between the means of the samples.




Fig. 6. Efficiency of green energy usage for the system with and without the adaptive mechanism. The system with no homeostatic control is only 79.5% efficient throughout.

Next, we assume that only a proportion of the population has storage capability (i.e., they own storage devices capable of storing 4kWh with charging and discharging efficiency uniformly distributed between 0.3kW and 0.5kW to model current storage devices [Vytelingum et al. 2010]). The consumers with storage capabilities are thus able to optimize their storage profile to reduce their costs (and implicitly reduce their carbon emissions as a result of our carbon-based pricing scheme). Given a control signal from the supplier, we empirically demonstrate in this section how the system efficiency improves in terms of carbon emissions (i.e., by minimizing energy used from the grid and maximizing green energy use).

6.2. Efficiency of Homeostatic Control

First, we analyse the daily effect of homeostatic control on the system. We do so by comparing the green energy usage efficiency (i.e., the percentage of green energy that is directly used by the consumers contracted to the supplier) and the cost of electricity used (based on the carbon-based pricing scheme and the *Band 0* pricing of the supplier) in a system with and without homeostatic control. Furthermore, we also demonstrate the effectiveness of our adaptive mechanism and how it allows the system to converge to an optimal solution (given in Section 5.3).

From Figure 6, given a population with 50% having storage capability, we can see a daily convergence of our homeostatic control system with an adaptive mechanism, for different learning rates β_S , to the optimal solution at 91.8%. As can be seen,

the nonadaptive version is not very efficient (i.e., only 88.3%) as it overestimates the proportion of the population with storage and understates γ that represents the amount of green energy available to consumers with storage. When the supplier adopts the adaptive mechanism using a higher learning rate (i.e., $\beta_S = 0.05$), this leads to a faster, though less smooth, convergence as opposed to a slower, but smoother convergence to the lower learning rate (i.e., $\beta_S = 0.005$). In general, this demonstrates that the emergent behavior of individually utility-maximizing agents is indeed efficient such that the system maximizes its use of green energy, avoiding having to acquire electricity from the grid as the consumers change their demand profiles to follow green supply. Also, an efficient outcome means that the consumers have lower carbon emissions overall since they maximize their use of renewable energy (and thus, minimize their use of grid energy). Because the objectives of reducing carbon emissions and minimizing costs are aligned through the carbon-based pricing scheme, the selfish utility-maximizing behavior of a consumer allows it to be as green as it could be. Next, we analyse the system at an equilibrium state (i.e., once the system has converged after a number of days), for different proportions of the population that have storage capability.

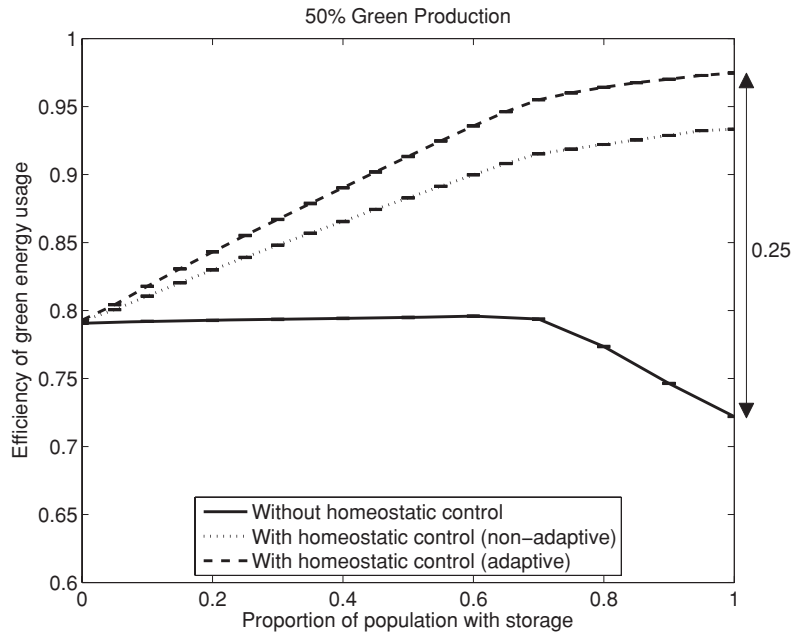
6.3. Effect of the Population Size with Storage Capability

Based on the previous experiment, here we set the learning rate at $\beta_S = 0.005$ (to ensure convergence) for our homeostatic control mechanism and vary the proportion of the population with storage capability from 0 to 1. Figure 7(a) shows the green energy usage efficiency (as a percentage of the generation capacity) for different storage penetration in the population (at the end of 100 simulation days). First, we observe that when too many consumers adopt storage, the system with no homeostatic control wastes more green energy (when the proportion is greater than 0.6 in particular). This is because too many consumers are blindly optimizing their demands based only on market prices, unaware of the green production, with efficiency decreasing to 72%.

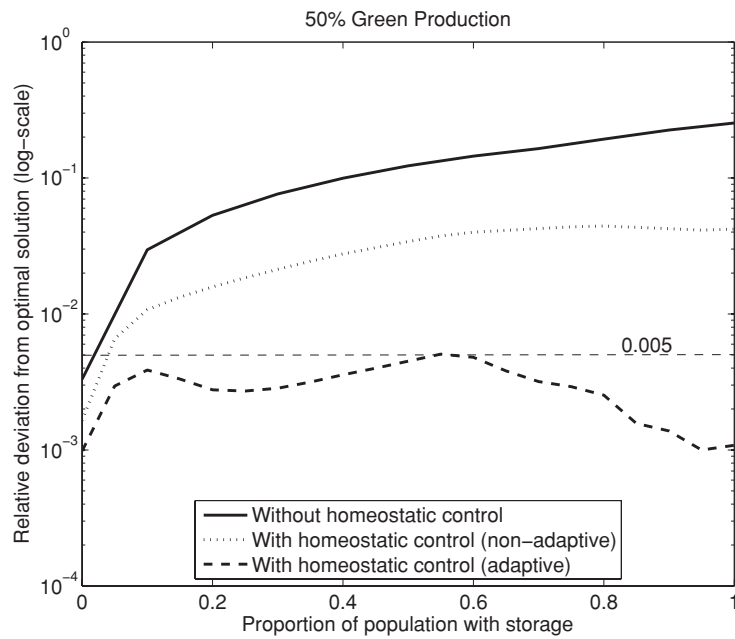
Now, when given a signal about green production (as in the homeostatic control mechanisms), the consumers are significantly greener as they are optimizing their demand based on how much and when green energy is available to them. With no storage in the system (i.e., no storage capability), green energy usage efficiency is the same with and without homeostatic control as a signal is then redundant. As storage penetration increases, the homeostatic control system becomes greener (having to use less electricity from the grid and more of the supplier's green energy). When storage penetration is at its highest, we can see how the effect of homeostatic control peaks with up to a 25% increase in green energy usage efficiency (i.e., a 25% decrease in carbon emissions) compared to without homeostatic control. In addition, Figure 7(b) (log-scale) also compares the system, at equilibrium, with and without homeostatic control against the optimal solution (i.e., we are comparing the relative error between the optimal solution and where the system converges to). We can see that our adaptive mechanism allows the system to converge to within 0.5% of the optimal solution as opposed to up to 4.4% without an adaptive mechanism and up to 25% without homeostatic control. Next, we investigate the effect of different penetrations of storage in the population on the cost of electricity.

6.4. The Cost of Electricity

Here we use our carbon-based pricing scheme to evaluate the cost of electricity given that different proportions of the population are able to utilize the full capacity of the green supplier. We consider cases where the cost of storage is added to the overall costs of the consumer or not. Thus, if cost of storage is fully subsidized, we observe from



(a)



(b)

Fig. 7. Green energy usage efficiency with and without homeostatic control for different proportions of the population with storage. The relative deviation from the optimal solution (as computed in Eq. (2)) is given in (b).

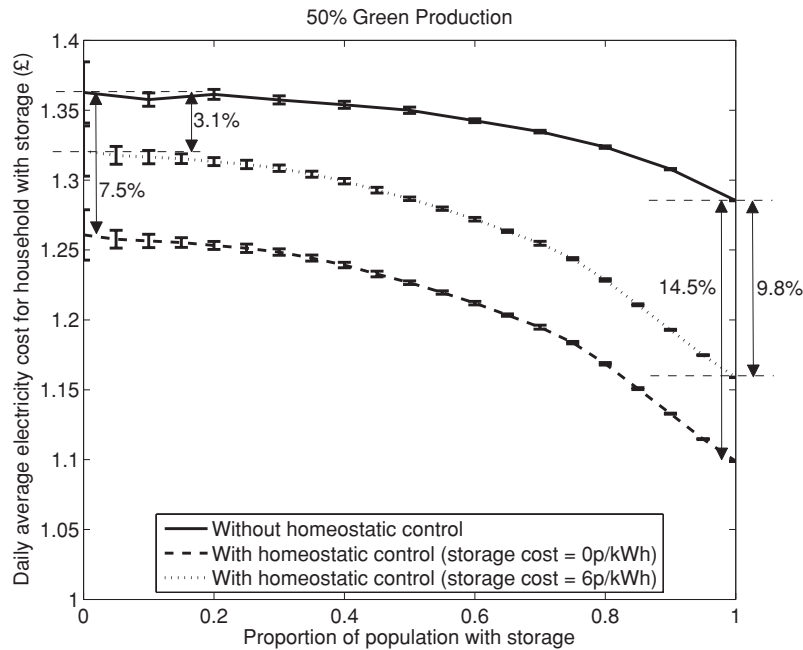


Fig. 8. Average cost of electricity for storage-capable consumers with and without homeostatic control for different proportions of the population with storage.

Figure 8 the decreasing trend which results from more consumers being able to change their demand to use green and less expensive energy. Furthermore, the cost without homeostatic control is marginally higher than with homeostatic control. The individual savings of a consumer with storage are maximized when the demand response of the population is maximized. At this point, it results in up to 14.5% lower costs because it can optimize its demand for greener and less expensively priced electricity. Conversely, when there is no storage in the system, a single consumer deciding whether to adopt storage would be able to make 7.5% savings even if she were the only one to have storage in the system. Indeed, this points to a significant incentive for consumers to adopt storage. Moreover, because the payoff of owning storage (rather than not) is always higher, it is likely that all consumers will eventually acquire storage devices and capitalize on the supplier's control signal of green energy.²⁸

Next, we assume that there is a cost to storage. Specifically, we assume a storage device typically costing £200 per kWh and a startup cost of £200 (i.e., installation cost, wiring, and voltage inverter), with a lifetime of 10 years. When the costs are spread over the lifetime and with a capacity of 4kWh, this equals to a cost of 6p/kWh per day. From Figure 8, we can observe an increase in the cost of storage with homeostatic control (as we now factor in the cost of storage). However, a system with homeostatic control still remains less costly than without, changing from 3.1% less expensive (when the

²⁸Note that the setting we analyse here is different from Vytelingum et al.'s [2010] where they describe how eventually the system converges to an equilibrium where only 38% require storage. In our context, we are looking at a single supplier and its consumers (with changes in total demand profile small enough to have no effect on wholesale market prices), while they consider the whole grid where changes in market demand affect wholesale market prices.

proportion of the population with storage is low) to 9.8% (when the whole population has storage).

6.5. Effect of Green Energy Prediction Error


Given that the supplier's generation capacity is dependent on weather conditions and hence uncertain, we investigate the sensitivity of our system with respect to prediction errors in green energy production (since the signal is a prediction of green energy to be produced in a near future, typically 3–4 hours). We do so to evaluate the sensitivity of our homeostatic control mechanism to prediction errors and do so by analysing the efficiency of our system in the presence of prediction errors.

From Figures 9(a) and 9(b), we can see how the system performs if the supplier overstates or understates its green production. In particular, we can observe that efficiency increases slightly (by up to 0.5%) when the supplier overstates γ . This is because storage-capable consumers are recommended to use more green energy than available ensuring that all the green energy will be used. Thus the green energy-focused supplier (i.e., motivated to use all the green energy it produces to supply its consumers) is more inclined to slightly overstate γ . However, in so doing, the total demand will be higher than the green energy available at peak times (when consumers are most incentivized to use green energy because of the greater cost of grid electricity) such that the supplier will need to buy from the grid to cover the difference, leading to higher costs. When the supplier slightly overstates γ , we can clearly see that cost increases exponentially (see Figure 9(b)) while efficiency only increases by up to a negligible 0.5% (see Figure 9(a)).


On the other hand, when the supplier understates γ , the cost increases while efficiency decreases, such that the system is worse off both in terms of cost and efficiency. Thus, crucially to minimize cost (and trade off only up to 0.5%), the supplier is incentivized to predict accurately since poor prediction (whether understating or overstating) invariably results in rapidly increasing costs for its consumers. However, when the absolute prediction error is small, the increase in costs is also small such that the system is relatively nonsensitive to small prediction errors. Specifically, when the supplier understates γ by up to 10%, the cost only increases by up to 0.5% and, by up to 0.08% when it overstates γ by up to 10%. This suggests that our system is relatively robust against small errors in prediction.

7. CONCLUSIONS

In this article we have presented a novel decentralized homeostatic control mechanism for a green supplier and its consumers connected through the grid. Our approach is underpinned by the use of autonomous software agents endowed with algorithms to optimize their owners' energy consumption. We fully utilize the real-time carbon intensity signal provided by existing information and communication technologies coupled with a flexible carbon-based pricing in order to implement real-time pricing in the smart grid for any type of supplier (green or brown). Given this, we have shown how agents can optimize their consumption of energy by storing electricity at times when it is inexpensive, thereby maximizing their individual utility and minimizing their carbon footprint. Building upon the feedback from the population of agents, the homeostatic control mechanism implements a control signal that the supplier learns to adapt to maximize its green energy usage by incentivizing the portion of the population with storage capability to optimize their storage of electricity in such a way that their individual selfish behavior is aligned with the greater good of the system. We demonstrate that such an adaptive mechanism is essential for an efficient system,



(a)



(b)

Fig. 9. Effect of inaccurate green generation prediction. With no homeostatic control, cost is 10% more than the optimal cost throughout, as is the efficiency of green usage at 79.5%.

and indeed, we empirically demonstrate that it allows the system to converge to an optimal solution. Thus, we are able to show that our mechanism can achieve up to 25% greater efficiency for the supplier and up to 14.5% savings for the consumer over other single-signal-based mechanisms. We also demonstrate that the system is not sensitive to small prediction errors and that a supplier is incentivized to predict accurately its green production to minimize costs.

In general, our work demonstrates the strengths of closer synergies between power systems and information and communication technologies by showing how real-time consumption and carbon emissions data can be used to make electricity distribution more efficient. Moreover, our approach departs from the traditional approaches that did not consider the self-interested behaviors of consumers with different preferences, constraints (on storage capacity), and consumption profiles. Hence, we believe this is the first realistic step in building deployable multiagent systems in the grid. However, this article also highlights the challenges that lie in building robust communication technologies and autonomous agents for home energy management. In particular, while robust communications are needed to guarantee that signals are provided in a timely fashion to all actors in the system (i.e., γ for the agents and real-time consumption data for the supplier), intelligent agents need to be designed to learn the patterns of behavior of their owners either by analysing their usage profile or interacting with them through new user interfaces such as those of AlertMe or Google PowerMeter.

In future we aim to investigate the properties of our control mechanism within a more theoretical framework in order to prove its convergence properties (under specific settings) and expand our optimization model to consider other types of constraints (e.g., bounds on energy storage to extend battery life or decay in battery storage capacity). Moreover, to improve the practical performance of our mechanism we aim develop better prediction algorithms. While in our simulations we assumed that the renewable power generation does not vary significantly across days, this may not always be the case. Hence, we will investigate the robustness of our system given different volatilities for renewable power generation. Indeed, it is important to note that our approach is highly dependent on the predictions that both suppliers and consumers can make about the generation and consumption of energy, respectively. While we have studied how prediction errors on the supplier side affect the system, we have simplified the model of the consumers to avoid significant prediction errors. The main challenge here is that consumer behavior tends to vary significantly between days, across seasons, and on specific days (e.g., sporting events or holidays) and predicting the daily consumption is a challenging task that is likely to require advanced machine learning solutions and human-agent interaction mechanisms that will allow agents to understand the users' future consumption intentions. Indeed, a major assumption of our approach is that users will be happy to delegate the task of shifting energy to their agents in the hope to get the best deal. While this may be fine for some people, others may not trust their agents to act on their behalf in such a sensitive domain. Given this, work is needed to develop good feedback interfaces and interaction mechanisms for humans to trust their agents.

Another important future direction involves extending the model to consider multiple green suppliers operating in the grid such that they constitute the majority of the energy producers. This, we believe, will require more advanced control techniques given that carbon intensity signals may disappear altogether and a key objective of suppliers will change from maximizing green energy usage to ensuring that demand always follows supply throughout the grid in order to avoid brownouts and blackouts.

APPENDIX

A. NOMENCLATURE

Half-hourly time-slots	$T = \{1, 2, \dots, 48\}$
Agent	$a \in \mathcal{A}$
Agent's load profile	$\ell_i^a \forall i \in T$
Agent's storage profile	$b_i^a \forall i \in T$
Agent's demand profile	$l_i^a \forall i \in T$
Agent's demand profile priced as green energy	$l_i^{a,green} \forall i \in T$
Agent's demand profile priced as grid	$\tilde{l}_i^{a,green} \forall i \in T$
Agent's storage charging profile	$b_i^{a+} \forall i \in T$
Agent's storage discharging profile	$b_i^{a-} \forall i \in T$
Agent's storage charging capacity	b_+^a
Agent's storage discharging capacity	b_-^a
Agent's storage efficiency	e^a
Agent's storage capacity	c^a
Aggregate demand profile	$l_i^A \forall i \in T$
Aggregate demand profile from green production	$l_i^{A,green} \forall i \in T$
Aggregate demand profile from grid	$\tilde{l}_i^{A,grid} \forall i \in T$
Grid price bands	$p_n^{grid} \in \mathbb{R}^+$ where $n \in \{1, 2, 3, 4, 5\}$
Carbon intensity	$\rho_i \in \mathbb{R}^+ \forall i \in T$
Green production price	$p^{green} \in \mathbb{R}^+$
Grid price	$p_i^{grid} \in \mathbb{R}^+ \forall i \in T$
Green supplier production	$q_i \in \mathbb{R}^+ \forall i \in T$
Supplier's control signal	$\gamma_i \in \mathbb{R}^+ \forall i \in T$
Learning rate of supplier's adaptive mechanism	$\beta_S \in [0, 1]$

ACKNOWLEDGMENTS

We wish to thank the anonymous reviewers for their valuable comments and our industrial partners for their input and direction.

REFERENCES

- ALPAYDIN, E. 2004. *Introduction to Machine Learning*. The MIT Press.
- BATHURST, G. N. AND STRBAC, G. 2003. Value of combining energy storage and wind in short-term energy and balancing markets. *Electric Power Syst. Res.* 67, 1, 1–8.
- BURGIO, A., MENNITI, D., PINNARELLI, A., AND SORRENTINO, N. 2009. Reliability studies of a PV-WG hybrid system in presence of multi-micro storage systems. In *Proceedings of PowerTech*. 1–5.
- BUSSMANN, S., JENNINGS, N. R., AND WOOLDRIDGE, M. J. 2004. *Multi-Agent Systems for Manufacturing Control: A Design Methodology*. Springer-Verlag, Berlin.
- DECC. 2009a. Smarter grids: The opportunity. Tech. rep., Department of Energy and Climate Change (DECC), UK Government.
- DECC. 2009b. The UK low carbon transition plan: National strategy for climate and energy. Tech. rep., Department of Energy and Climate Change (DECC), UK Government.

- FARUQUI, A. AND GEORGE, S. 2005. Quantifying customer response to dynamic pricing. *Electricity J.* 18, 4, 53–63.
- GAJBHIYE, R. AND SOMAN, S. A. 2010. Facilitating emission trade within power exchange: Development of conceptual platform. In *Proceedings of the 7th International Conference on the European Energy Market*. 1–6.
- GERDING, E. H., ROBU, V., STEIN, S., PARKES, D., ROGERS, A., AND JENNINGS, N. R. 2011. Online mechanism design for electric vehicle charging. In *Proceedings of the 10th International Conference on Autonomous Agents And Multi-Agent Systems*.
- GOMES, C. 2009. Computational sustainability: Computational methods for a sustainable environment, economy, and society. *Bridge* 39, 4, 5–13.
- HAMMERSTROM, D., AMBROSIO, R., CARLON, T., DESTESE, J., HORST, G., KAJFASZ, R., KIESLING, L., MICHIE, P., PRATT, R., YAO, M., ET AL. 2008. Pacific northwest gridwise testbed demonstration projects; Part I. Olympic Peninsula Project. Tech. rep., PNNL-17167, Pacific Northwest National Laboratory (PNNL), Richland, WA.
- HARRIS, C. 2005. *Electricity Markets: Pricing, Structures, and Economics*. Wiley.
- HOBBS, B. F., BUSHNELL, J., AND WOLAK, F. A. 2010. Is load-based carbon trading less costly than source-based trading? In *Proceedings of the 7th International Conference on the European Energy Market*. 1–6.
- INFIELD, D., SHORT, J., HOME, C., AND FRERIS, L. 2007. Potential for domestic dynamic demand-side management in the UK. In *Proceedings of the Power Engineering Society General Meeting*. 1–6.
- JENNINGS, N. R., CORERA, J., LARESGOITI, I., MAMDANI, E. H., PERRIOLAT, F., SKAREK, P., AND VARGA, L. Z. 1996. Using ARCHON to develop real-world DAI applications for electricity transportation management and particle accelerator control. *IEEE Expert Syst.* 11, 6, 60–88.
- KIRSCHEN, D. AND STRBAC, G. 2004. *Fundamentals of Power System Economics*. Wiley.
- KOCKAR, I., CONEJO, A. J., AND McDONALD, J. R. 2009. Influence of emissions trading scheme on market clearing and prices. In *Proceedings of the Power Energy Society General Meeting*. 1–5.
- KOK, K. AND VENEKAMP, G. 2010. Market-based control in decentralized power systems. In *Proceedings of the 1st International Workshop on Agent Technology for Energy Systems*.
- KORPAAS, M., HOLEN, A. T., AND HILDRUM, R. 2003. Operation and sizing of energy storage for wind power plants in a market system. *Int. J. Electrical Power Energy Syst.* 25, 8, 599–606.
- LI, H. AND TESHATSION, L. 2009. Development of open source software for power market research: The AMES test bed. *J. Energy Markets* 2, 2, 111–128.
- LI, J., POULTON, G., AND JAMES, G. 2010. Coordination of distributed energy resource agents. *J. Appl. Artif. Intell.* 24, 5, 351–380.
- LUBOSNY, Z. AND BIALEK, J. W. 2007. Supervisory control of a wind farm. *IEEE Trans. Power Sys.* 22, 3, 985–994.
- MAKAROV, Y., PENGWEI, D., KINTNER-MEYER, M. C. W., CHUNLIAN, J., AND ILLIAN, H. 2010. Optimal size of energy storage to accommodate high penetration of renewable resources in WECC system. In *Proceedings of the Conference on Innovative Smart Grid Technologies*. 1–5.
- MARIEB, E. AND HOEHN, K. 2007. *Human Anatomy and Physiology*. Pearson Education.
- MAZER, A. 2007. *Electric Power Planning for Regulated and Deregulated Markets*. Wiley-IEEE.
- MILLIGAN, M., PORTER, K., DEMEO, E., DENHOLM, P., HOLTINEN, H., KIRBY, B., MILLER, N., MILLS, A., O'MALLEY, M., SCHUERGER, M., AND SODER, L. 2009. Wind power myths debunked. *IEEE Power Energy Mag.* 7, 6, 89–99.
- MORALES, J., CONEJO, A. J., AND PEREZ-RUIZ, J. 2010. Short-Term trading for a wind power producer. *IEEE Trans. Power Syst.* 25, 1, 554–564.
- RAMCHURN, S., VYTELINGUM, P., ROGERS, A., AND JENNINGS, N. R. 2011. Agent-Based control for decentralised demand side management in the smart grid. In *Proceedings of the 10th International Conference on Autonomous Agents And Multi-Agent Systems*.
- SCHWEPPE, F., DARYANIAN, B., AND TABORS, R. 1989. Algorithms for a spot price responding residential load controller. *IEEE Power Engin. Rev.* 9, 5, 49–50.
- SCHWEPPE, F., TABORS, R., KIRTLEY, J., OUTHRED, H., PICKEL, F., AND COX, A. 1980. Homeostatic utility control. *IEEE Trans. Power Apparatus Syst.* 99, 3, 1151–1163.
- SMITH, K. 2010. Energy demand research project - Fifth progress report. Tech. rep., OFGEM, UK.
- SUN, J. AND TESHATSION, L. 2007. DC optimal power flow formulation and solution using QuadProgJ. In ISU Economics Working Paper no. 06014.
- US DEPARTMENT OF ENERGY. 2003. Grid 2030: A national vision For electricity's second 100 years. <http://www.oe.energy.gov/DocumentsandMedia/Electric.Vision.Document.pdf>

- VYTELINGUM, P., VOICE, T. D., RAMCHURN, S., ROGERS, A., AND JENNINGS, N. R. 2010. Agent-Based micro-storage management for the smart grid. In *Proceedings of the 9th International Conference on Autonomous Agents and Multi-Agent Systems*. 39–46.
- WONG, S., BHATTACHARYA, K., AND FULLER, J. 2010. Long-Term effects of feed-in tariffs and carbon taxes on distribution systems. *IEEE Trans. Power Syst.* 25, 3, 1241–1253.
- YGGE, F., AKKERMANS, J. M., ANDERSSON, A., KREJIC, M., AND BOERTJES, E. 1999. The HOMEBOTS system and field test: A multi-commodity market for predictive power load management. In *Proceedings of the 4th International Conference on the Practical Application of Intelligent Agents and Multi-Agent Technology*. vol. 1. 363–382.

Received October 2010; revised December 2010; accepted February 2011